Expected Utility and Catastrophic Risk in a Stochastic Economy-Climate Model*

January 30, 2017

Masako Ikefuji
International College of Liberal Arts, Yamanashi Gakuin University

Roger J. A. Laeven
Department of Quantitative Economics,
University of Amsterdam, CentER and EURANDOM

Jan R. Magnus
Department of Econometrics & Operations Research,
Vrije Universiteit Amsterdam and Tinbergen Institute

Chris Muris
Department of Economics, Simon Fraser University

*We are grateful to Graciela Chichilnisky, John Einmahl, Johan Eyckmans, Reyer Gerlagh, Christian Groth, David Hendry, John Knowles, Sjak Smulders, Peter Wakker, Aart de Zeeuw, and Amos Zemel for their comments and suggestions. This research was funded in part by the Japan Society for the Promotion of Science (JSPS) under grant C-22530177 (Ikefuji), by the Netherlands Organization for Scientific Research (NWO) under grant Vidi-2009 (Laeven) and by the Social Sciences and Humanities Research Council's Insight Development Grant 430-2015-00073 (Muris). E-mail addresses: m-ikefuji@ygu.ac.jp (Ikefuji), R.J.A.Laeven@uva.nl (Laeven), jan@janmagnus.nl (Magnus), cmuris@sfu.ca (Muris).

Corresponding author:
Jan R. Magnus
Department of Econometrics & Operations Research
Vrije Universiteit Amsterdam
de Boelelaan 1105
1081 HV Amsterdam
The Netherlands

Phone: +31 20 598 6010
E-mail: jan@janmagnus.nl

Abstract: We analyze a stochastic dynamic finite-horizon economic model with climate change, in which the social planner faces uncertainty about future climate change and its economic damages. Our model (SSICE) is a simplified version of Nordhaus' deterministic DICE model, but it incorporates, possibly heavy-tailed, stochasticity. We develop a regression-based numerical method for solving a general class of dynamic finite-horizon economy-climate models with potentially heavy-tailed uncertainty and general utility functions. We then apply this method to SSICE and examine the effects of light- and heavy-tailed uncertainty. The results indicate that the effects can be substantial.

JEL Classification: E2; Q5.

Keywords: Economy-climate models; Economy-climate policy; Expected utility; Heavy tails; Uncertainty.

1 Introduction

The current economy-climate debate raises many difficult issues. Only one of these issues is discussed in the current paper, namely the question if and how abatement, consumption, and investment policies are affected by catastrophic risk. Economy-climate policies are typically analyzed using Integrated Assessment Models (IAMs) that describe the complex interplay between climate and the economy. Our paper augments a widely-used deterministic IAM by incorporating (potentially heavy-tailed) risk related to future climate change and its associated economic damage, and analyzes its impact on the policy variables.

Our model is based on Nordhaus' (2008, 2013) dynamic integrated model of climate and the economy (DICE), which has become an important benchmark IAM, not only in the theoretical literature but also serving as a tool for economy-climate policy analysis by the US government. The deterministic version of our model is a simplified version of DICE, and we shall refer to it as SICE (= simplified DICE). The simplifications occur primarily in the specification of the dynamics for carbon dioxide and temperature and are chosen to achieve maximum parsimony. The resulting equations thus contain the bare minimum required to analyze the relevant economy-climate issues. Despite its simplicity the SICE model yields optimal policies that closely resemble those of Nordhaus.

The SICE model serves as our starting point but, like DICE, it is deterministic. To represent uncertainty and motivated by the developments in Manne and Richels (1992), Nordhaus (1994), Roughgarden and Schneider (1999), Kelly and Kolstad (1999), Keller *et al.* (2004), Mastrandrea and Schneider (2004), Leach (2007), Weitzman (2009), and in particular Ackerman *et al.* (2010), we introduce to SICE random shocks featuring potentially heavy-tailed risk. We refer to the resulting model as stochastic SICE (SSICE). We initially focus attention on uncertainty through the damage function and, later, in an extension of this base stochastic model, we shall also account for an uncertain emissions-to-output ratio and uncertainty in technological efficiency.

To solve the stochastic dynamic economy-climate model thus obtained, we embed the associated optimization problem into a general class of stochastic dynamic finite-horizon optimization problems. We next develop a regression-based method for solving such problems. Our solution method is flexible in the sense that it allows for a wide class of utility functions and that it imposes only weak assumptions on the stochasticity, e.g., permitting both light- and heavy-tailed risks, and stochastic parameters.

In the context of SSICE we show formally that heavy-tailed risk is only compatible with *some* utility functions, and in particular that it is not compatible with power utility. To do so, we invoke the general decision-theoretic results of Ikefuji *et al.* (2015) and apply these to the current setting. We propose to use the Pareto utility function to represent preferences in the presence of heavy-tailed risk. This utility function was introduced by Ikefuji *et al.* (2013) and advocated by Cerreia-Vioglio *et al.* (2015). Pareto utility avoids the drawbacks 'near the edges' that standard families of utility functions such as power and exponential utility exhibit, and is particularly suited for heavy-tailed risk analysis.

Our four main findings are as follows. First, the introduction of light-tailed uncertainty through the damage function of SICE leads to a reduction of consumption and an increase of investment in the early periods. Conditional upon the shocks realizing their expected value, namely zero, in the first few periods, we find the opposite pattern in later periods. This applies to both power and Pareto utility and is the result of precautionary savings (Kimball, 1990). Similarly, abatement increases in the early periods, leading to lower concentration levels. This can be viewed as the equivalent of precautionary savings in the climate component of our model. The changes in the optimal policy variables are monotone in the variance of the shock. The changes remain small as long as the shocks take values close to their expectation, that is, in the 'center' of the distribution.

Second, when the light-tailed shocks take larger negative values, the optimal policies are more affected: pronounced differences occur in the optimal policy and state variables at the 'edges', both within and between models. In particular, because a power utility maximizer has a stronger motive to smooth consumption than a Pareto utility maximizer, he or she has a stronger desire to keep up consumption in adverse scenarios at the cost of precautionary action. This effect is the result of a trade-off between maintaining current consumption and taking intensified precautionary action under adverse circumstances. We find that the Pareto utility maximizer tends to favor a larger substitution from current consumption to intensified precautionary action when compared to conventional power utility. This effect is distinct from the regular precautionary savings motive — the intertemporal substitution effect from current certain to future uncertain consumption — which it more (Pareto) or less (power) amplifies. Thus, compared to a Pareto utility maximizer, a power utility social planner consumes too much and abates too little under adverse circumstances.

Third, allowing for heavy-tailed uncertainty making catastrophic risk scenarios more pronounced, our first main finding broadly remains valid under Pareto utility, while our second main finding gets reinforced, with power utility becoming incompatible in this case. Indeed, for a power utility maximizer, the expectation of the intertemporal marginal rate of substitution becomes infinite when considering heavy-tailed uncertainty in the SSICE model.

Fourth, the impact of uncertainty in the damage function dominates the impact of an uncertain emissions-to-output ratio and closely resembles the impact of uncertainty through technological efficiency in the center of the distribution. At the edges, however, when adverse scenarios for technological efficiency realize, optimal abatement is suppressed compared to the adverse scenarios in which economic damages are relatively large.

Although there are many papers on climate policy under uncertainty, the literature on the interplay between climate and the economy under uncertainty is much smaller. The existing IAMs which explicitly include uncertainty can be divided in three classes: (i) stochastic dynamic IAMs with learning, but no consideration of catastrophe; (ii) deterministic IAMs considering catastrophe; and (iii) stochastic dynamic IAMs considering tipping points.

