SUPPLEMENT

ON THE PHYSICS OF THREE INTEGRATED ASSESSMENT MODELS

Raphael Calel and David A. Stainforth

This document is a supplement to "On the Physics of Three Integrated Assessment Models," by Raphael Calel and David A. Stainforth (*Bull. Amer. Meteor. Soc.*, **98,** 1199–1216) • ©2017 American Meteorological Society • Corresponding author: Raphael Calel, raphael.calel@georgetown.edu • DOI:10.1175/BAMS-D-16-0034.2

FROM PHYSICS TO IAMS. IAMs simulate the global economic impacts of climate change under a variety of mitigation scenarios. They assess the trade-off between the costs of mitigation action and the damages resulting from the remaining changes in physical climate. In the three IAMs considered here—DICE, FUND, and PAGE—these damages are expressed as functions of the change in global-mean, annual-mean surface temperature anomaly T (even this simple statement is not obvious, as we will see below), and the current value of a particular mitigation policy therefore depends on the time path of T that results. A central task of the climate component in IAMs is therefore the calculation of the trajectory of T over time.

To understand the physical basis for this, we begin with some simple descriptions of the climate system in terms commonly found in the physical science literature. These provide a foundation from which we can derive the temperature equations in the three IAMs analyzed in this paper.

Some simple climate physics. Arguably the simplest model of climate change considers the climate system as a single box in which the rate of change in energy content is equated to changes in incoming radiation balanced by changes in outgoing radiation. This is often represented by Eq. (ES1) (Andrews and Allen 2008; Senior and Mitchell 2000; Dickinson 1986) in which the change in incoming radiation is taken to be the radiative forcing due principally to changes in atmospheric greenhouse gas concentrations F, and the change in outgoing radiation is taken to be proportional to changes in global-mean, annual-mean surface temperature λT:

$$C_{\text{eff}} \frac{dT}{dt} = F - \lambda T \,, \qquad \text{(ES1)}$$

where C_{eff} is effective heat capacity of the climate system, $T(t)$ is surface temperature change from some equilibrium state, $F(t)$ is radiative forcing, λ is feedback parameter, and t is time.

The effective heat capacity in Eq. (ES1) is largely a result of the heat reservoir of the upper oceans. A simple extension of this model is to allow for the diffusion of heat to the deep oceans by adding a second box, the "deep oceans," which exchanges heat with the surface (or upper ocean) according to a one-dimensional heat transfer equation. This gives us the system

$$\begin{aligned} C_{\text{up}} \frac{dT}{dt} &= F - \lambda T - \beta\left(T - T^{\text{LO}}\right) \\ C_{\text{deep}} \frac{dT^{\text{LO}}}{dt} &= \beta\left(T - T^{\text{LO}}\right) \end{aligned} \,, \qquad \text{(ES2)}$$

where C_{up} is effective heat capacity of the upper oceans, land surface, and atmosphere; C_{deep} is effective heat capacity of the deep oceans; T^{LO} is deep ocean temperature change from some equilibrium state; and β is heat transfer coefficient.

Having honed the whole of physical climate science down to a simple physical model such as one of the above, the developer of an IAM like DICE, FUND, or PAGE must select a method for numerically representing such a model on a computer. Let us therefore look now at what modeling choices the developers of these three main IAMs have made.

DICE. The simplest numerical approach is a forward Euler discretization in which the value one time step in the future is calculated in terms of the present values. Applying this technique to the model represented by the system of equations in Eq. (ES2) returns the system of equations

$$T_t = T_{t-1} + \frac{\Delta t}{C_{up}}\left[F_{t-1} - \lambda T_{t-1} - \beta\left(T_{t-1} - T^{LO}_{t-1}\right)\right]$$
$$T^{LO}_t = T^{LO}_{t-1} + \frac{\beta\Delta t}{C_{deep}}\left(T_{t-1} - T^{LO}_{t-1}\right), \quad \text{(ES3)}$$

where t is the number of the time step (not continuous time).

The terms in Eq. (ES3) have been rearranged compared to Eq. (ES2) to facilitate comparison with the equations used in DICE. It is now plain to see that the system of Eq. (ES3) corresponds almost exactly to the temperature equations in DICE, with the following notational changes:

$$\xi_1 = \frac{\Delta t}{C_{up}} \text{ and } \xi_2 = \lambda \text{ and } \xi_3 = \beta \text{ and } \xi_4 = \frac{\beta\Delta t}{C_{deep}}. \quad \text{(ES4)}$$

The only remaining difference is that DICE assumes that forcing has been constant at the current level (F_t) over the last time step, as opposed to assuming that it has remained constant at the level at the beginning of the period (F_{t-1}). The practical distinction between these assumptions depends on the length of the time step; the two converge as the time step is reduced. DICE-2013R uses a 5-yr time step, while earlier versions of DICE use a 10-yr time step. All versions have hard coded the length of the time step in other parameter values, except for some economic equations in DICE-2013R.

FUND. Applying Euler's discretization method to the model represented by Eq. (ES1) and rearranging for T_t yields

$$T_t = \left(1 - \frac{\lambda\Delta t}{C_{eff}}\right)T_{t-1} + \frac{\Delta t}{C_{eff}}F_{t-1}. \quad \text{(ES5)}$$

Notice that T_t is a linear combination of forcing and temperature. The temperature equation in FUND takes the same form, although FUND (like DICE) uses forcing at the end of the period (F_t) instead of at the beginning (F_{t-1}). Notationally, let us write

$$\phi = \frac{C_{eff}}{\lambda\Delta t} \text{ and } \lambda = \frac{5.35\ln(2)}{CS}, \quad \text{(ES6)}$$

where ϕ is e-folding time of temperature in units of time steps and CS is climate sensitivity. Substituting these expressions back into Eq. (ES5), and replacing F_{t-1} with F_t, directly yields the equation for global mean temperature in FUND:

$$T_t = \left(1 - \frac{1}{\phi}\right)T_{t-1} + \frac{1}{\phi}\frac{CS}{5.35\ln(2)}F_t. \quad \text{(ES7)}$$

Note that the two equalities in Eq. (ES6) are correspondences, not assumptions. This point is made clear by an example. A simple manipulation of Eq. (ES1), on which this model is based, gives $CS = F_{2\times CO_2}/\lambda$, which yields Eq. (ES6) for γ. The generally accepted value of $F_{2\times CO_2}$ is 3.7 W m^{-2}, which corresponds to the value assumed in FUND of 5.35ln(2) W m^{-2}.

