Environ Resource Econ (2013) 56:73–84
DOI 10.1007/s10640-013-9645-z

Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy?

Frank Ackerman · Elizabeth A. Stanton ·
Ramón Bueno

Accepted: 25 February 2013 / Published online: 16 March 2013

Abstract Climate change involves uncertain probabilities of catastrophic risks, and very longterm consequences of current actions. Climate economics, therefore, is centrally concerned with the treatment of risk and time. Yet conventional assumptions about utility and optimal economic growth create a perverse connection between risk aversion and time preference, such that more aversion to current risks implies less concern for future outcomes, and vice versa. The same conflation of risk aversion and time preference leads to the equity premium puzzle in finance. A promising response to the equity premium puzzle, the recursive utility of Epstein and Zin, allows separation of risk aversion and time preference—at the cost of considerable analytic complexity. We introduce an accessible implementation of Epstein–Zin utility into the DICE model of climate economics, creating a hybrid "EZ-DICE" model. Using Epstein–Zin parameters from the finance literature and climate uncertainty parameters from the science literature, we find that the optimal climate policy in EZ-DICE calls for rapid abatement of carbon emissions; it is similar to standard DICE results with the discount rate set to equal the risk-free rate of return. EZ-DICE solutions are sensitive to the intertemporal elasticity of substitution, but remarkably insensitive to risk aversion. Insensitivity to risk aversion may reflect the difficulty of modeling catastrophic risks within DICE. Implicit in DICE are strong assumptions about the cost of climate stabilization and the certainty and speed of success; under these assumptions, risk aversion would in fact be unimportant. A more realistic analysis will require a subtler treatment of catastrophic climate risk.

Keywords Climate economics · Risk aversion · Epstein–Zin utility · DICE · Integrated assessment modeling · Climate sensitivity · Catastrophic risk

---

F. Ackerman (✉)· E. A. Stanton · R. Bueno
Synapse Energy Economics, Cambridge, MA, USA
e-mail: frankackerman12@gmail.com

E. A. Stanton
e-mail: stantonea88@gmail.com

R. Bueno
Stockholm Environment Institute-US Center, Somerville, MA, USA
e-mail: Ramon.Bueno@sei-us.org

1 The Need for a New Utility Function

Economic analysis of climate change is often framed in cost-benefit terms. On the one hand, spending money now on reducing emissions means that less is available for current consumption and ordinary investment. On the other hand, not spending money now on reducing emissions means that less will be available for future consumption and investment as climate damages increase. An optimal economic growth path can be calculated, subject to these two constraints; among other results, it includes an optimal climate policy, or pace of emission reduction.

This much is standard economic methodology, equally applicable to many different problems. Climate change, however, is unique both in the magnitude of the uncertain risks it poses, and the long time spans, often exceeding a century, between physical causes and effects. The attempt to stretch standard methods to cover such extremes of risk and time reveals a hole in the theoretical fabric.

The growth model commonly used in climate economics (and elsewhere) can be traced back to Ramsey (1928), who demonstrated that along an optimal growth path, the discount rate for consumption equals the productivity of capital. Later analysis (Cass 1965; Koopmans 1965) formalized this conclusion in what is now known as the "Ramsey equation":

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

Here r is the discount rate for consumption, δ is the rate of pure time preference (or equivalently, the discount rate for utility), η is the elasticity of the marginal utility of consumption (or equivalently, the inverse of the intertemporal elasticity of substitution), and g is the rate of growth of per capita consumption, which is assumed to be known with certainty.[1] Informally, growth means that future generations will be richer and the marginal utility of increasing their consumption will be lower, so their consumption is discounted relative to ours today. The larger the value of η, the more we are focused on our own, relatively lower-income circumstances, rather than on our wealthier descendants.

The model is developed in an expected-utility framework. Although many utility functions could be used, it is mathematically convenient to adopt the constant relative risk aversion (CRRA) function. Under this assumption, the utility obtained from current consumption, c, is[2]

$$u\,(c) = \frac{c^{1-\eta} - 1}{1-\eta} \tag{2}$$

Here η, the same parameter as in (1), is the coefficient of relative risk aversion; it determines the curvature of the utility function. The larger the value of η, the more we are focused on avoiding risks of losses rather than gambling for uncertain gains.