In class (i), Kelly and Kolstad (1999) explore Bayesian learning about the relationship between greenhouse gas levels and global mean temperature changes, analyze when uncertainty is resolved, and show that the expected learning time depends on the variance of the temperature realisations and varies directly with the emission policy. Extensions of Kelly and Kolstad (1999) are provided in Keller *et al.* (2004), Leach (2007), and Traeger (2014). Jensen and Traeger (2014) study the effects of climate sensitivity uncertainty, learning, and temperature stochasticity separately, and find precautionary savings in the presence of the stochasticity of temperature, while Bayesian learning about climate sensitivity raises the abatement rate and hence the optimal carbon tax.

In class (ii), Mastrandrea and Schneider (2004), Ackerman *et al.* (2010), Dietz (2011), Hwang *et al.* (2013), and Gillingham *et al.* (2015) study the implication of catastrophic risks in IAMs. These papers focus on examining the shape of the damage function and the climate sensitivity parameter. We mention in particular the relevant contribution by Ackerman *et al.* (2010), who analyze the impact of parameter uncertainty in the specification of the damage function and/or in the temperature equation on the optimal policies. Their approach consists in first simulating the parameter(s) of interest by drawing from a pre-specified probability distribution, and then deterministically solving DICE for each realization of the parameter(s), thus obtaining a 'distribution' of the optimal policies. This approach provides an assessment of the sensitivity and robustness of the optimal policies to parameter assumptions within the context of a deterministic economy-climate model. Also, Gillingham *et al.* (2015) conduct an extensive Monte Carlo analysis for six IAMs, to analyze how model output responds to model misspecification due to parameter uncertainty, by estimating surface-response functions. The current paper, in contrast, solves a stochastic optimization problem. Our social planner takes potentially heavy-tailed stochasticity in the damage function (and the emissions-to-output ratio and technological efficiency, in extensions of the model) already into account when solving for the optimal policies.

In class (iii), Lemoine and Traeger (2014), Lontzek *et al.* (2015), Cai *et al.* (2015, 2016), and Berger *et al.* (2016) explore how the risk of stochastically uncertain environmental tipping points affects climate policy, using a stochastic IAM based on the DICE model. Berger *et al.* (2016) adopt non-expected utility preferences to accommodate aversion to both risk and ambiguity when analyzing tipping elements in climate change, employing a two-period model in which uncertainty resolves in 2100. The paper by Cai *et al.* (2015) is particularly relevant for us. They extend conventional economy-climate analysis based on deterministic IAMs to allow for a range of stochastic features. In particular, they conduct an extensive analysis of carbon emission policies in a stochastic environment. A key distinction between their work and ours is that they only allow shocks with a bounded probability distribution, thus ruling out the normal or the Student distribution, in order to avoid catastrophic risk scenarios ('tail events'). In contrast, risks with unbounded support and potentially featuring heavy tails, as well as the catastrophic risk scenarios they may induce, are at the heart of our analysis.

Our paper also relates to the literature on numerical methods for dynamic programming and stochastic optimal control. The algorithm that we develop for solving SSICE is inspired by the Least Squares Monte Carlo (LSMC) approach introduced by Longstaff and Schwartz (2001); see also Carriere (1996), Clément *et al.* (2002), and Powell (2011) for further details, including convergence results. LSMC has been successfully applied to a variety of problems in financial economics and operations research; see e.g., Brandt *et al.* (2005), who use LSMC to solve a portfolio choice problem with non-standard preferences, Laeven and Stadje (2014), who solve problems of optimal portfolio choice and indifference valuation under general asset price dynamics and in the presence of model uncertainty using LSMC, and Krätschmer *et al.* (2015), who employ LSMC to analyze model uncertainty in optimal stopping.

The paper is organized as follows. In Section 2 we introduce SICE, an economy-climate model of the DICE type (Nordhaus, 2013). In Section 3 we introduce uncertainty into SICE. In Section 4 we provide a formal description of a general class of stochastic dynamic finite-horizon economy-climate models, allowing for heavy-tailed uncertainty and general utility functions and embedding SSICE as a special case, and develop a regression-based method for solving such models. In Section 5 we show, in the context of our model, that heavy-tailed uncertainty is not compatible with all utility functions, in particular power utility, and propose an alternative utility function: Pareto utility. In Section 6 we present the analysis of our SSICE model and discuss its implications, while some extensions are presented in Section 7. Section 8 concludes.

2 A simplified economy-climate model

We present a simple economy-climate model in the spirit of Nordhaus and Yang (1996) and Nordhaus (2008, 2013). Nordhaus' DICE models have two main versions: 2007 and 2013, which are described in Nordhaus (2008) and Nordhaus (2013), respectively, and in user manuals available from Nordhaus' website. Our model retains all essential features of both versions of DICE, but it contains much fewer equations (primarily because of simplifications in the carbon dioxide concentrations and temperature equations), and is therefore more transparent and numerically easier to handle. While the methods developed in this paper can readily accommodate more complex economy-climate models and their stochastic extensions, we prefer maximum parsimony when analyzing the fundamental question of how catastrophic risk impacts on optimal abatement, consumption, and investment.

Everybody works. In period t, the labor force L_t together with the capital stock K_t generate GDP Y_t through a Cobb-Douglas production function

$$Y_t = A_t K_t^{\gamma} L_t^{1-\gamma} \qquad (0 < \gamma < 1), \tag{1}$$

where A_t represents technological efficiency and γ is the elasticity of capital. Capital is accumulated through

$$K_{t+1} = (1-\delta)K_t + I_t \qquad (0 < \delta < 1), \tag{2}$$

where I_t denotes investment and δ is the depreciation rate of capital. Production generates carbon dioxide (CO_2) emissions E_t:

$$E_t = \sigma_t(1-\mu_t)Y_t, \tag{3}$$

where σ_t denotes the emissions-to-output ratio for CO_2, and μ_t is the abatement fraction for CO_2. Total CO_2 emissions consist of industrial emissions E_t and non-industrial ('land-use') emissions. We denote the latter type by E_t^0 and consider it to be exogenous to our model, as in Nordhaus. The associated CO_2 concentration M_t accumulates through

$$M_{t+1} = (1-\phi)M_t + E_t^0 + E_t \qquad (0 < \phi < 1), \tag{4}$$

where ϕ is the depreciation rate of CO_2 (rate of removal from the atmosphere). Temperature H_t develops according to

$$H_{t+1} = \eta_0 + \eta_1 H_t + \eta_2 \log(M_{t+1}) \qquad (\eta_1 > 0,\ \eta_2 > 0). \tag{5}$$

The temperature and climate systems in (4) and (5) are simplified versions of those in Nordhaus (2013), and this represents the main difference between our model and DICE.

In each period t, the fraction of GDP not spent on abatement or 'damage' is either consumed (C_t) or invested (I_t) along the budget constraint

$$(1 - \omega_t) D_t Y_t = C_t + I_t. \tag{6}$$

The damage function D_t depends only on temperature and satisfies $0 < D_t \leq 1$, where $D_t = 1$ represents the optimal temperature for the economy. Deviations from the optimal temperature cause damage. As in Nordhaus (2013) we specify D_t as

$$D_t = \frac{1}{1 + \xi H_t^2} \qquad (\xi > 0). \tag{7}$$

For very high and very low temperatures D_t approaches zero. The optimal value $D_t = 1$ occurs at $H_t = 0$, the temperature in 1900. Of course, other forms of the damage function are possible; see Stern (2007), Weitzman (2009), and Ackerman *et al.* (2010).

A fraction ω_t of $D_t Y_t$ is spent on abatement, and we specify the abatement cost fraction as

$$\omega_t = \psi_t \mu_t^\theta \qquad (\theta > 1). \tag{8}$$

When μ_t increases then so does ω_t, and a larger fraction of GDP will be spent on abatement. These equations capture the essence of the DICE models. Notice that many theoretical models treat the labor force L_t as a flow variable. Here, however, labor force equals population, which is a stock variable.