It should be noted, further, that FUND calculates ϕ by the following equation:

$$\phi = \max(\alpha + \beta CS + \gamma CS^2, 1).$$

The parameter values are chosen so as to reflect a "best guess" that $\phi = 44$. This general type of behavior, although not necessarily its specific form, would follow if one believes that the climate sensitivity and the effective heat capacity are constant over time and can be deduced from twentieth-century observations. In that case a higher climate sensitivity must be matched with a higher effective heat capacity to maintain a similarly good fit to twentieth-century observed changes. This functional relationship between ϕ and CS captures this behaviour.

PAGE. Observe that Eq. (ES1) has an analytic solution for a fixed value of F, say F'. The corresponding equilibrium value of T, call it ET, is given by $ET = F'/\lambda$. Substituting this expression into Eq. (ES1) and rearranging yields

$$\frac{d(ET - T)}{ET - T} = -\frac{\lambda}{C_{eff}}dt.$$

Integrating on both sides and rearranging returns the following expression:

$$T = ET - Ae^{-\lambda t/C_{eff}}, \quad \text{(ES8)}$$

where A is the constant of integration. In words, the temperature approaches the new equilibrium as an exponential function of time. In the case of time-dependent forcing, as is the situation studied with an IAM, Eq. (ES8) can be used to calculate the change in T over a time step Δt. Assuming the system starts in an equilibrium state (i.e., $T = 0$ when $t = 0$), we get

$$T_t = T_{t-1} + \left(\mathrm{ET}_t - T_{t-1}\right)\left(1 - e^{-\lambda \Delta t / C_{\mathrm{eff}}}\right). \quad \text{(ES9)}$$

By making the following notational changes, one can see that this equation is equivalent to that used in PAGE:

$$\frac{\mathrm{SENS}}{\mathrm{FSLOPE}_1 \ln(2)} F_t = \mathrm{ET}_t \text{ and}$$

$$\mathrm{OCEAN} = \frac{C_{\mathrm{eff}}}{\lambda} \text{and} \, Y_t - Y_{t-1} = \Delta t \,. \quad \text{(ES10)}$$

As in the other IAMs, PAGE takes the forcing to be F_t during the time step.

PAGE does not calculate the change in global mean temperature directly from this model, though. PAGE applies this temperature equation separately to eight world regions and then computes the change in global mean temperature as an area-weighted sum of these eight regional mean temperature anomalies. The result is an equation mirroring the one above, so for the purpose of forecasting the change in global mean temperature, the above derivation of the global temperature equation is equivalent to deriving eight regional temperature equations and averaging them.

To see this, we must distinguish between the global temperature anomaly, GT, and the temperature anomaly in region r, RT_r, the latter being determined from the above equation. In PAGE the global temperature anomaly is taken as a weighted average of regional temperature anomalies;

$$\mathrm{GT}_t = \sum_r w_r \mathrm{RT}_{rt} \,,$$

where w_r is the weight of region r, with $\sum_r w_r = 1$. For a system of two regions, this gives

$$\begin{aligned} \mathrm{GT}_t &= w_1 \mathrm{RT}_{1t} + w_2 \mathrm{RT}_{2t} \\ &= w_1 \mathrm{RT}_{1t-1} + w_2 \mathrm{RT}_{2t-1} + \left[\frac{\mathrm{SENS}}{\mathrm{FSLOPE}_1 \ln(2)}\left(w_1 F_1 + w_2 F_2\right) - \left(w_1 \mathrm{RT}_{1t-1} + w_2 \mathrm{RT}_{2t-1}\right)\right] \\ &\quad \times \left(1 - e^{-(Y_t - Y_{t-1})/\mathrm{OCEAN}}\right). \end{aligned}$$

Since $w_1 + w_2 = 1$, we have $\mathrm{GT}_{t-1} = w_1 \mathrm{RT}_{1t-1} + w_2 \mathrm{RT}_{2t-1}$ and $F_t = w_1 F_{1t} + w_2 F_{2t}$. The above equation therefore simplifies to

$$\mathrm{GT}_t = \mathrm{GT}_{t-1} - \left[\frac{\mathrm{SENS}}{\mathrm{FSLOPE}_1 \ln(2)} F_t - \mathrm{GT}_{t-1}\right]\left(1 - e^{-(Y_t - Y_{t-1})/\mathrm{OCEAN}}\right),$$

which is identical to the temperature change equation derived earlier. The same argument holds for n regions. For the purposes of forecasting changes in global temperature, therefore, the procedure of averaging eight regions in PAGE is equivalent to deriving a global temperature change equation from a global energy balance model.

A note on MAGICC. Our article concentrates on three IAMs—DICE, FUND, and PAGE. As noted in the main text there are a number of other IAMs, and many of them use the Model for the Assessment of Greenhouse Gas Induced Climate Change (MAGICC) as their representation of physical climate. It is therefore worth looking at MAGICC in a bit more detail and thinking about whether the physics of MAGICC could be fruitfully brought into our discussion.

MAGICC is essentially used as an emulator of more complicated atmosphere–ocean global circulation models (AOGCMs). Indeed, the peer-reviewed paper describing the latest version of the model is entitled "Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6" (Meinshausen et al. 2011). It is a climate model of intermediate complexity built to facilitate the study of a wider range of emission trajectories than is possible with those more complicated models directly. Three important things follow from this.

First, MAGICC is a stand-alone climate model, not built for the purpose of calculating the economic consequences of climate change. It has been used together with economic models to address certain questions about climate policy (Clarke et al. 2014), but this does not include the sort of global cost–benefit assessments done with DICE, FUND, and PAGE (Stern 2007; Watkiss and Hope 2011; Greenstone et al. 2013; Interagency Working Group 2015).

Second, MAGICC was developed to emulate AOGCMs rather than to cope with the additional computational demands attendant to economic analysis. As such, it can and does offer a much more complex representation of the climate system than does DICE, FUND, or PAGE. MAGICC has an upwelling-diffusion-entrainment ocean that represents the Northern Hemisphere (NH) and the Southern Hemisphere (SH) separately. It also allows for different feedback parameters over ocean and over land—again