While these two interpretations of η might appear to parallel each other—larger η means greater aversion to lower consumption in each case—they are actually on a collision course. Expressing more aversion to climate risks, via larger η in (2), implies a higher discount rate and greater disinterest in the future in (1). Economists are thus apparently condemned to choose between a low discount rate, reflecting heightened concern for future generations but a very low level of current risk aversion (as in Cline 1992; Stern 2006), or a high discount rate, reflecting more aversion to current risks at the expense of greater indifference to future

---

[1] If g is normally distributed with variance σ^2, and the CRRA utility function (2) is assumed to apply, then (1) becomes $\mathrm{r} = \delta + \eta\, \mathrm{g} - 1/2\, \eta^2\, \sigma^2$ (see, e.g., Ackerman et al. 2009).

[2] When $\eta = 1$, (2) is replaced by $u(c) = \ln c$—which is the limit of (2) as η approaches 1.

generations (as in Nordhaus 2008 and many others). There is no natural way to model a combination of strong risk aversion and strong future orientation.[3] A similar problem can be seen in survey research, which finds that the same people express attitudes toward risk and toward time preference that imply very different values of η (Atkinson et al. 2009).

Indeed, the expected-utility framework in general may be an obstacle to a realistic analysis of risk aversion. Within such a framework, the assumptions of plausible levels of risk aversion toward large risks and nonzero aversion toward small risks have been shown to be incompatible (Rabin 2000).

2 Epstein–Zin Utility

Similar problems arise in other areas of economics. The same expected-utility framework, including the same conflation of risk aversion and time preference, is unable to explain why rates of return are so high on equity and so low on risk-free assets (Mehra and Prescott 1985). More than 25 years of discussion of this "equity premium puzzle" has not yet produced a consensus on the solution, but has given rise to a number of possible explanations (DeLong and Magin 2009). Many of them involve new theories of investor behavior, challenging or expanding the traditional understanding of economic rationality.

One of the most promising responses to the equity premium puzzle is the recursive utility model of Epstein and Zin (1989, 1991). They propose a utility function with an additional degree of freedom, allowing separate calibration of risk aversion and time preference. Epstein–Zin utility at time t, or U_t, depends both on current consumption, c_t, and on the certainty-equivalent at time t of future utility, $\mu_t(U_{t+1})$,:

$$U_t = \lceil (1-\beta)\, c_t^{\rho} + \beta \left(\mu_t \left[U_{t+1}\right]\right)^{\rho} \rceil^{1/\rho} \tag{3}$$

Here β is the discount factor for utility;[4] the rate of pure time preference is $\delta = (1-\beta)/\beta$. Time preference is also affected by ρ; the intertemporal elasticity of substitution (IES) is $\psi = 1/(1-\rho)$. The certainty-equivalent of future utility is specified as

$$\mu_t\left(U_{t+1}\right) = \left(E_t\left[U_{t+1}^{\alpha}\right]\right)^{1/\alpha} \tag{4}$$

Here E_t is the expected value operator at time t. Risk aversion is measured by α; the coefficient of relative risk aversion is $\gamma = 1 - \alpha$. In the special case where $\rho = \alpha$, (3) becomes equivalent to the present value (over the unlimited future) of (2), with $\eta = \gamma = 1/\psi$. However, this special case appears unrealistic, since both risk aversion (γ) and the IES (ψ) must be greater than one to fit observed financial market patterns (Bansal and Yaron 2004); matching such patterns is possible with (3) but not with (2).

There have been a few applications of Epstein–Zin utility to climate economics; the first may be Ha-Duong and Treich (2004), using a simplified four-period model. Traeger and several colleagues are exploring the theoretical implications of Epstein–Zin utility for climate analysis, with applications to a model derived from the well-known DICE model (Crost and Traeger 2010; Jensen and Traeger 2011).[5] In DICE and similar models, the constraints of the

---

[3] It is theoretically possible to simultaneously model strong risk aversion with a large η, and strong future orientation with a negative δ—but it is difficult to develop a plausible argument for $\delta < 0$.

[4] The symbols α, β, and ρ are introduced solely in order to make (3) and (4) as transparent as possible; as explained in the text, they are simple transformations of, respectively, γ (coefficient of relative risk aversion), δ (rate of pure time preference), and ψ (intertemporal elasticity of substitution).

[5] DICE, developed by Nordhaus, is described in Nordhaus (2008).

expected-utility framework lead to the paradoxical result that a higher coefficient of relative risk aversion can lead to reduced abatement efforts; Kaufman (2012), using a simplified DICE-like model, shows that Epstein–Zin utility eliminates this paradox.

The existing studies that apply Epstein–Zin utility to climate economics demonstrate the importance of this approach. They show that it leads to optimal solutions and policy recommendations that are significantly different from the conventional approach based on the Ramsey equation and the expected-utility framework. The paradoxes that result from the traditional entanglement of risk aversion and time preference are neatly resolved.