The eight equations (1)–(8) imply a condensed system consisting of three dynamic equations (in the state variables K, M, and H) in terms of the policy variables I and μ and exogenous variables and parameters:

$$\begin{cases} K_{t+1} = (1-\delta) K_t + I_t, & (9) \\ H_{t+1} = \eta_0 + \eta_1 H_t + \eta_2 \log(M_{t+1}), & (10) \\ M_{t+1} = (1-\phi) M_t + E_t^0 + \sigma_t (1-\mu_t) A_t K_t^\gamma L_t^{1-\gamma}, & (11) \end{cases}$$

and, through the budget constraint

$$C_t + I_t = \frac{(1 - \psi_t \mu_t^\theta) A_t K_t^\gamma L_t^{1-\gamma}}{1 + \xi H_t^2}, \tag{12}$$

a third policy variable C in terms of I and μ and the state variables (and exogenous variables and parameters).

As in Nordhaus (2013) one period is five years. Period 0 refers to the time interval 2010–2014, period 1 to 2015–2019, and so on. Stock variables are measured at the beginning of the period; for example, K_0 denotes capital in the year 2010. We choose the exogenous variables such that $L_t > 0$, $A_t > 0$, $E_t^0 > 0$, $\sigma_t > 0$, and $0 < \psi_t < 1$. The policy variables must satisfy

$$C_t \geq 0, \quad I_t \geq 0, \quad 0 \leq \mu_t \leq 1. \tag{13}$$

Given a utility function U we define welfare in period t as

$$W_t = L_t U(C_t/L_t). \tag{14}$$

The policy maker has a finite horizon, and maximizes total discounted welfare,

$$W = \sum_{t=0}^{T} \frac{W_t}{(1+\rho)^t} \qquad (0 < \rho < 1), \tag{15}$$

where ρ denotes the discount rate. Letting x denote per capita consumption, the utility function $U(x)$ is assumed to be defined and strictly concave for all $x > 0$. There are many such functions, but a popular choice is

$$U(x) = \frac{x^{1-\alpha} - 1}{1-\alpha} \qquad (\alpha > 0), \tag{16}$$

where α denotes the elasticity of marginal utility of consumption. This is the so-called *power* function. Many authors, including Nordhaus, select this function. In earlier versions of the DICE model, Nordhaus (2008) chooses $\alpha = 2$ in which case $U(x) = 1 - 1/x$. Also popular is $\alpha = 1$; see Kelly and Kolstad (1999) and Stern (2007). We use the value maintained in the 2013 version of the DICE model, namely $\alpha = 1.45$.

Table 1: Comparison of state variables in DICE and SICE models

	K_t/Y_t		M_t/Y_t		H_t	
	DICE	SICE	DICE	SICE	DICE	SICE
2010	2.1246	2.1244	12.8889	12.8876	0.8300	0.8300
2035	2.2017	2.2124	6.8168	6.8446	1.4037	1.4279
2060	2.2887	2.3031	4.2907	4.4201	2.0673	2.1219
2085	2.3627	2.3759	2.8942	3.1746	2.6864	2.8467

The parameter values and initial levels are presented and discussed in Table 4 in Appendix A. Our simple model (hereafter, SICE = simplified DICE) produces optimal values over 25 periods ($t = 0, \dots, 24$). Since the final periods are not representative because of the terminal condition, we only report the first 16 time periods ($t = 0, \dots, 15$) covering 75 years. The results during those 75 years (obtained under power utility, like Nordhaus) are close, although not identical, to those of Nordhaus, as shown in Table 1. A more detailed comparison is provided in Appendix A.

3 Stochastic SICE

We now introduce uncertainty in the SICE model, focussing on the uncertainty about the economic impact of future climate change. Thus we obtain a stylized stochastic integrated assessment model of climate economics, to which we refer as SSICE (= stochastic SICE).

Stochasticity is introduced in SICE by adding two random shocks $u_{1,t}$ and $u_{2,t}$ to the condensed system defined by the three dynamic equations (9)–(11) and the budget constraint (12). More precisely, we adjust (11) and (12) to

$$\begin{cases} M_{t+1} = (1-\phi)M_t + E_t^0 + (1-\mu_t)B_{1,t}u_{1,t}, & (17) \\ C_t + I_t = (1-\psi_t\mu_t^\theta)B_{2,t}u_{2,t}, & (18) \end{cases}$$

respectively, where

$$B_{1,t} = \sigma_t A_t K_t^\gamma L_t^{1-\gamma}, \qquad B_{2,t} = \frac{A_t K_t^\gamma L_t^{1-\gamma}}{1+\xi H_t^2}, \tag{19}$$

and

$$u_{1,t} = e^{-\tau_1^2/2}e^{\tau_1\epsilon_{1,t}}, \qquad u_{2,t} = e^{-\tau_2^2/2}e^{\tau_2\epsilon_{2,t}}. \tag{20}$$

The normalizing constants $e^{-\tau_1^2/2}$ and $e^{-\tau_2^2/2}$ are chosen such that $\mathbb{E}(u_{1,t}) = \mathbb{E}(u_{2,t}) = 1$ when $\epsilon_{1,t} \sim \mathrm{N}(0,1)$ and $\epsilon_{2,t} \sim \mathrm{N}(0,1)$. Special cases include:

- $\tau_1 = 0$, $\tau_2 > 0$: uncertainty through the damage function by modifying D_t in (7) to $\bar{D}_t = D_t u_{2,t}$. This is the case we shall emphasize in particular;
- $\tau_1 > 0$, $\tau_2 = 0$: uncertainty through the emissions-to-output ratio by modifying σ_t in (3) to $\bar{\sigma}_t = \sigma_t u_{1,t}$;
- $\tau_1 = \tau_2 > 0$ and $\epsilon_{1,t} = \epsilon_{2,t}$: uncertainty through technological efficiency by modifying A_t in (1) to $\bar{A}_t = A_t u_{2,t}$. This is the case emphasized in Cai *et al.* (2015).

The general formulation in (17) and (18) thus encompasses a rich spectrum of uncertainties: we can use it to examine how damage and mitigation uncertainty interacts with climate change policies, but also how uncertainty in productivity or in the emissions-to-output ratio interacts with such policies. We shall first focus on damage and mitigation uncertainty through $\epsilon_{2,t}$, but later we shall discuss all cases.

We notice for later reference that both $B_{1,t}$ and $B_{2,t}$ are positive for all t. This follows because the exogenous variables A_t, L_t, and σ_t are assumed to be positive for all t; the parameters δ and ξ satisfy $0 < \delta < 1$ and $\xi > 0$; the initial condition $K_0 > 0$ holds; and $K_t \geq (1-\delta)^t K_0 > 0$ since $I_t \geq 0$. We also notice that consumption is bounded by

$$C_t \leq C_t + I_t = (1 - \psi_t \mu_t^\theta) B_{2,t} u_{2,t} \leq B_{2,t} u_{2,t} = B_{2,t} e^{-\tau_2^2/2} e^{\tau_2 \epsilon_{2,t}}, \tag{21}$$

since $I_t \geq 0$, $\psi_t > 0$, $\mu_t \geq 0$, $B_{2,t} > 0$, and $u_{2,t} > 0$ (with probability one).

We shall consider both light- and heavy-tailed risk. If the $\epsilon_{2,t}$ are independent and identically distributed (iid) and follow a normal distribution $\mathrm{N}(0,1)$, then the moments of $u_{2,t}$ exist, and we have $\mathbb{E}(u_{2,t}) = 1$ and $\mathrm{var}(u_{2,t}) = e^{\tau_2^2} - 1$. Since the distribution of $u_{2,t}$ is heavily skewed, more uncertainty (higher τ_2) implies more probability mass of $u_{2,t}$ close to zero. If, however, we move only one step away from the normal distribution and assume, e.g., that $\epsilon_{2,t}$ follows a Student distribution with *any* (finite) degrees of freedom, then the expectation is infinite (Geweke, 2001). The analysis in, among others, Dietz (2011), Pindyck (2011), Buchholz and Schymura (2012), and Hwang *et al.* (2013) suggests that heavy-tailed risk plays an important role in the economics of climate change.