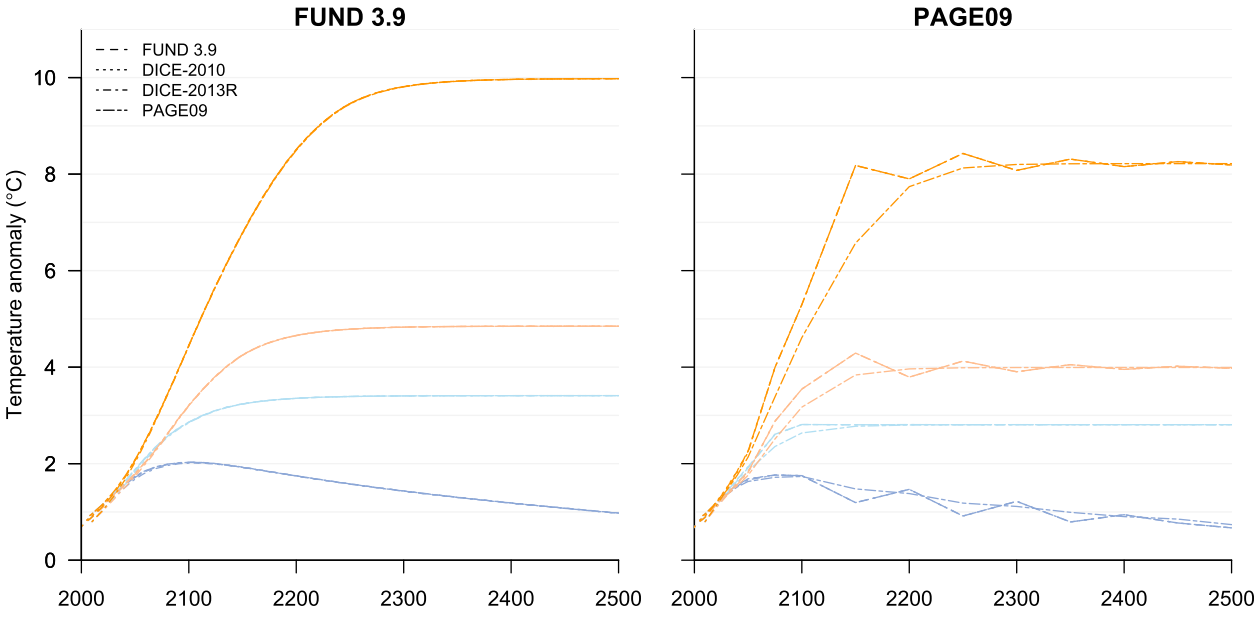

Fig. ES1. Temperatures with time-varying heat capacities: (left) Replicating the leftmost panel of Fig. 2 in the main text, but lets the heat capacities in DICE vary over time so as to produce the same constant value of the effective heat capacity as is assumed in FUND. (right) Replicating the rightmost panel of Fig. 2 in the main text, now letting the heat capacities in DICE vary over time so as to produce the same constant value of the effective heat capacity as is assumed in PAGE.

with a separation of NH and SH. In its latest incarnation, MAGICC6, eight parameters are tuned to achieve a good emulation of a variety of AOGCMs (Meinshausen et al. 2011). Some of its parameters enable emulation of the time variations of the feedback parameter and ocean heat uptake that are seen in AOGCMs. The flexibility provided by these parameters makes it an extremely powerful and useful tool for emulating some of the behavior of AOGCMs but also makes MAGICC largely irrelevant in a discussion of the physics of DICE, FUND, and PAGE. The two sets of models embody different levels of physical complexity.

Third, MAGICC's design and use as an AOGCM emulator means that its physical assumptions are essentially those of the emulated AOGCMs. The typical parameter values used in MAGICC simply reflect, by design, the behavior of the major AOGCMs. In the main text we contrast the temperature responses with those found in the CMIP5 ensemble—nothing would be added to this by examining MAGICC output. Because MAGICC is an emulator, a discussion about the physical assumptions in MAGICC would add little if anything to the substantial literature comparing many different aspects of AOGCMs and exploring uncertainty in their formulation (e.g., multimodel and perturbed-physics ensembles).

For these reasons, we have set aside MAGICC and other intermediate-complexity climate models and focused our discussion on the impact of different physical assumptions in the simpler IAMs—DICE, FUND, and PAGE. This allows us to go beyond simply documenting the models and to consider how the best physical science understanding can be brought to bear in future development and use of these models. It may be interesting to study how the conclusions of MAGICC are sensitive to each of the eight tuned parameters (or to more of the parameters in the model) and to relate them to observational or physical constraints. This would extend the existing model uncertainty and perturbed-parameter literature but is substantially beyond the scope of the present work.

REPRESENTING HEAT CAPACITIES. While FUND and PAGE define a set of parameters, including *e*-folding times, that in combination define the effective heat capacity of the climate system C_{eff}, DICE defines a set of parameters that define the effective heat capacities of two different parts of the system, C_{up} and C_{deep}. There is obviously a relationship between these two descriptions of the climate system: a given value of C_{eff} can be related to a pair of values for C_{up} and β given particular values of the temperature of the upper and lower ocean at some point in time. At any point in time one can therefore calculate the value of C_{up} implied by a value of C_{eff}, and vice versa. Specifically, we have the following relationship:

$$C_{\mathrm{up},t} = C_{\mathrm{eff},t}\,\frac{F_t - \lambda T_t - \beta\left(T_t - T_t^{\mathrm{LO}}\right)}{F_t - \lambda T_t}.$$

In this paper we assume that β is constant, as DICE does. Thus, whenever the system is out of equilibrium a constant value of C_{eff} implies a time-varying C_{up} and, conversely, a constant value C_{up} implies a time-varying C_{eff}. Hence, if the model parameters are constant we can only ensure that C_{up} is "equivalent" to a given C_{eff} (and vice versa) in a single period—that is, that PAGE, FUND, and DICE have equivalent thermal inertia in a single period. The comparable parameter simulations presented in this paper assume that this equivalence occurs in the first period for which all models produce a forecast.

It is easy to demonstrate that these mutually exclusive constancy restrictions are the cause of many of the observed differences between DICE and FUND/PAGE. As Fig. ES1 illustrates (also in the main text), the systematic difference between the IAMs vanishes when we remove these restrictions by recalculating either C_{eff} or C_{up} at each time step, thereby forcing one model structure to mirror the response of the other, while also ensuring parametric equivalence by setting the feedback parameters to be the same.

The main text already highlights some important consequences of representing the climate's thermal inertia through *e*-folding times or through heat capacities, to which we would add two further points. First, in physical terms the *e*-folding time equates to C_{eff}/λ, as can be seen from Eq. (ES8). Radiative feedbacks (as represented by ϕ) and effective heat capacity involve largely independent physical processes so there is little reason to expect them to be physically correlated, even though their statistical quantification from the observational record is indeed correlated (see main text). There may therefore be advantages, particularly when considering how these variables may vary in the future, to express thermal inertia by C_{eff}, which is independent of γ (and effective CS), rather than *e*-folding time, for which some relationship must be specified. FUND and PAGE both do explicitly define a relationship between the *e*-folding time and the climate sensitivity, but these embedded relationships go beyond what is encapsulated in the underlying physical climate model.