This accomplishment, however, comes at a great analytical cost. The applications of Epstein–Zin utility to date have required a level of mathematical sophistication and complexity well beyond the norm for the integrated assessment literature, yet have analyzed only very simplified climate models—often using ad hoc models based on or resembling DICE, but with substantial simplifications of its structure (Crost and Traeger 2010; Kaufman 2012). Even with the best mathematical techniques, there are inherent obstacles to direct application of Epstein–Zin utility to a model such as DICE.

The problem results from the recursive nature of Epstein–Zin utility. The separation of risk aversion and time preference in Eq. (3) makes current utility depend not only on current consumption, but also on expectations about the next period's utility, which in turn depends on the following period, and so on into the indefinite future. Thus utility is defined only on the complete branching tree of possible futures growing out of the present moment. The DICE model analyzes futures over 60 time periods (decades); for use in this article we have truncated it at 40 periods. If one binary choice is made in each of 40 time periods, the tree of possible futures has 2^{40}, or roughly a trillion branches. To calculate Epstein–Zin utility in the first period, it would be necessary to follow every one of those trillion branches to its endpoint.

One response to this problem is to model very few time periods, thereby losing the long-term modeling of climate and economic dynamics found in integrated assessment models (Ha-Duong and Treich 2004; Kaufman 2012). Another approach is taken by Crost and Traeger (2010), who analyze a model based on DICE over an infinite time horizon, using the Bellman equation to reduce the recursive dynamic programming problem to a single-period problem. This is only a limited success, however: in order to obtain an approximate numerical solution to the Bellman equation, it was necessary to replace the DICE treatment of the carbon cycle and climate dynamics with a much simpler and physically less realistic alternative (Crost and Traeger 2010, p. 6).

3 The EZ-DICE Model

Our goal is to demonstrate the applicability of Epstein–Zin utility to a model at the level of complexity of DICE—which, it should be remembered, is among the simplest of the integrated assessment models. The existing Epstein–Zin climate analyses have retained the full theoretical treatment of uncertainty but applied it to a stylized simplification of a climate model; we adopt the opposite strategy, retaining the full structure of DICE but combining it with a stylized simplification of uncertainty. Specifically, we achieve a drastic pruning of the tree of future consequences by modeling only one form of uncertainty, regarding the climate sensitivity parameter.

Climate sensitivity is the long-term global average temperature increase caused by a doubling of the atmospheric concentration of carbon dioxide. It is a crucial determinant of the expected pace of climate change. Although it remains uncertain, perhaps inescapably so, there is an emerging near-consensus on how to model its probability distribution (Roe and Baker 2007). The true value of climate sensitivity will, of course, eventually become known or knowable in retrospect, when it may well be too late to do anything about climate change.

We introduce the following form of uncertainty into the DICE model: climate sensitivity, at the outset, adopts one of five possible values. At first, only the probabilities of these five values are known. The actual value becomes known at a specific date in the future. Our interest is in the optimal decisions made under uncertainty, in the years before the true value of climate sensitivity is discovered.

The date for the resolution of uncertainty should, on general principles, be in the later years of this century. If it is much less than 50 years into the future, then the prior period of decision-making under uncertainty becomes too short for useful analysis. If, on the other hand, it is more than a century from now, then climate outcomes are all but determined prior to the resolution of uncertainty—and, with a century of additional data, it seems unlikely that climate sensitivity would remain completely uncertain.[6]

For the numerical results presented here, we chose the midpoint of the second half of the century—2075, or 70 years after the model's base year—for the resolution of uncertainty. Experiments with the model (not presented here) showed that other dates in the second half of this century would yield qualitatively similar results. An additional result of the specific choice of 2075 is explained below.

To implement this picture of uncertainty, we run the DICE model simultaneously for five possible states of nature, differing only in climate sensitivity. The five states are constrained to make identical investment and abatement choices for the first 70 years, since it is not yet known which state prevails; in all other respects, the five states are independent of each other. The solutions to DICE in all five states, both before and after the discovery of true climate sensitivity in 2075, are chosen to maximize first-period Epstein–Zin utility:

$$U_1 = \left\lceil (1-\beta)\, c_1^{\rho} + \beta\, (\mu_1 [U_2])^{\rho} \right\rceil^{1/\rho} \tag{5}$$

Here the first-period certainty-equivalent value of next-period utility, $\mu_1(U_2)$ is calculated across the five states of nature, according to (4).