If $\tau_2 > 0$ then the assumption of iid distributed errors $\epsilon_{2,t}$ is sufficient to generate the possibility of incompatibility between preferences and distributional assumptions, as discussed and proved in Section 5. The algorithm we propose in Section 4 can also handle more sophisticated error assumptions.

4 Optimization problem and solution algorithm

In this section we discuss a class of stochastic dynamic finite-horizon optimization problems to which the SSICE model in Section 3 belongs as a special case, and develop a numerical method to solve such problems. We first introduce some notation, define our general class of optimization problems, and show how it encompasses SSICE as a special case. Then we design a regression-based algorithm to numerically solve optimization problems in this class. The optimization problem that we consider is a challenging one, because of the nonlinearities induced by the economy-climate model, the desired generality of preferences and beliefs, and the aim to accurately capture tail-risk behavior far away from a rapidly evolving steady state. We shall indicate how our solution algorithm deals with each of these challenges.

4.1 SSICE as a stochastic dynamic finite-horizon optimization problem

The social planner in SSICE faces a discrete-time stochastic dynamic finite-horizon programming problem, which consists of maximizing expected total discounted welfare W given in (15) subject to the SSICE model specified in the condensed system of equations (9)–(10) and (17)–(18).

To facilitate the discussion of our optimization problem, we express it using the nomenclature of dynamic optimization; see, e.g., Bertsekas (2005) or Powell (2011). We start by considering a discrete-time setup with control variables, state variables, and stochastic drivers, adapted to an underlying filtered probability space. The time-t state variables are stacked into a vector x_t, the time-t control variables into a vector z_t and the time-t stochastic drivers into a vector ϵ_t. In SSICE, the prime state variables are the capital stock K_t, temperature H_t, and carbon dioxide concentrations M_t, while the prime control variables are consumption C_t and the abatement fraction μ_t. Stochasticity enters SSICE through $\epsilon_t = (\epsilon_{1,t}, \epsilon_{2,t})$.

For given values of (x_t, z_t, ϵ_t), and given the exogenous variables and parameters at time t, all other endogenous variables in the model are supposed to be known at time t. In SSICE, the exogenous variables and parameters are: the initial values of the state variables, (K_0, H_0, M_0), the time-varying exogenous stock variables and parameters $(A_t, L_t, \psi_t, \sigma_t)$, the time-invariant parameters $(\gamma, \delta, \rho, \phi, \xi, \theta, \eta_0, \eta_1, \eta_2)$, and the stochasticity parameters (τ_1, τ_2). The remaining endogenous variables in the model — investment I_t, emissions E_t, the abatement cost fraction ω_t, welfare W_t, output Y_t, and damage D_t — are determined by the prime control and state variables, the exogenous variables and the parameters. For example, the control variable investment, I_t, is obtained from the budget constraint (18) and the state variable GDP, Y_t, follows from the identity (1). Similarly, explicit expressions depending on the prime state and control variables and stochastic drivers are obtained for all other state and control variables that are not contained in x_t or z_t under SSICE. Given the controls, the discrete-time process of state variables is assumed to be a controlled Markov process. The Markov property is essential in our development. For ease of exposition we assume in the present section that the $\{\epsilon_t\}$ are independent over time, but this assumption can be relaxed. (Removing the requirement that the $\{\epsilon_t\}$ are independent over time means that the value function and its approximation at time t introduced below will explicitly depend on the stochasticity vector ϵ_{t-1}.)

In general, the system of prime state variables evolves dynamically according to a sequence of vector functions f_t taking values on the support of x_t:

$$x_{t+1} = f_t(x_t, z_t, \epsilon_t).$$

By allowing f_t to be a time-varying function, we accommodate arbitrary time paths for exogenous variables and parameters. In particular, the function f_t in the SSICE model is given by (9)–(10) and (17).

The decision maker seeks to implement the optimal policy, while satisfying the constraints imposed by the model. The constraints on the time-t control variables z_t are represented by a time-varying set $\mathcal{Z}_t(x_t, \epsilon_t)$ that depends in particular on the current value of the state vector x_t and the stochastic driver ϵ_t. Maximization is then over

$$z_t \in \mathcal{Z}_t(x_t, \epsilon_t). \tag{22}$$

For the SSICE model this set of constraints specializes to:

$$0 \le C_t \le \frac{(1 - \psi_t \mu_t^\theta)\, A_t K_t^\gamma L_t^{1-\gamma}\, e^{-\tau_2^2/2}\, e^{\tau_2 \epsilon_{2,t}}}{1 + \xi H_t^2}$$

and

$$0 \le \mu_t \le 1.$$

The decision maker's objective is to maximize his/her evaluation of a stream of payoffs (or rewards) by optimally selecting the control variables. Denote by V_s $(0 \le s \le T)$ the maximum of the evaluation of the payoff stream collected in periods s through T, given all the information available at time $s-1$ and subject to the constraints in (22):

$$V_s(x_s) = \max_{z_s, \ldots, z_T} \mathbb{E}_{s-1}\left[\sum_{t=s}^{T} g_t(z_t)\right] \tag{23}$$

subject to

$$\begin{cases} z_t & \in \mathcal{Z}_t(x_t, \epsilon_t) \quad (s \le t \le T), \\ x_{t+1} & = f_t(x_t, z_t, \epsilon_t) \quad (s \le t \le T-1), \end{cases}$$

where g_t is the decision maker's time-t specific objective function and $\mathbb{E}_s$ is short-hand notation for the conditional expectation with respect to the filtration at time s. The function V_s is referred to as the value function. The corresponding Bellman equation is given by

$$V_t(x_t) = \max_{z_t \in \mathcal{Z}_t(x_t, \epsilon_t)} \mathbb{E}_{t-1} \left[g_t(z_t) + \beta V_{t+1}(f_t(x_t, z_t, \epsilon_t)) \right], \tag{24}$$

where β is a discount factor ($0 < \beta < 1$). The time-t objective function $g_t(z_t)$ in the SSICE model is given by

$$g_t(z_t) = \frac{L_t\, U(C_t/L_t)}{(1+\rho)^t},$$

where U is the utility function. The discount factor in SSICE is equal to $\beta = 1/(1+\rho)$.

If the value function in period $t+1$ is known, then Equation (24) is a static optimization problem in the time-t control variables z_t. The connection between the value function at time t and the value function at time $t+1$, as stipulated by the Bellman equation, allows the decision maker to maximize his/her evaluation recursively by backward induction. Indeed, because all variables including the realizations of stochasticity are observed in each period, the decision maker first determines the optimal control variables in the final period, depending on the other variables and parameters in the model at that time. Then the decision maker maximizes the sum of that part of the evaluation that pertains to time $T-1$ and the discounted future value function, thus proceeding backwards in time.

4.2 Solution algorithm: generic description

We will solve the Bellman equation numerically. Our approach to the computation of the optimal policies is inspired by the Least Squares Monte Carlo (LSMC) approach introduced by Longstaff and Schwartz (2001) in the context of optimal stopping for American-style derivatives and adapted here to our discrete-time dynamic stochastic finite-horizon optimization problem; see also Carriere (1996) and Tsitsiklis and Van Roy (1999).