Second, as the climate changes we would expect to see changes in both ocean circulation patterns (Gregory et al. 2005; Collins et al. 2013) and radiative feedbacks (Senior and Mitchell 2000), which will change both the effective heat capacity of the system and the relevant value of the climate sensitivity to use within these models. The effect of the former could be modeled by introducing a more complicated ocean—such as the two-equation model used in DICE—or more simply by allowing either the effective heat capacity or *e*-folding time to vary over time. A time-dependent effective heat capacity can be readily interpreted in terms of changing ocean circulation, but a time-dependent *e*-folding time undermines the concept on which *e*-folding times are based; that is, that the system is relaxing toward a new equilibrium with a single characteristic time scale.

To discuss responses in terms of heat capacities may be alien to economists, but the thermal inertia of the climate system is better expressed this way because it is closer to quantities that can be constrained by historical observations, it is independent of radiative feedback uncertainties, and it is more amenable

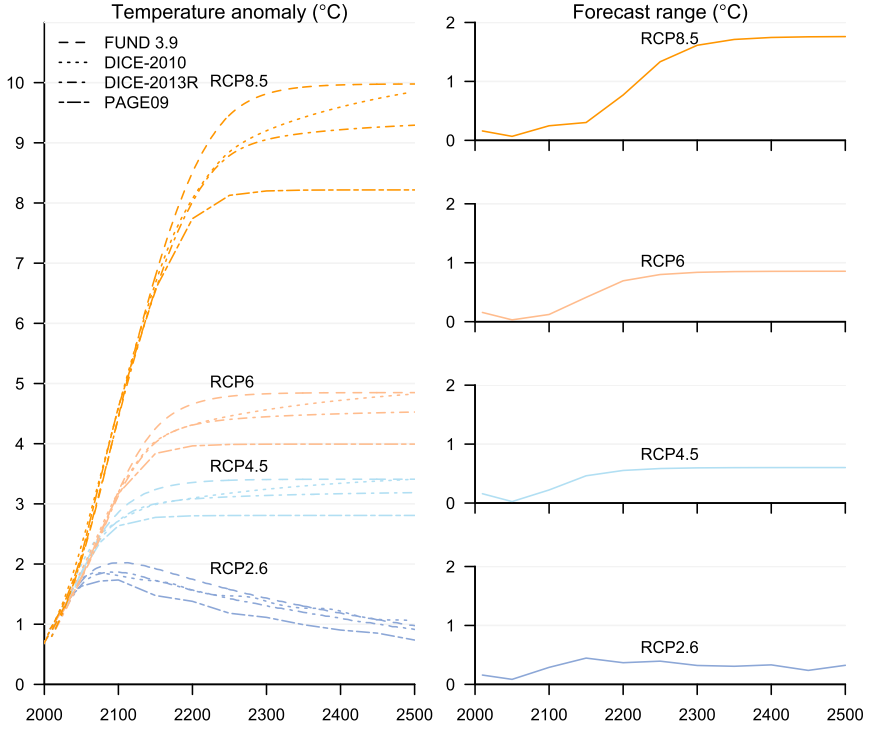

Fig. ES2. Temperatures under RCP forcing scenarios up to 2500: extends the time axis of Fig. 1 in the main text up to the year 2500.

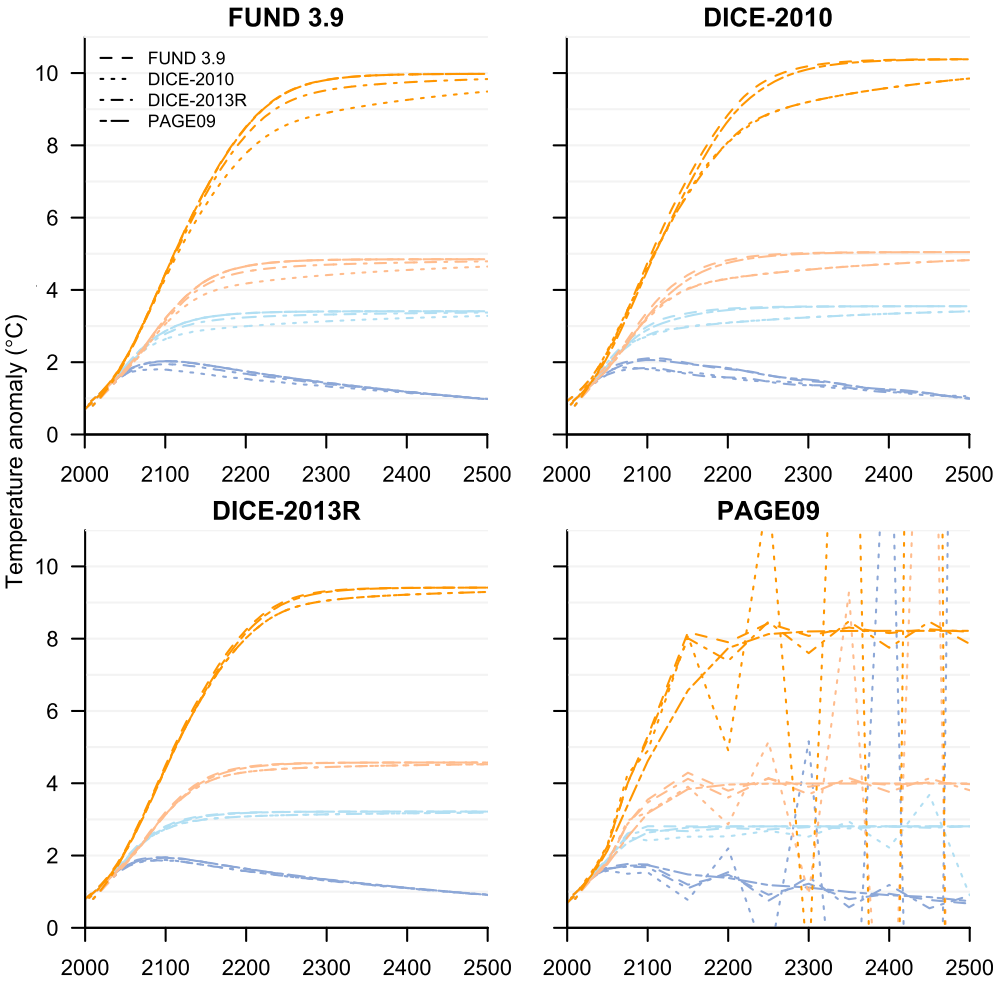

Fig. ES3. Temperatures with comparable parameter values up to 2500: extends the time axis of Fig. 2 in the main text up to the year 2500.

to physics-based speculation regarding how it might change in the future.

NUMERICAL REPRESENTATION. DICE and FUND are more vulnerable to numerical instabilities. PAGE's comparative stability comes about because it uses an analytic solution to the energy balance equation to forecast temperatures while DICE and FUND use a forward Euler method, as we discussed earlier.