In this model of uncertainty, the tree of possible futures branches only once, in the first period; after that point there are five separate, unbranching trunks (although they are tied together for the first 70 years). That is, the future is deterministic once climate sensitivity has been picked, even though no one knows, for quite a while, which of the possible futures is occurring. In the second period and thereafter, each state of nature has a single value of climate sensitivity, hence no risk, no role for risk aversion, and no need for calculation of certainty-equivalent values. Second-period utility in each state reduces to

$$U_2 = \left\lceil (1-\beta)\, c_2^{\rho} + \beta U_3^{\rho} \right\rceil^{1/\rho} \tag{6}$$

This is equivalent to the present value, over all remaining time periods, of the much simpler utility function (2), with $\rho = 1 - \eta$.

So in our model, the full Epstein–Zin calculation, with its risk-averse evaluation of five possible futures, is needed only once, in the initial period. This simplifies Epstein–Zin utility to an extent that makes it feasible to introduce it into DICE. The modified model, EZ-DICE, chooses the single investment plan and climate policy that applies to all states until 2075,

---

[6] If precise, noise-free measurements were possible—as the simple, deterministic treatment of climate variables in DICE may suggest—then it would be possible to calculate the value of climate sensitivity much sooner. In reality, precision is difficult to achieve: both year-to-year climate variability and the positive-feedback nature of the relevant climate dynamics lead to great uncertainty in estimates of climate sensitivity (Roe and Baker 2007).

Table 1 Climate sensitivity parameters for five states of nature

State label	S_1	S_2	S_3	S_4	S_5
Length of interval (%)	50	40	5	3	2
Climate sensitivity at midpoint	2.43	3.67	6.05	8.20	16.15

and the five separate investment plans and climate policies, one for each state, after 2075. These choices are made to maximize Epstein–Zin utility across all five states in the initial period, i.e. U_1 as defined in (5)—which, thanks to recursion, includes the evaluation of all future consequences in all five states.

4 Model Calibration

For climate sensitivity, we use the Roe and Baker (2007) distribution, calibrated to match the results of Murphy et al. (2004).[7] We partition the probability distribution into five unequal intervals, as shown in Table 1, using the climate sensitivity at the midpoint of each interval as the value for that state of nature, and the length of the interval as the probability of that state. For example, the 25th percentile climate sensitivity value, 2.43, is assumed to occur in state S_1 with 50 % probability; the 99th percentile value, 16.15, occurs in state S_5 with 2 % probability. In this distribution, the median value of climate sensitivity is 3.00, and the mean is 3.55. States S_3, S_4, and S_5, all of them above the 90th percentile, are chosen to allow examination of the role of dangerously high climate sensitivity.

For the Epstein–Zin parameters, there is not yet a consensus on the correct values for modeling financial markets. Two major studies, however, have estimated the IES at 1.5 and the coefficient of relative risk aversion at 9.5 to 10 (Bansal and Yaron 2004; Vissing-Jørgensen and Attanasio 2003). It is not clear that financial market parameters are applicable to the different risks encountered in a climate model, but there is no obvious source for more appropriate parameters. It may be of some interest, moreover, to learn what the financial market parameters would imply for optimal responses to climate risks. On this basis, we assume $\psi = 1.5$ and $\gamma = 10$.

There is one more parameter in the utility function (3), influencing the rate of time preference—the discount factor, β, based on the rate of pure time preference, δ. This can be calibrated to the risk-free rate of return: in the absence of risk, (1) should apply, with r representing the risk-free rate and $\eta = 1/\psi$. The other element of (1). is the growth rate of per capita consumption, g. In DICE scenarios, g typically declines over time; in the default scenario, g averages 1.3 % per year for the first 150 years—a reasonable value, and the same as the average value of g in the Stern Review modeling (Stern 2006). Here we are ignoring the variance in g, which might be significant in reality, and would imply a lower discount rate (see footnote 1).

We assume a long-run average risk-free rate of return of 2.0 %. This is higher than some empirical estimates, which are often closer to 1.0 %.[8] Under our assumptions, (1) implies that $\delta = .0113$—well above the Stern Review rate of .001, although lower than the rates often assumed in DICE and other conventional models.

---

[7] See Roe and Baker (2007), Figure 4B.

[8] For example, the real rate of return on U.S. treasury bonds has averaged 1.1 % since World War II (DeLong and Magin 2009).