While it may seem natural to consider all potential future paths of the variables in our model when conducting optimization, this readily becomes ineffective in multiple dimensions and over longer time spans, which is the situation we face in our application. We therefore propose a method based on Monte Carlo where we simulate a set of future paths of the state variables and stochastic drivers, and then invoke regression to obtain estimates of the value function in a recursive fashion. By relying on forward-simulated paths, our method is relatively efficient. Moreover, because of the use of regression methods, the method does not require nested simulation, and hence is computationally fast.

We start in the final period T where the value function is given by

$$V_T(x_T) = \max_{z_T \in \mathcal{Z}_T(x_T, \epsilon_T)} \mathbb{E}_{T-1}[g_T(z_T)]. \tag{25}$$

Indeed, for our finite-horizon model, the payoff in periods after time T is equal to zero. This implies that $V_{T+1}(x_{T+1}) = 0$, so that the Bellman equation in (24) simplifies to the Bellman equation in (25) at time T.

At time T (and similarly for earlier time periods), our algorithm then consists of four steps, as follows.

(a) First we use a random number generator to draw R values (x_T^r, ϵ_T^r) for $r = 1, \ldots, R$. For SSICE, ϵ_T^r is drawn from a probability distribution pre-specified in the model, while the value of the state vector x_T^r is drawn from a uniform distribution with a wide support. This support is centered at the optimal value of the state vector in the deterministic version of the model (SICE). The two (multivariate) draws are independent.

(b) Next, for each r, we compute the deterministic quantity v_T^r as the maximum value of the period-T objective function given the r-th draw (x_T^r, ϵ_T^r). This specific optimization problem is typically straightforward at time T. For example, in the SSICE model, consumption is set equal to the available budget in the final period, and abatement is set to zero.

(c) We then use regression to approximate the function $V_T(x_T)$. To obtain the approximation, we assume that there exists a set of basis functions $\phi_j(x_T)$ and coefficients $\beta_{j,T}$ $(j = 0, 1, 2, \ldots)$ such that

$$V_T(x_T) = \sum_{j=0}^{\infty} \beta_{j,T}\phi_j(x_T), \quad \text{and} \quad V_T(x_T) \approx \sum_{j=0}^{J} \beta_{j,T}\phi_j(x_T), \quad J \in \mathbb{N}_{>0},$$

can serve as an approximation, and, for each r, we decompose the deterministic maximum v_T^r into the sum of this approximation and an (r)-specific disturbance $\nu_{r,T}$, that is,

$$v_T^r = \sum_{j=0}^{J} \beta_{j,T}\phi_j(x_T^r) + \nu_{r,T}.$$

Note that setting $\phi_0(x_T) = 1$ corresponds to including a constant term $\beta_{0,T}$ in the approximation.

(d) Finally, we obtain least-squares estimates of the coefficients in this approximation, which we denote by $\hat{\beta}_{j,T}$ ($j = 0, \ldots, J$), and we define

$$\widehat{V}_T(x_T) = \sum_{j=0}^{J} \hat{\beta}_{j,T} \phi_j(x_T)$$

as our approximation to the value function at time T.

Now consider period $T-1$. The corresponding Bellman equation is

$$V_{T-1}(x_{T-1}) = \max_{z_{T-1} \in \mathcal{Z}_{T-1}(x_{T-1}, \epsilon_{T-1})} \mathbb{E}_{T-2} \left[g_{T-1}(z_{T-1}) + \beta V_T(f_{T-1}(x_{T-1}, z_{T-1}, \epsilon_{T-1})) \right].$$

The algorithm now proceeds as above in four steps: (a) Generate draws $(x^r_{T-1}, \epsilon^r_{T-1})$ for $r = 1, \ldots, R$; (b) Given the r-th draw $(x^r_{T-1}, \epsilon^r_{T-1})$, compute the deterministic maximum v^r_{T-1}, using the approximation $\widehat{V}_T$ obtained above; (c) Obtain estimates for the coefficients $\beta_{j,T-1}$ in

$$v^r_{T-1} = \sum_{j=0}^{J} \beta_{j,T-1} \phi_j(x^r_{T-1}) + \nu_{r,T-1};$$

and (d) Define the value function approximation $\widehat{V}_{T-1}(x_{T-1})$ to the value function at time $T-1$ as

$$\widehat{V}_{T-1}(x_{T-1}) = \sum_{j=0}^{J} \hat{\beta}_{j,T-1} \phi_j(x_{T-1}).$$

Next, we approximate the value function in period $T-2$, and so on. In this way we define, recursively, the value function $\widehat{V}_t$ for all the time periods $t = 0, \ldots, T$ in the model. We thus obtain a flexible least-squares Monte-Carlo-based approach, which accommodates general preferences and beliefs, is easy to implement, and is effective and efficient.

Partial convergence results for Least Squares Monte Carlo in the context of optimal stopping and American option pricing are provided by Longstaff and Schwartz (2001); see also Tsitsiklis and Van Roy (1999). These results are significantly expanded by Clément *et al.* (2002); see also Egloff (2005) and Egloff *et al.* (2007). Their formal results can be adapted to our discrete-time optimal control setting, and this allows us to conclude that the regression-based approximations to the optimal control variables resulting from our approach converge to the optimal control variables as the number of simulations and the number of basis functions (in this order) tend to infinity. The proof is somewhat tedious but conceptually straightforward, and proceeds by showing first that the regression estimates converge using standard asymptotic regression theory, and next (more tedious) that the error propagation resulting from the backward induction procedure vanishes asymptotically.

4.3 Some practical aspects of the algorithm

The previous subsection provides a generic description of a numerically efficient algorithm, which computes the solution to a class of discrete-time stochastic dynamic finite-horizon optimization problems. In our application of this algorithm to the SSICE model, our goal is to accurately capture the nonlinear behavior of the model as well as its tail risk behavior potentially far away from a rapidly evolving steady state. For this reason, the support from which we generate values for the state variables must be sufficiently wide. In addition, we need a flexible approximation to V_t over this wide support. We now describe some further details specific to the implementation of our algorithm.

The support. We need to specify the support from which we draw x_t^r. This support must be wide enough to capture optimal policies away from the steady state, because we are specifically interested in optimal policies in the presence of large negative shocks, i.e., under catastrophic risk, and we want our approximation to the optimal policies to be accurate in such scenarios. Let x_t^* denote the state vector under the optimal solution to the deterministic version of the model (SICE). We draw x_t^r from a uniform distribution with support $[0.6x_t^*, 1.5x_t^*]$. Considering even wider supports leads to deteriorating numerical stability. The first period is of particular importance, as we will investigate in detail the distribution of the optimal policies under large negative shocks in that period. In period 0, capital is equal to $K_0 = 135$. The specified support for K_1 is now given by $[97, 242]$. Even under very large negative shocks in period 0, the specified lower bound of 97 is never binding for the optimal choice of K_1.

Number of draws. We must also specify the number of simulation draws R to be drawn in every time period (each time period consists of five years). The value of R was determined by trial and error. We started with $R = 1{,}000$ simulations per period, and then checked whether the solution is sensitive to increases in R by steps of 1,000. After $R = 5{,}000$, the change in optimal consumption is less than 0.01. We then conservatively set $R = 10{,}000$. Such a large value for R is feasible because our regression-based approach avoids nested simulation. This is not only useful to capture tail risk behavior but also to accommodate general preferences and beliefs.