The instabilities in DICE and FUND are evident when their forecasts overshoot the analytically derived temperature, after which they fall back down and can begin to oscillate. If the initial overshoot is small, the oscillations dissipate in subsequent periods, but if it is sufficiently large the oscillations can grow over time. This is a very well-known instability in numerical analysis and is commonly found in Euler methods. It is seen in the lower-right panel of Fig. 2 in the main paper and Fig. ES1 here because of applying the long time steps used in PAGE.

This behavior is not unique to large time steps, however. It can also arise with shorter time steps if the other parameters take more outlying values in their uncertainty distributions. An example of this can be found in Calel et al. (2015), where DICE is run with lower, but still scientifically plausible, values of heat capacity and climate sensitivity. This combination can result in overshooting, and subsequent numerical instability, and DICE does indeed become numerically unstable under some such parameter assumptions, even with the time step used in DICE.

One could also imagine other parameter combinations that give rise to this behavior. If the length of the time step was not hard coded into IAMs, such instabilities could be easily removed by reducing the length of the time step, which would allow researchers to more easily explore these regions of the parameter space. An alternative discretization of DICE now exists (Cai et al. 2012), though it remains to be seen which will be taken forward and commonly developed.

FROM TEMPERATURE FORECASTS TO ECONOMIC DAMAGES. In this paper we use the temperature equations from DICE, PAGE, and FUND to forecast the path of global-mean, annual-mean surface temperature anomalies T under four RCP radiative forcing scenarios. Even though the RCP scenarios end in 2500, the figures in the paper show the forecasts only up to the year 2300 to match the time horizon used by the U.S. Interagency Working Group on the Social Cost of Carbon. This also tends to put greater visual emphasis on those temperature differences that are more economically significant. For completeness, we reproduce the figures here up to 2500, which shows both the adjustment and equilibrium behavior of the different representations of the climate system (Figs. ES3 and ES4).

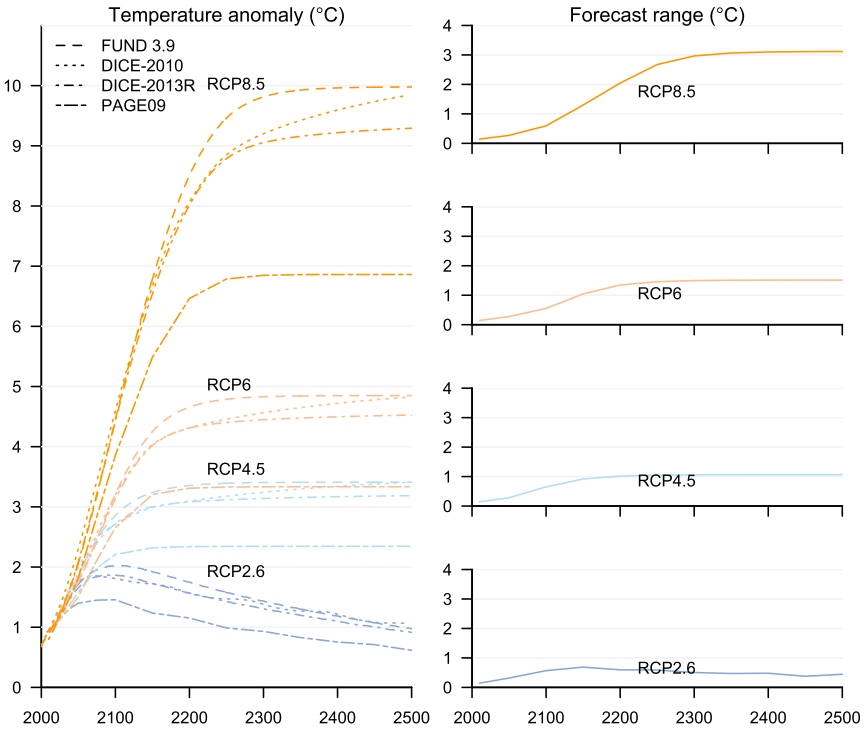

FIG. ES4. Temperatures with a different modal climate sensitivity in PAGE: replicates Fig. 1 in the main text, except that the equilibrium climate sensitivity in PAGE is set to 2.1°C, which corresponds to setting the transient climate response and the e-folding time to their modal values.

It is straightforward to compute the difference between forecasts at a given point in time, as in Fig. ES2. But we require more information to convert these forecast differences into economic summary measures.

It is of course incredibly difficult to anticipate the loss of economic output associated with a given T—it requires an accounting of altered agricultural conditions, of stresses on social and political institutions, of destruction of people and property from extreme weather events and the spread of disease vectors, and so on. While a large and growing literature examines these relationships, urgency has demanded that integrated modelers adopt a provisional solution. An instantaneous damage function, of the kind commonly found in the literature, takes T and returns a fraction of gross economic output that is lost to climate-induced damages, $D(T)$. We consider both Nordhaus's quadratic damage function (Nordhaus and Sztorc 2013) and Weitzman's polynomial damage function (Weitzman 2012). If gross economic output (i.e., output in the absence of climate damages) is Y, net economic output is $Y[1 - D(T)]$.

There are different ways to measure the utility derived from this economic output. Generally, one applies a discount to dollars that accrue to wealthier individuals and to dollars that accrue further into the future. In a simple optimal growth model, the discount rate can be written as the sum of the pure rate of time preference ρ and the product of economic growth g and the elasticity of marginal utility η: $r = \eta g + \rho$; η measures society's aversion to inequality (and risk), while ρ measures society's aversion to waiting. There is little agreement about what value the social discount

TABLE ES1. Differences in economic damages: assuming g = 1%.

r =	2%	3%	5%	2%	3%	5%
	RCP2.6			RCP4.5		
Nordhaus	0.17	0.12	0.08	0.41	0.20	0.09
Weitzman	0.17	0.11	0.07	0.80	0.35	0.11
	RCP6			RCP8.5		
Nordhaus	0.64	0.26	0.09	1.89	0.56	0.12
Weitzman	3.31	1.02	0.16	7.11	2.64	0.50

rate r should take, but a recent survey of experts found that over 90% felt most comfortable with r somewhere in the range from 1% to 3%, with a central value of 2% (Drupp et al. 2015). The U.S. Interagency Working Group on the Social Cost of Carbon uses social discount rates of 2.5%, 3%, and 5% (Interagency Working Group 2015). Here we consider values of 2%, 3%, and 5%.