Table 2 Results using DICE utility function

	Description	δ	η	Climate sensitivity	SCC, 2015 (2010 \$ /tCO$_2$)	Abatement, 2075	Implicit risk-free rate of return ($\delta + \eta g$)
D1	DICE defaults	.0150	2.0	3.00	\$12.70	.333	.041
D2	Higher climate sensitivity	.0150	2.0	3.55	\$14.95	.368	.041
D3	Lower discount rate	.0031	1.3	3.00	\$48.03	.647	.020
D4	Lower discount rate, higher climate sensitivity	.0031	1.3	3.55	\$58.76	.721	.020
D5	5 states, expected value of DICE utility	.0150	2.0	(5 values)	\$13.95	.346	.041

5 Results

Our principal results consist of five runs using the DICE utility function (2), shown in Table 2, and five runs using the Epstein–Zin utility function (3), shown in Table 3. We compare model runs on the basis of two principal statistics: the social cost of carbon (SCC), or marginal damages per tonne of CO$_2$, for emissions in 2015; and the fraction of global emissions abated by 2075, the final year of decision-making under uncertainty. To facilitate comparison to the finance literature, we include the risk-free rate of return implied by the parameters in each run.[9]

Running DICE with its default values for all parameters and inputs (run D1, Table 2) yields a SCC of about \$13, with one-third of emissions abated by 2075, broadly consistent with other DICE analyses.[10] Switching from the median to the mean value of climate sensitivity (D2) causes a slight increase in both the SCC and the rate of abatement.

Lowering the discount rate, to match the 2 % risk-free rate assumed in our Epstein–Zin analysis, has a large effect on the results: it roughly doubles the rate of abatement and quadruples the SCC (compare D3 and D4 to D1 and D2, respectively).[11]

In a final calculation using DICE utility (D5), we use DICE default parameters, but adopt our five-state model of uncertainty about climate sensitivity. Maximizing the expected value of DICE utility, from Eq. (2), across the five states, yields results comparable to DICE defaults with median or mean climate sensitivity (i.e., D5 results fall between D1 and D2). Thus the introduction of the five-state model per se does not cause a major change in results.

In contrast, using the same five-state model of uncertainty but maximizing Epstein–Zin utility (run E1, Table 3) yields a SCC more than four times as high as the DICE default, with two-thirds of emissions abated by 2075. Consistent with other applications of Epstein–Zin utility to climate models, we find that the optimal policy involves much higher carbon prices and more rapid abatement when maximizing Epstein–Zin utility.

---

[9] The risk-free rate is calculated using Eq. (1), with $g = .013$; in Table 3, we assume $\eta = 1/\psi$. The calculation of the implied equity premium is more complex, and requires information beyond the scope of our results.

[10] Nordhaus (2008) uses DICE to find a SCC of \$30 per tonne of carbon emissions in "today's prices," equivalent to \$8 per tonne of CO$_2$; adjusted to the 2010 prices used in this article, this would be \$9 or \$10 per tonne of CO$_2$. The year of emissions for the Nordhaus calculation is not specified, but is likely earlier than 2015, which would tend to make the SCC lower than our value.

[11] The same reduction in the DICE discount rate can be achieved with many combinations of δ and η—with similar results.

Table 3 Results using Epstein–Zin utility function

	Description	δ	ψ	γ	SCC, 2015 (2010 \$/t CO_2)	batement, 2075	Implicit risk-free rate of return ($\delta + g/\psi$)
E1	Epstein–Zin defaults	.0113	1.5	10	\$57.49	.673	.020
E2	IES = 2.0	.0113	2.0	10	\$71.73	.734	.018
E3	Zero risk aversion	.0113	1.5	0	\$57.26	.658	.010
E4	High risk aversion	.0113	1.5	20	\$58.07	.690	.020
E5	Risk-free rate = 1 %	.0013	1.5	10	\$165.73	.995	.010

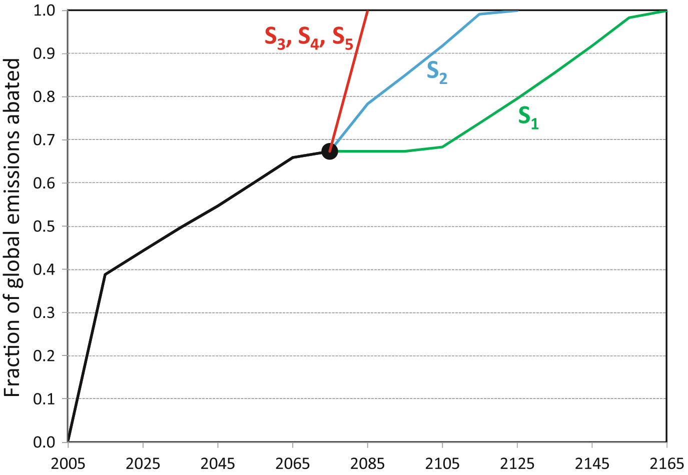

Fig. 1 Rate of abatement in five states of nature, run E1

The additional runs are sensitivity analyses, to explore the properties of this solution. Changes in the IES make a noticeable difference (E2). As expected, a larger IES implies that greater weight is given to future outcomes, raising the SCC and the pace of abatement.