Basis functions. Then we must specify the choice and number of basis functions in the approximation to the value function. To guide our choice described below, we inspect the fit of our approximating model. The approximating model is estimated by a regression of the value function computed at $V_t(x_t^r)$ on the simulated values of x_t^r at which it is computed. If the residuals of that regression do not vary systematically with x_t^r, then the approximating model provides an adequate description of the value function. Visual inspection of residual plots and model specification tests were used to determine appropriate candidate models. The value function is approximately separable in the state variables, so that we can express the approximate value function as

$$\widehat{V}_t(x_t) = \widehat{V}_{t,k}(K_t) + \widehat{V}_{t,m}(M_t) + \widehat{V}_{t,h}(H_t).$$

This separability reduces one nonparametric regression problem with three continuous variables to three nonparametric regression problems with one continuous variable. We found that natural splines of degree 5 (for K) and 3 (M, H) provide a good approximation to the value function. Other choices such as Chebyshev polynomials tend to perform worse when considering wide supports for the state variables. A more flexible approximation does not improve the fit, but could affect the stability of the solution. The wide support that we consider for the state variables implies that we rarely have to extrapolate outside the domain of the observed values of (K, M, H). In those few cases, the natural splines extrapolate linearly from the lower and upper bound of the supports discussed above.

Code and testing details. The stochastic dynamic optimization problem is solved over 25 periods, and we report for periods 0 through 15, capturing 75 years. The code is written in `R` and is available from the authors upon request. It was tested with `R` version 3.3.1, on a desktop computer with Core i7-2600 architecture running Ubuntu 16.04.

5 Compatibility of preferences and stochasticity

Considerable care is required when combining the expected utility paradigm with distributional assumptions, a fact known since Bernoulli (1738) and Menger (1934). The numerical methods developed in Section 4 are valid, in principle, for general expected utility preferences, but this is only true if these preferences are compatible with the assumed stochasticity. If not, then expected utility or expected marginal utility can become infinite, a situation which we wish to avoid. Hence, if only weak assumptions on the stochasticity are imposed, then some compatibility conditions are required to ensure that our model's stochastic optimization problem is well posed. In fact, we shall place no restrictions on the stochasticity and allow for arbitrarily heavy-tailed risks. Not all families of utility functions are then compatible and this raises the question: which families of utility functions are and which are not compatible with arbitrarily heavy-tailed risks? To answer this question we invoke the general decision-theoretic results of Ikefuji *et al.* (2015) and apply these to SSICE, using backward induction.

We know from Section 3 that $B_{2,t} = A_t K_t^{\gamma} L_t^{1-\gamma}/(1+\xi H_t^2) > 0$, and hence, using (21), that

$$0 < C_t \leq B_{2,t} e^{-\tau_2^2/2} e^{\tau_2 \epsilon_{2,t}}. \tag{26}$$

Now, A_t and L_t are exogenous, and K_t and H_t are deterministic given all information at time $t-1$, since K_t depends on K_{t-1} and I_{t-1}, while H_t depends on K_{t-1}, H_{t-1}, M_{t-1}, μ_{t-1}, $\epsilon_{1,t-1}$, and exogenous variables. Hence, $B_{2,t}$ is deterministic given all information at time $t-1$.

Since the social planner in our setup has time-additive expected utility preferences, the inequality (26) implies that inequality (2) in Ikefuji *et al.* (2015) would be satisfied if C_t were the only choice variable. In fact, there are three choice variables: I_t, μ_t, and C_t. It is obvious, however, that in the final period zero abatement and zero investment are optimal: $I_T^* = \mu_T^* = 0$. Hence, in the final period there is only one choice variable, namely C_T, and hence the desired inequality is satisfied at time T.

We can now invoke Proposition 5.2 of Ikefuji *et al.* (2015), apply it to the final two periods in our setup, and conclude that if the probability distribution of $\epsilon_{2,T}$ is heavy-tailed to the left, then expected marginal utility (or the expected intertemporal marginal rate of substitution) pertaining to time T is infinite whenever the utility function belongs to the power family. Thus, if we move only slightly away from normality and allow $\epsilon_{2,T}$ to follow, e.g., a Student distribution with any degrees of freedom, then expected marginal utility explodes under power utility. A similar result is true for expected utility, but we shall not expand on this.

The fragility of expected power utility to heavy-tailed distributional assumptions was noted earlier, e.g. by Geweke (2001). More recently, in the context of catastrophic climate change, Weitzman (2009) pointed out that not only expected utility but also expected marginal utility, and hence the intertemporal marginal rate of substitution, may become infinite with power utility and heavy-tailed log consumption, inducing unacceptable conclusions in cost-benefit analyses.

Because of the incompatibility of power utility we need to look for a different family of utility functions to represent preferences over heavy-tailed risks in SSICE. The Pareto family, introduced by Ikefuji *et al.* (2013) and given by

$$U(x) = 1 - \left(1 + \frac{x}{\lambda}\right)^{-k} \qquad (k > 0,\ \lambda > 0), \tag{27}$$

enjoys a combinations of appealing properties especially relevant in heavy-tailed risk analysis. Let

$$\mathrm{ARA}(x) = -\frac{U''(x)}{U'(x)}, \qquad \mathrm{RRA}(x) = -\frac{xU''(x)}{U'(x)} \tag{28}$$

denote the local indexes of absolute and relative risk aversion. Under power utility, often referred to as constant RRA utility, we have $\mathrm{ARA}(0) = \infty$, $0 < \mathrm{RRA}(0) < \infty$, and $\mathrm{RRA}(x)$ is bounded (in fact constant). Under exponential utility, given by $U(x) = 1 - e^{x/\lambda}$ $(\lambda > 0)$ and often referred to as constant ARA utility, we have $0 < \mathrm{ARA}(0) < \infty$, $\mathrm{RRA}(0) = 0$, and $\mathrm{RRA}(x)$ is unbounded for large values of x. By contrast, under Pareto utility,

$$\mathrm{ARA}(x) = \frac{k+1}{x+\lambda}, \qquad \mathrm{RRA}(x) = \frac{x(k+1)}{x+\lambda}, \tag{29}$$

so that $0 < \mathrm{ARA}(0) < \infty$ and $\mathrm{ARA}(x)$ is non-negative decreasing and convex, while $\mathrm{RRA}(0) = 0$ and $\mathrm{RRA}(x)$ is increasing concave and bounded between 0 and $k + 1$. Notice that the property that $\mathrm{RRA}(0) = 0$ does not imply risk-neutrality at $x = 0$, since $\mathrm{ARA}(0) = (k + 1)/\lambda > 0$.

The family of Pareto utility functions is parsimonious yet flexible. Pareto utility avoids the drawbacks that the popular families of power (constant RRA) and exponential (constant ARA) utility exhibit 'near the edges'. This includes both the extreme behavior of power utility near the origin, where ARA becomes infinite, and the extreme behavior of exponential utility for large x, where RRA increases without bound. In view of Propositions 5.1–5.3 in Ikefuji *et al.* (2015), Pareto utility is particularly appropriate for heavy-tailed risk analysis. It ensures finiteness of both expected utility and expected marginal utility, irrespective of distributional assumptions; see also the discussion in Cerreia-Vioglio *et al.* (2015).

In particular, Proposition 5.2 (or 5.3) of Ikefuji *et al.* (2015) implies that, under Pareto utility, expected marginal utility remains finite for any t. Hence, the expected intertemporal marginal rate of substitution that trades off current and future consumption remains finite under Pareto utility. Because of the boundedness of Pareto utility (cf. Proposition 5.1 of Ikefuji *et al.*, 2015), we see that expected utility also remains finite under Pareto utility, irrespective of distributional assumptions. We conclude that the Pareto family represents a suitable choice of utility functions when analyzing heavy-tailed risk in SSICE.

6 Main findings

We now have developed a stochastic economy-climate framework and a solution method, and this permits a variety of applications and analyses, including exploring fundamental questions such as whether the social planner would abate and invest more or less, and how much, in the presence of uncertainty or under the manifestation of catastrophic risk.

When interpreting the results, it is important to understand whether the results obtained from IAMs have a normative or a descriptive meaning. While climate models are typically interpreted descriptively, the use of optimization suggests a normative perspective. Gordon *et al.* (1987) noted, however, that the results derived from IAMs provide an approximation to an economically efficient market equilibrium, and therefore don't have a normative meaning *per se*.