Table ES2. Differences in economic damages: assuming g = 3%.

r =	2%	3%	5%	2%	3%	5%
	RCP2.6			RCP4.5		
Nordhaus	0.14	0.19	0.12	1.01	0.81	0.21
Weitzman	0.13	0.18	0.12	2.25	1.73	0.36
	RCP6			RCP8.5		
Nordhaus	2.08	1.54	0.27	10.80	6.59	0.59
Weitzman	15.32	10.20	1.07	145.88	28.65	2.74

Conventionally, η is set somewhere in the range from 1 to 1.5, which tends to produce a social discount rate of circa 5% when ρ is around 2%–3%. The nonlinearity in the utility function (which is represented by η being greater than 0) means that the conversion from economic output to utility depends not only on the rate of economic growth g but also on who gets how much, and when. In the presence of this nonlinearity, the conversion becomes a function of additional assumptions about population growth and economic inequality, which substantially complicates computation and interpretation. To make interpretation as straightforward as possible, the results shown in the paper assume that every dollar is worth as much as the next. In the jargon of economics, we use a linear utility function, or an isoelastic utility function (common in IAMs) with no aversion to inequality (i.e., η= 0). In this simple case, the rate at which future consumption is discounted is equal to the pure rate of time preference, η. We choose values of η to achieve comparable social discount rates r to those used in more complex assessments in the literature.

Under this assumption, we can write the net present value (NPV) of the flow of gross economic output up to some terminal time S as

$$\mathbb{Y} = \sum_{s=0}^{S} \frac{1}{\left(1+\rho\right)^{s}} Y_0 \left(1+g\right)^{s},$$

where $s = 0,\ldots,S$ indexes time periods and Y_0 is current world output. Notice that the growth rate of gross economic output g enters into this expression. To compute the NPV of future output, we must know not only the rate at which we discount it, but also the rate at which it would grow in the absence of climate damages. Similarly, we can write the NPV of the flow of damages as

$$\mathbb{D} = \sum_{s=0}^{S} \frac{1}{\left(1+\rho\right)^{s}} Y_0 D\left[T\left(s\right)\right]\left(1+g\right)^{s}.$$

In a moment we shall see how these formulas are used to express differences between temperature forecasts in economic terms. But first, we should note the complications arising from the appearance of g in these formulas.

By design, our analysis takes the time series of radiative forcing as given in order to isolate the behavior of the climate components of DICE, PAGE, and FUND. Since we do not run the economic components explicitly, though, we cannot check what value g would have to be in each IAM in order for that IAM to produce the particular forcing time series that we use. Fortunately, this does not put us at a dead end. Since there will in principle be many combinations of economic parameter values and initial conditions that produce these particular forcing time series, we could imagine picking the values in order to constrain the rate of economic growth to equal the g of our choice. The apparent obstacle of an unknown g is therefore resolved by reference to the many degrees of freedom.

Table ES3. Differences in economic damages: assuming g = 2% and values of {ρ,η} indicated below.

	{0.1%,0.95}	{1%,1}	{2%,1.5}	{0.1%,0.95}	{1%,1}	{2%,1.5}
	RCP2.6			RCP4.5		
Nordhaus	0.04	0.04	0.03	0.16	0.10	0.05
Weitzman	0.04	0.04	0.03	0.35	0.20	0.07
	RCP6			RCP8.5		
Nordhaus	0.31	0.16	0.05	1.39	0.53	0.10
Weitzman	2.06	0.88	0.17	23.56	7.83	1.97

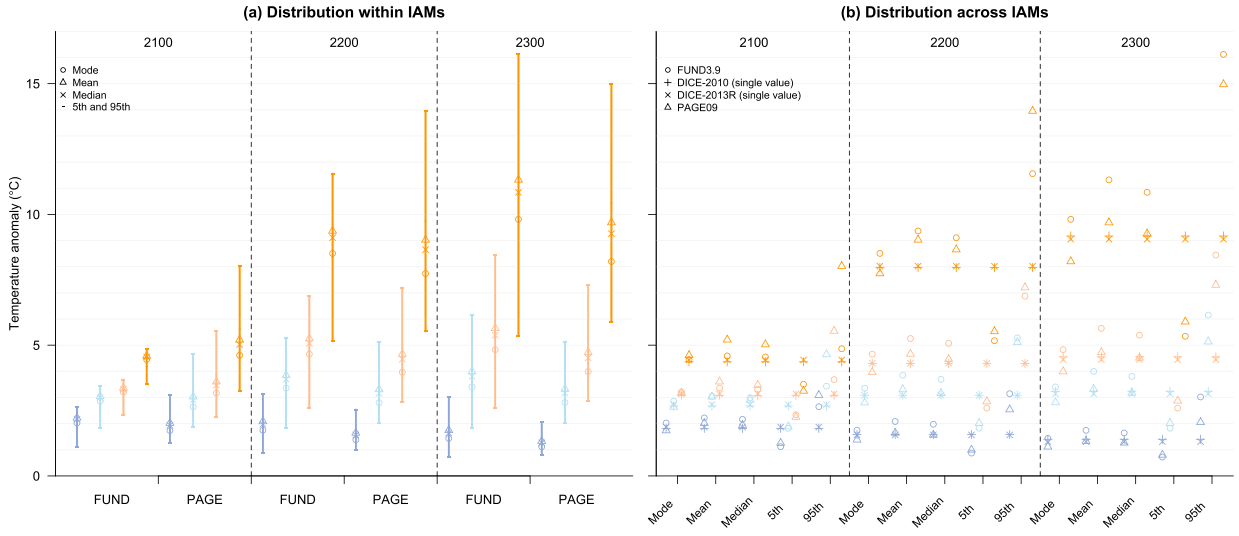

FIG. ES5. Temperatures with mode, mean, median, and extreme parameter values: (a) The temperature forecasts for FUND and PAGE when the mode, mean, median, 5th percentile, and 95th percentile parameter values are chosen. In PAGE, the climate sensitivity and the e-folding time are both random parameters, but the temperature forecast is generally increasing in the former and decreasing in the latter. To show the 95th percentile, therefore, we have opted to pair the 95th percentile of climate sensitivity distribution with the 5th percentile of the e-folding time distribution (conditional on the climate sensitivity), and vice versa for the 5th percentile. (b) Range of the temperature forecasts across the set IAMs, including DICE, when evaluated at different points along the parameter distributions. It shows more clearly that the range of forecasts is generally larger when we do not use the modal parameter values.

Returning to the issue at hand, we will now use the expressions above to compute an economic summary measure of the difference between forecasts. In essence, we measure the range of NPV damages across forecasts as a proportion of the average NPV of net output. Letting $\overline{\mathbb{D}}$ denote the set of damages $\mathbb{D}$ that correspond to the set of temperature forecasts in a particular forcing scenario, we compute

$$\frac{\max\left(\overline{\mathbb{D}}\right)-\min\left(\overline{\mathbb{D}}\right)}{\mathbb{Y}-\text{mean}\left(\overline{\mathbb{D}}\right)}.$$

The numbers presented in the paper are calculated under the assumption that $g = 2\%$. Tables ES1 and ES2 show the results for $g = 1\%$ and $g = 3\%$, which span IPCC's projections that the world average growth rate of income per capita will be between 1.3% and 2.8% for the next century, on average (Drupp et al. 2015). The complicated interaction between the social discount rate, damage function, and growth rate make it difficult to discern general patterns in the economic damages. Economic damages often increase with the growth rate, but not always. What is more, temperature differences on shorter time scales are sometimes emphasized by a higher social discount rate, as discussed in the main text. The main point here is merely that there is nothing special about $g = 2\%$. Physically based differences between these IAMs can be large enough to have economic significance when one operates within the range of plausible economic assumptions.