On the other hand, fairly large changes in risk aversion have very small effects on the results (E3 and E4). Switching from $\gamma = 10$ to either 0 or 20 changes the SCC by 1 % or less, and changes the extent of abatement in 2075 by less than 2 percentage points. The effect is in the expected direction—greater risk aversion increases the SCC and the rate of abatement—but it appears much less significant than the effect of the IES, consistent with the findings of Crost and Traeger (2010). In fact, our initial numerical experiments with the model showed that the choice of 2075 for the resolution of uncertainty maximizes the (always small) effect of risk aversion on the outcomes.

6 Interpretation of Results

A possible explanation of the insensitivity to risk aversion is suggested by Fig. 1, showing the rate of abatement over time in the Epstein–Zin default run (E1). EZ-DICE calculates a single rate of abatement for all five states until 2075, when the true value of climate sensitivity is discovered; thereafter, it calculates optimal abatement paths separately for each state. After

Table 4 Temperature and climate damages in 2075, five states of nature, run E1

State	S_1	S_2	S_3	S_4	S_5
Climate sensitivity	2.43	3.67	6.05	8.20	16.15
Temperature increase in 2075 (°C above 1900)	1.6	2.1	2.6	2.9	3.4
Climate damages in 2075 (% of output)	0.7	1.2	1.9	2.3	3.1

the paths separate, state S_1, with the lowest climate sensitivity, takes 90 additional years to reach 100 % abatement. State S_2 needs 50 years to reach the same goal. States S_3, S_4, and S_5, however, all jump to 100 % abatement in a single decade.

DICE, both in its original form and in our modification, offers the option of any rate of abatement, up to 100 %, at any time. Over time, the unit costs of abatement decline while the damages from unabated climate change increase, making it increasingly attractive to abate. In EZ-DICE, the potential future damages from climate change are so great in the high climate-sensitivity states, S_3, S_4, and S_5, that the optimal solution is to jump to 100 % abatement as soon as the uncertainty is resolved. In effect, DICE will always avoid the worst extremes of possible future damages, by instead opting for immediate 100 % abatement.

The only difference, therefore, between the perceived damages in states S_3, S_4, and S_5, despite their very different climate sensitivity values, results from the modest impacts that occur before 2075. In run E1, the temperature increase in 2075, relative to the level in 1900, differs by less than 2 °C between S_1 and S_5, and less than 1 °C between S_3 and S_5 (Table 4). Using the DICE damage function, these temperatures imply climate losses in 2075 of only 0.7 % of world output in S_1, ranging up to 3.1 % in S_5. (For discussion of the DICE damage function and the importance of alternative damage estimates, see Ackerman and Stanton 2012.)

Even in S_5, with a climate sensitivity that is disastrously high (implying disastrously rapid warming) by many standards, DICE projects losses of just over 3 % of output by the end of the long period of uncertainty—comparable to a moderate-sized business cycle downturn (Table 4). The much greater potential damages after 2075 are avoided by the 100 % abatement option. Numerical experiments, not shown here, reveal that an immediate jump to 100 % abatement (from the point reached by E1 in 2075) occurs whenever climate sensitivity exceeds 5.25, which is the 89th percentile value in the probability distribution. Anything above 5.25 looks about the same to DICE; that is, the model is unable to distinguish differences within the upper 11 % of the probability distribution. The problems of tail risk, associated with extreme values from the right-hand tail of the distribution, cannot be analyzed in a model that does not "see" those values.

If risk aversion is almost immaterial in EZ-DICE, then similar results should be attainable in a simpler, risk-free environment. In the absence of risk, our Epstein–Zin calibration amounts to using the assumed risk-free rate of 2 % as the consumption discount rate. Confirming this interpretation, our Epstein–Zin results (E1, Table 3) are quite similar to the DICE results with a 2 % risk-free rate (D3 and D4, Table 2).