In the current section, we numerically solve and analyze the base SSICE model with a simple iid specification of stochasticity. In Section 6.1 we discuss the parameter choices pertaining to the preferences (i.e., the social planner's utility function) and beliefs (i.e., the probability distribution of the shocks), and in Section 6.2 we compare the results of the deterministic SICE model to those of the DICE model.

Then we introduce stochasticity. In Section 6.3 we analyze the effects of uncertainty on the optimal abatement, consumption, and investment policies, focusing on optimal policies along the expected trajectory of the shocks, i.e., in the 'center' of the probability distribution. In Section 6.4 we explore the effects at the 'edges' of the probability distribution, that is, we ask what happens to the optimal policies upon the manifestation of a large negative shock. In Section 6.5 we analyze the effect of heavy-tailed versus light-tailed uncertainty.

In Section 7 we shall consider extensions to the base SSICE model, allowing in particular for uncertainty in the emissions-to-output ratio and uncertainty through technological efficiency.

6.1 Setting and base parameters

We shall consider both light-tailed and heavy-tailed probability distributions for the error terms $\epsilon_{1,t}$ and $\epsilon_{2,t}$. Following our discussion in Section 3, we consider both a normal distribution (light tails) and a Student distribution (heavy tails). Under normality, the damage function $\bar{D}_t = D_t u_{2,t}$ has a finite expectation. Under a Student distribution, which may be interpreted as the posterior predictive distribution of a normal distribution with uncertain standard deviation, its expectation is infinite.

We need to specify adequate values for the uncertainty parameters τ_1 and τ_2 and for the number of degrees of freedom of the Student distribution. Suppose $\epsilon_{1,t} = \epsilon_{2,t} \equiv \epsilon_t$ and $\tau_1 = \tau_2 \equiv \tau$. The stochasticity as generated by ϵ_t captures uncertainty about technological efficiency affecting GDP. Historical variation in GDP may therefore serve as a sensible proxy for τ. Barro (2009) calibrates the standard deviation of log GDP to a value of 0.02 on an annual basis, which corresponds to about 0.045 over a five-year horizon. We will therefore consider values of τ_1 and τ_2 in the range of $0.03 \leq \tau_1, \tau_2 \leq 0.06$. Throughout this section we focus on uncertainty through the damage function ($\tau_1 = 0$, $\tau_2 > 0$) and assume the errors to be iid. Other assumptions on τ_1 and τ_2, in particular ($\tau_1 > 0,\ \tau_2 = 0$) and ($\tau_1 = \tau_2 > 0,\ \epsilon_{1,t} = \epsilon_{2,t}$), are postponed to Section 7.

We also need to consider the question of whether or not the stochasticity is light- or heavy-tailed. A (partial) answer to this question is contained in Ursúa (2010), who claims that the growth rate of GDP features heavy tails. We choose the number of degrees of freedom of the Student distribution equal to 10. Our parameter choices then ensure that the summary statistics, including the 'tail index', of output growth rates generated by our model resemble those observed in empirical data.

Finally, we need to specify values for the parameters of the utility functions. In the 2013 version of the DICE model, Nordhaus uses a power utility function with constant relative risk aversion coefficient equal to $\alpha = 1.45$. For comparability, we choose the same value of α when we employ power utility. When we consider the Pareto utility function, we wish to mimic power utility along the expected trajectory of $\epsilon_{2,t}$, i.e., in the center of the probability distribution. With this objective in mind we calibrate the parameters of the Pareto utility function to $\kappa = 1.322$ and $\lambda = 0.0108$.

6.2 SICE versus DICE

In a nonstochastic world we find that the optimal policy and state variables under SICE with power utility closely match their counterparts under DICE, in agreement with our remarks in Section 3. This is true in particular for the variables pertaining to the economy part of the two models (with a maximum difference of 0.6% over the periods that we consider). The main differences between SICE and DICE are contained in the climate part of the two models, but the discrepancies remain small.

When we move from power utility to Pareto utility, we find that the optimal policy and state variables under SICE with Pareto utility match their counterparts under SICE with power utility quite closely, and that this applies to both the economy and climate parts of the SICE model.

Appendix A provides more details. It contains, in particular, results for the control variables in Table 5 and for the state variables in Table 6.

6.3 Light tails in the center

We now introduce stochasticity and consider the SSICE model with iid normally distributed errors $\epsilon_{2,t}$ (i.e., light tails), for different values of the degree of uncertainty τ_2, and under both power and Pareto utility. (Recall that we assume $\tau_1 = 0$ in this section.)

We focus on the 'center' of the distribution by considering shocks along the expected trajectory of $\epsilon_{2,t}$. Specifically, the results reported here are derived by solving for the optimal initial ($t = 0$) policies under uncertainty, and then computing the optimal policies over the following periods 1 to 15 under uncertainty, by assuming that the realized shocks in the previous periods are equal to zero.

Table 2: SSICE with normal errors — power versus Pareto

	power			Pareto		
$t\backslash\tau_2$	0.00	0.03	0.06	0.00	0.03	0.06
Consumption C_t						
2010	46.89	46.75	46.36	44.31	44.08	43.91
2035	104.42	104.05	103.81	106.10	106.01	105.48
2060	189.58	187.64	186.02	191.24	189.53	188.02
2085	301.63	304.19	305.68	301.81	302.47	302.77
Investment I_t						
2010	16.53	16.63	16.93	19.09	19.30	19.37
2035	34.74	34.35	34.06	36.24	36.56	37.20
2060	61.61	61.93	61.40	59.88	60.14	59.69
2085	96.83	94.02	92.15	90.15	89.90	89.61
Abatement μ_t						
2010	0.1523	0.1568	0.1605	0.1642	0.1690	0.1738
2035	0.2435	0.2401	0.2370	0.2371	0.2334	0.2300
2060	0.3277	0.3338	0.3400	0.3004	0.3074	0.3143
2085	0.3728	0.3652	0.3551	0.3383	0.3283	0.3176

The three panels in Table 2 present the results for optimal consumption, investment, and abatement, respectively. Our benchmark is $\tau_2 = 0$, which is the case without uncertainty, that is, SICE. The introduction of light-tailed uncertainty to SICE leads to a reduction of consumption in the early periods. This effect is caused by precautionary savings (Kimball, 1990). Conditional upon the shocks realizing their expected value, that is zero, in the first few periods, we find the opposite pattern in later periods. The pattern reverses around period 15.

In the initial periods, the social planner chooses higher levels of abatement under uncertainty for both power and Pareto utility. This corresponds to choosing lower levels of concentration, and because concentration has a negative propagation effect in our model, this behavior is consistent with precautionary savings. When faced with uncertainty, the social planner also chooses higher levels of investment in the initial periods, a direct result of the precautionary savings motive. Our finding that consumption decreases and initial investment and abatement increase in the presence of uncertainty is consistent with the results in Cai *et al.* (2015) in a more complex stochastic IAM.

The corresponding results for capital K_t, carbon concentration M_t, and temperature H_t are presented in Table 7 in Appendix B. The patterns in these tables are consistent with the results of the control variables.

Overall, the effect of uncertainty on the optimal policies is relatively small when considering a social planner at the center of the probability distribution. Indeed, we find reasonably small changes in the optimal policy variables as long as the shocks take values along their expected trajectory. The changes in optimal control and state variables are virtually always 'monotone' in the variance of the shock as represented by τ_2.

6.4 Light tails at the edges

In the previous subsection we evaluated the effect of uncertainty on the optimal policies in the center of the distribution. Now we analyze the optimal policies at the 'edges', under the manifestation of catastrophic risk (that is, tail events).