We can also repeat our main economic analysis using a nonlinear utility function, although this means we must make additional assumptions about population growth and economic inequality. Table ES3 repeats the analysis in Table 2 in the main paper with $\eta > 0$, assuming that world population grows linearly to 9 billion by 2050 and then stabilizes [this closely approximates the "middle scenario" in the U.N. (2004) long-range population projections] and a uniform distribution of consumption within each period (as in DICE). We obtain the same social discount rates as in the paper by choosing greater values of η and lower values of ρ. The resulting numbers are often, though not always, a bit lower than before, but it remains true that the physical differences between IAMs appear to be economically significant under some plausible economic assumptions.

ANALYSIS UNDER UNCERTAINTY. DICE solves for the optimal abatement policy using a single set of parameter values, while FUND and PAGE are designed to run many times with different parameter values to provide a distribution of consequences for a given policy. In this context, what is the best way

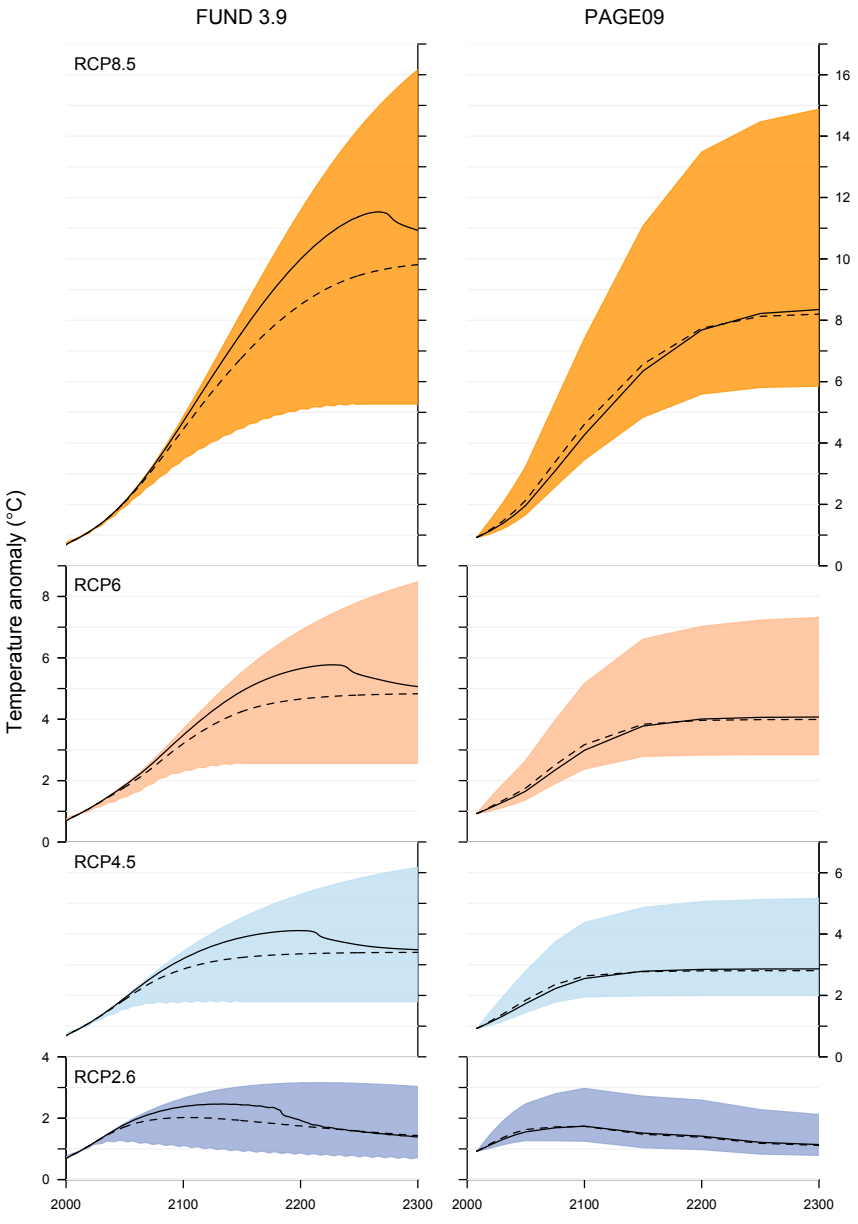

Fig. ES6. Temperature trajectory ensembles: The black solid lines trace the modes over time of 10,000-member ensembles of (left) FUND and (right) PAGE. For comparison, the dashed lines show the forecasts using modal parameter values, as in Fig. 1 of the main paper. The edges of the color-shaded areas trace the 5th and 95th percentiles of the ensembles. The colors correspond to the RCP scenarios in the AR5.

to study the consequences of their different physical structures and parameter values?

Perhaps the simplest and most transparent approach is to run each model once using modal parameter values (as we do in the paper). Even this, it turns out, is not entirely without difficulty. In PAGE, the climate sensitivity is a function of two random variables: transient climate response (mode = 1.3) and the half-life of global warming or *e*-folding time (mode = 30). The climate sensitivity corresponding to the modal values of these two random variables is roughly 2.1°C. The mode of the distribution of the climate sensitivity, however, is 2.54°C. It is thus impossible to set all three parameters equal to their modal values simultaneously. Since the transient climate response does not enter directly into the temperature forecasting equation, but the other two quantities do, we have opted to set the *e*-folding time to 30 yr, and the climate sensitivity to 2.54°C, which implies a transient climate response in PAGE of 1.56.[ES1] This turns out to be the more conservative comparison. If we instead let the *e*-folding time be 30 yr, and the climate sensitivity be 2.1°C, the IAMs produce a wider range of temperature forecasts (see Fig. ES4).

One might still be worried that this approach would exaggerate model differences if the parameter distributions were most dissimilar at their modes. This can be addressed most directly by looking at the model differences when the parameter values are chosen at other points along their respective distributions. Figure ES5 displays the range of forecasts when the parameters are set equal to the modes, means, medians, 5th percentiles, and 95th percentiles of the relevant distributions. When the parameter values are chosen at these alternative points along the parameter distributions, the temperature trajectories differ even more sharply than for the modal values (see Fig. ES5).