From this perspective, the principal contribution of Epstein–Zin utility in a model that truncates extreme risk, such as DICE, is its rationale for discounting at the risk-free rate, which is lower than discount rates adopted in many climate analyses (see arguments by Howarth 2003 along similar lines). If we were to assume a risk-free rate of 1 %, consistent with some readings of the historical record, then the results would be more extreme, as shown in run E5: worldwide emissions would be all but eliminated by 2075. This would make for a less interesting analysis of the Epstein–Zin methodology, since the response in the early decades is so strongly precautionary that it leaves little scope for variation among the five states or response to parameter changes. Yet it might be the logical result of a strict application of findings from financial markets.

7 Conclusion

We have successfully introduced Epstein–Zin utility into the DICE model, and confirmed the findings of similar analyses with simpler climate models: with Epstein–Zin utility, the SCC is higher and the optimal pace of emission reduction is much faster than with the conventional DICE utility function; the results are sensitive to the intertemporal elasticity of substitution but remarkably insensitive to the level of risk aversion.

The latter finding, though, poses an unresolved puzzle. The threat of unchecked climate change is often described in terms of catastrophic risks of low but growing probability; how can an economic analysis of this phenomenon find that the level of risk aversion is unimportant? The apparent minimization of catastrophic risk in DICE is a logical consequence of the model's strong and optimistic assumptions about mitigation. The only climate risk that arises in DICE, the threat of escalating temperature increases, can be prevented with certainty, at limited cost, in a comparatively short period of time. If this were true of all climate risks, then the available, precautionary responses to the worst risks would clearly pass a cost-benefit test; rational policymakers would accept and act on this analysis, and the climate problem would be solved.

The optimistic assumptions made by DICE are not unique to that model, but are shared by many economic analyses. Many climate modelers have developed scenarios for rapid abatement, often leading to phasing out virtually all carbon emission by the end of this century. Ackerman and Stanton (2012) demonstrate that under plausible hypotheses about future uncertainties, the SCC is at or above the marginal cost of several rapid abatement scenarios. If such abatement scenarios could reliably eliminate all catastrophic climate risks, then DICE's optimism would be well-founded, complete mitigation would pass a cost-benefit test, and the level of risk aversion would indeed be of secondary importance. In such a world, catastrophic future outcomes would be avoided by immediate mitigation, as seen in Fig. 1, above, for states S_3, S_4, and S_5—with roughly the same response at many levels of risk aversion.

Unfortunately, we may not live in such a world. Climate science has emphasized the potential for "tipping points" at which the earth's climate may make a relatively abrupt, and perhaps irreversible, transition to a different state; once the tipping point has been identified with certainty, it may be too late to avoid it (Lenton et al. 2008). A deeper explanation of the unimportance of risk aversion in this paper's results is that the most important risks may be overlooked by models such as DICE. If catastrophic risks cannot be prevented with certainty, if rival policy responses differ widely in costs as well as probabilities of success, or if there are long delays between the initiation of policy responses and the reduction in risk, then there is a need for a different analysis in which risk and risk aversion play a central role. Incorporating such approaches to risk into an integrated assessment model is a challenge for future modeling efforts.

The issues raised here in the context of DICE may also be relevant to more complex integrated assessment models of climate change. While other models may offer greater detail and precision than DICE in modeling climate and economic dynamics, they face similar dilemmas to those identified here if they rely on the expected-utility framework. Use of

CRRA utility functions leads to the contradictory treatment of risk discussed above, which the Epstein–Zin framework seeks to address. Epstein–Zin utility combined with the simplified model of learning under uncertainty used in this paper may be of even greater importance in complex models with more state variables, where more sophisticated models of learning would introduce an infeasible computational burden.

New approaches to modeling risk need to be approached carefully; any treatment of truly catastrophic risk has the potential to overwhelm ordinary economic analysis. Weitzman's "dismal theorem" (Weitzman 2009) demonstrates that if knowledge of the climate system is necessarily incomplete, and the disutility of worst-case outcomes (such as extinction of the human race) is unbounded, then the marginal benefit of emission reduction is literally infinite. This implies the implausible conclusion that willingness to pay for emission reduction should approach 100 % of income. The problem is not unique to the dismal theorem; in a pattern that has been dubbed the "tyranny of catastrophic risk," any unbounded risk, no matter how small its (non-zero) probability of occurrence, will lead to the same implausible result (Buchholz and Schymura 2012). Since limitless willingness to pay for risk avoidance seems implausible, there must be a limit to the disutility of even the worst risks.

In the end, we are left with an agenda for further research, falling somewhere in the gap between the optimistic, risk-minimizing assumptions of DICE and similar models on the one hand, and the risk-obsessed framework of the dismal theorem and the tyranny of catastrophic risk on the other hand. There is a need to develop subtler analyses of catastrophic climate risks that cannot or will not be quickly prevented with certainty, applying realistic measures of societal risk aversion to calibrate appropriate responses.

Acknowledgments Thanks to Robert Litterman and Kent Daniel for many insightful discussions, and to the Litterman Family Foundation for financial support.

References

Ackerman, F, Stanton, EA (2012) Climate risks and carbon prices: revising the social cost of carbon. Econ E-J 6:2012–10. doi:10.5018/economics-ejournal.ja.2012-10

Ackerman F, DeCanio SJ, Howarth RB, Sheeran K (2009) Limitations of integrated assessment models of climate change. Clim Change 95(3–4):297–315. doi:10.1007/s10584-009-9570-x

Atkinson G, Dietz S, Helgeson J, Hepburn C, Sælen H (2009) Siblings, not triplets: social preferences for risk, inequality and time in discounting climate change. Econ E-J 3:2009–2026: doi:10.5018/economics-ejournal.ja.2009-26

Bansal R, Yaron A (2004) Risks for the long run: a potential resolution of asset pricing puzzles. J Financ 59:1481–1509. doi:10.1111/j.1540-6261.2004.00670.x

Buchholz W, Schymura M (2012) Expected utility theory and the tyranny of catastrophic risks. Ecol Econ 77:234–239. doi:10.1016/j.ecolecon.2012.03.007

Cass D (1965) Optimum growth in an aggregative model of capital accumulation. Rev Econ Stud 32:233. doi:10.2307/2295827

Cline WR (1992) The economics of global warming. Institute of International Affairs, Washington, DC

Crost B, Traeger C (2010) Risk and aversion in the integrated assessment of climate change. University of California eScholarship. http://escholarship.org/uc/item/1562s275

DeLong JB, Magin K (2009) The U.S. equity return premium: past, present, and future. J Econ Perspect 23(1):193–208. doi:10.1257/jep.23.1.193

Epstein, LG, Zin SE (1989) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57(4):937–969. doi:10.2307/1913778

Epstein LG, Zin SE (1991) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis. J Polit Econ 99(2):263–286. doi:10.2307/2937681

Ha-Duong M, Treich N (2004) Risk aversion, intergenerational equity and climate change. Environ Resour Econ 28(2):195–207. doi:10.1023/B:EARE.0000029915.04325.25

Howarth RB (2003) Discounting and uncertainty in climate change policy analysis. Land Econ 79(3):369–381. doi:10.3368/le.79.3.369

Jensen S, Traeger C (2011) Growth and uncertainty in the integrated assessment of climate change. Ragnar Frisch Centre for Economic Research (Oslo) and University of California-Berkeley. http://www.webmeets.com/files/papers/AERE/2011/141/NCCS.pdf

Kaufman N (2012) The bias of integrated assessment models that ignore climate catastrophes. Clim Change 110(3):575–595. doi:10.1007/s10584-011-0140-7

Koopmans TC (1965) On the concept of optimal economic growth. Cowles Foundation Paper. http://cowles.econ.yale.edu/P/cp/p02a/p0238.pdf

Lenton TM, Held H, Kriegler E, Hall JW, Lucht W, Rahmstorf S, Schellnhuber HJ (2008) Tipping elements in the earth's climate system. Proc Natl Acad Sci USA 105(6):1786–1793. doi:10.1073/pnas.0705414105

Mehra R, Prescott EC (1985) The equity premium: a puzzle. J Monet Econ 15(2):145–161. doi:10.1016/0304-3932(85)90061-3

Murphy JM, David MH (2004) Quantification of modelling uncertainties in a large ensemble of climate change simulations. Nature 430(7001):768–772. doi:10.1038/nature02771

Nordhaus WD (2008) A question of balance: economic modeling of global warming. Yale University Press, New Haven, CT

Rabin M (2000) Risk aversion and expected-utility theory: a calibration theorem. Econometrica 68(5):1281–1292. doi:10.1111/1468-0262.00158.

Ramsey FP (1928) A mathematical theory of saving. Econ J 138(152):543–559

Roe GH, Baker MB (2007) Why is climate sensitivity so unpredictable? Science 318(5850):632. doi:10.1126/science.1144735

Stern N (2006) The economics of climate change: the Stern review. Cambridge University Press, Cambridge

Vissing-Jørgensen A, Orazio P (2003) Stock-market participation, intertemporal substitution, and risk-aversion. Am Econ Rev 93:383–391. doi:10.1257/000282803321947399

Weitzman ML (2009) On modeling and interpreting the economics of catastrophic climate change. Rev Econ Stat 91(1):1–19. doi:10.1162/rest.91.1.1