[ES1] The same issue also arises in both PAGE and FUND with other measures of central tendency, such as the mean.

An alternative approach would involve running ensembles of FUND and PAGE and comparing the resulting distributions of temperature forecasts. For our purposes, this approach has a few serious drawbacks, however. First and foremost, any summary statistic of the distribution of temperature forecasts (mode, median, mean, etc.) will not correspond to any particular set of parameter values. This severs the direct link between physical assumptions and the behavior of the temperature trajectory and, therefore, makes it all but impossible to give a clear physical interpretation of the differences between temperature forecasts. Second, the isoelastic utility function commonly used with IAMs has a single parameter to describe aversion to inequality and to risk (η). This means it treats inequality across space and time as formally identical to inequality across states of the world. Economic damages calculated from an ensemble of temperature trajectories therefore do not have the same interpretation as when they are calculated from a single model run. This makes it more difficult to compare FUND and PAGE with DICE. Alternative utility functions that separate aversion to inequality from aversion to risk instead present other interpretational and computational challenges (Epstein and Zin 1989), which would distract from our main objective here of understanding the physically based differences between models.

These drawbacks make it difficult to conduct a full replication of our work using the ensemble approach, but a look at some initial results suggests it is unlikely that we would end up concluding that the physical differences between IAMs are unimportant. Figure ES6 plots the mode and 5%–95% confidence region for 10,000-member ensembles of FUND and PAGE over time. Three features are worth remarking on. First, the mode of the temperature distribution in FUND exhibits quite different dynamic behavior from the temperature trajectory computed for modal parameter values. This is a consequence of the strong positive correlation between equilibrium climate sensitivity and effective heat capacity specified in FUND, which is a feature of the ensemble as a whole rather than a single run of the model. It is therefore difficult to give a clear physical interpretation of this behavior. In PAGE, which has a mild negative correlation between the equilibrium climate sensitivity and effective heat capacity, the ensemble median is generally slightly lower than when the model is run with median parameter values.

Second, and more to the point, we can see that the model runs using modal parameter values tend to understate the differences in the modes of the temperature distributions. The mode of the FUND ensemble is substantially higher than the forecast FUND produces with modal parameter values, while the mode of the PAGE ensemble is slightly lower. The differences are also pronounced at the 5th and 95th percentiles. These results show no indication that the output from different IAMs would suddenly appear very similar if only we looked at the distributions of temperature forecasts. The differences between IAMs would likely remain large enough that they will continue to matter under plausible economic assumptions.

REFERENCES

Andrews, D. G., and M. R. Allen, 2008: Diagnosis of climate models in terms of transient climate response and feedback response time. *Atmos. Sci. Lett.*, **9**, 7–12, doi:10.1002/asl.163.

Cai, Y., K. L. Judd, and T. S. Lontzek, 2012: Continuous-time methods for integrated assessment models. National Bureau of Economic Research Working Paper 18365, 44 pp., doi:10.3386/w18365.

Calel, R., D. A. Stainforth, and S. Dietz, 2015: Tall tales and fat tails: The science and economics of extreme warming. *Climatic Change*, **132**, 127–141, doi:10.1007/s10584-013-0911-4.

Clarke, L., and Coauthors, 2014: Assessing transformation pathways. *Climate Change 2014: Mitigation of Climate Change*, O. Edenhofer et al, Eds., Cambridge University Press, 413–510.

Collins, M., and Coauthors, 2013: Long-term climate change: Projections, commitments and irreversibility. *Climate Change 2013: The Physical Science Basis*, T. F. Stocker et al., Eds., Cambridge University Press, 1029–1136.

Dickinson, R. E., 1986: How will climate change? *The Greenhouse Effect, Climatic Change and Ecosystems*, B. Bolin et al., Eds., John Wiley and Sons, SCOPE Report, Vol. 29, 206–270.

Drupp, M., M. Freeman, B. Groom, and F. Nesje, 2015: Discounting disentangled: An expert survey on the determinants of the long-term social discount rate. Centre for Climate Change Economics and Policy Working Paper 195, Grantham Research Institute on Climate Change and the Environment Working Paper 172, 52 pp. [Available online at http://piketty.pse.ens.fr/files/DruppFreeman2015.pdf.]

Epstein, L. G., and S. E. Zin, 1989: Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, **57**, 937–969, doi:10.2307/1913778.

Greenstone, M., E. Kopits, and A. Wolverton, 2013: Developing a social cost of carbon for US regulatory

analysis: A methodology and interpretation. *Rev. Environ. Econ. Policy*, **7**, 23–46, doi:10.1093/reep/res015.

Gregory, J. M., and Coauthors, 2005: A model intercomparison of changes in the Atlantic thermohaline circulation in response to increasing atmospheric CO_2 concentration. *Geophys. Res. Lett.*, **32**, L12703, doi:10.1029/2005GL023209.

Interagency Working Group, 2015: Technical support document: Technical update of the social cost of carbon for regulatory impact analysis—Under Executive Order 12866. United States Government Interagency Working Group on Social Cost of Carbon Tech. Rep., 35 pp.

IPCC, 2013: Summary for policymakers. *Climate Change 2013: The Physical Science Basis*, T. F. Stocker et al., Eds., Cambridge University Press, 3–29.

Meinshausen, M., S. C. B. Raper, and T. M. L. Wigley, 2011: Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6—Part 1: Model description and calibration. *Atmos. Chem. Phys.*, **11**, 1417–1456, doi:10.5194/acp-11-1417-2011.

Nordhaus, W., and P. Sztorc, 2013: DICE 2013R: Introduction and user's manual. 2nd ed. 102 pp. [Available online at www.econ.yale.edu/~nordhaus/homepage/documents/DICE_Manual_100413r1.pdf.]

Senior, C. A., and J. F. B. Mitchell, 2000: The time-dependence of climate sensitivity. *Geophys. Res. Lett.*, **27**, 2685–2688, doi:10.1029/2000GL011373.

Stern, N., 2007: *The Economics of Climate Change: The Stern Review*. Cambridge University Press, 712 pp.

U.N., 2004: World population to 2300. United Nations Publ., 240 pp. [Available online at www.un.org/esa/population/publications/longrange2/WorldPop2300final.pdf.]

Watkiss, P., and C. Hope, 2011: Using the social cost of carbon in regulatory deliberations. *Wiley Interdiscip. Rev.: Climate Change*, **2**, 886–901, doi:10.1002/wcc.140.

Weitzman, M. L., 2012: GHG targets as insurance against catastrophic climate damages. *J. Public Econ. Theory*, **14**, 221–244, doi:10.1111/j.1467-9779.2011.01539.x.