Figures 1 and 2 present optimal consumption C_t and optimal abatement μ_t as a function of $\epsilon_{2,t}$ at time $t = 0$, for both the power and Pareto SSICE models. Considering SSICE against the benchmark given by SICE but now allowing the light-tailed shocks to take large negative values, we find that the optimal policy variables are more affected. Towards the edges we observe pronounced differences in the optimal policy variables, both within and between the SSICE models.

The lines in the figures are labeled such that 1, 2, and 3 refer to $\tau_2 = 0.00$, 0.03, and 0.06, respectively; and a and b refer to power and Pareto utility, respectively. As expected, optimal policy derived under certainty — lines

Figure 1: Consumption C_0: SSICE with normal errors and $\tau_2 = 0.00$, 0.03, and 0.06 — power versus Pareto utility

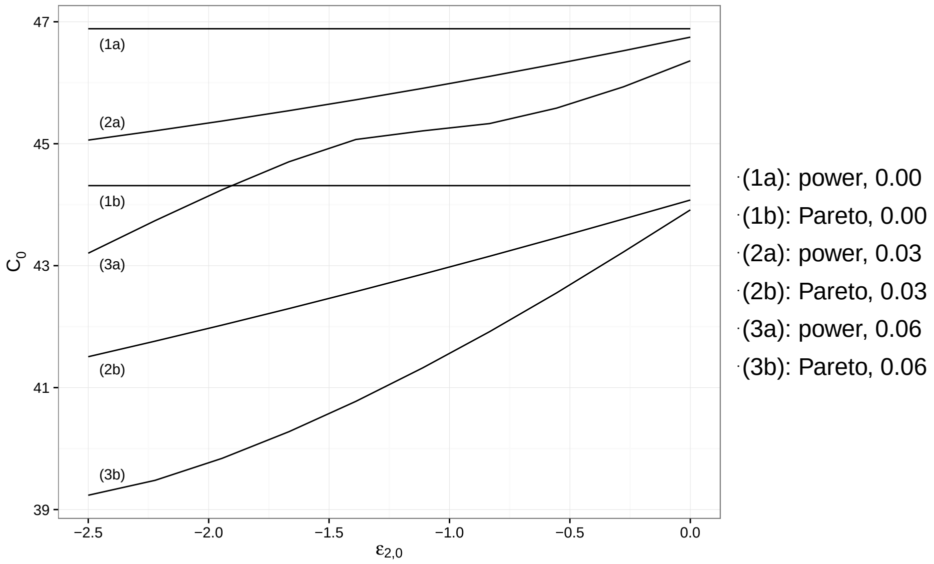

Figure 2: Abatement μ_0: SSICE with normal errors and $\tau_2 = 0.00$, 0.03, and 0.06 — power versus Pareto utility

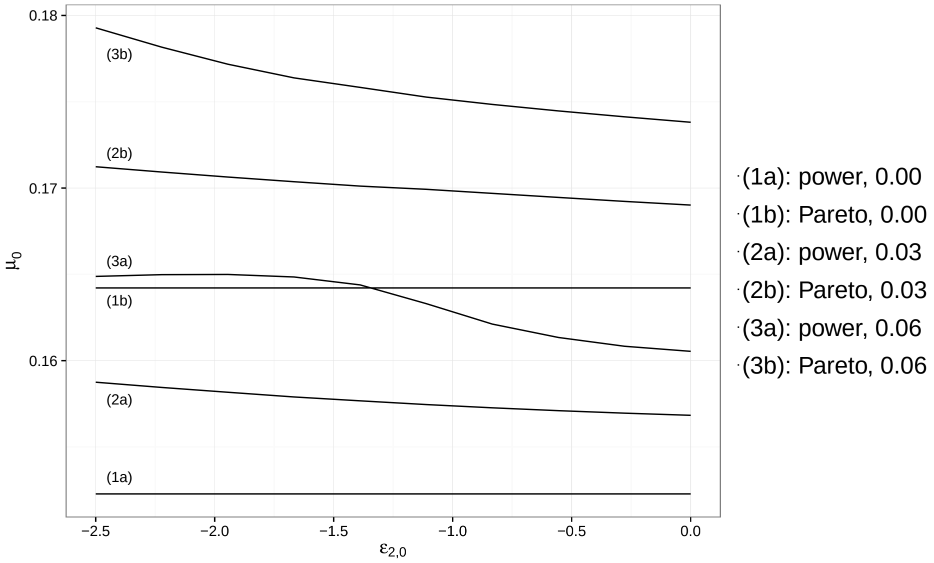

(1a) and (1b) — does not respond to negative shocks. Further, there is a clear ordering in (1a), (2a), (3a), and (1b), (2b), (3b), which tells us that optimal policy at the edges is monotonic in τ_2.

A power utility maximizer has a stronger motive to smooth consumption compared to a Pareto utility maximizer. Under adverse circumstances, the power utility maximizer keeps consumption at a substantial level, but this comes at the cost of lower abatement. As a result, the left tail of the distribution of optimal consumption appears to be lighter under power than under Pareto utility. While the presence of uncertainty increases abatement in the initial periods (as found in Section 6.3), a power utility maximizer puts (excessively) large emphasis on keeping up consumption in adverse circumstances, having a downward pressure on abatement. A Pareto utility maximizer in adverse circumstances will consume less and abate more than a power utility maximizer.

6.5 Heavy tails

Heavy-tailed risk is represented by a Student-t distribution. The random shock $\epsilon_{2,t}$ is not $\mathrm{N}(0, 1)$ anymore but rather follows a t-distribution with 10 degrees of freedom, so that $\mathrm{var}(\epsilon_{2,t}) = 1.25$. Power utility is not compatible with heavy-tailed risk: its expected intertemporal marginal rate of substitution trading off current and future uncertain consumption is infinite. Hence, we only consider Pareto utility.

The three panels in Table 3 report optimal values in the SSICE model under Pareto utility for consumption, investment, and abatement, both for light- and heavy-tailed uncertainty, and for different values of τ_2. In the center of the distribution, the changes are small when we compare the impact of heavy-tailed versus light-tailed uncertainty. As in Sections 6.3 and 6.4, we observe precautionary savings from initial consumption to investment, we find reasonably small changes in the optimal policy variables as long as the shocks take values close to or equal to their expectation (in the center of the distribution), and we see that the changes are 'monotone' in the variance of the shock.

We also report results at the 'edges'; see Figure 3. The changes in the optimal policy variables now become more pronounced, both within and between the models with light and heavy tails. When large negative shocks occur, we find that the social planner chooses lower values of consumption under heavy-tailed risk than under light-tailed risk, favoring spending on precautionary actions.

In summary, under heavy tails the main findings of Sections 6.3 broadly remain valid and those of Section 6.4 are reinforced.

Table 3: SSICE under Pareto utility — light- versus heavy-tailed

	$\tau_2 = 0.03$		$\tau_2 = 0.06$	
	light	heavy	light	heavy
Consumption C_t				
2010	44.08	44.21	43.91	44.11
2035	106.01	105.79	105.48	105.13
2060	189.53	192.06	188.02	192.25
2085	302.47	301.70	302.77	301.47
Investment I_t				
2010	19.30	19.17	19.37	19.19
2035	36.56	36.67	37.20	37.18
2060	60.14	60.18	59.69	60.51
2085	89.90	90.27	89.61	90.25
Abatement μ_t				
2010	0.1690	0.1605	0.1738	0.1567
2035	0.2334	0.2394	0.2300	0.2412
2060	0.3074	0.2966	0.3143	0.2923
2085	0.3283	0.3320	0.3176	0.3260

Figure 3: Consumption C_0: SSICE under Pareto utility — normal versus Student errors with $\tau_2 = 0.00$, 0.03, and 0.06

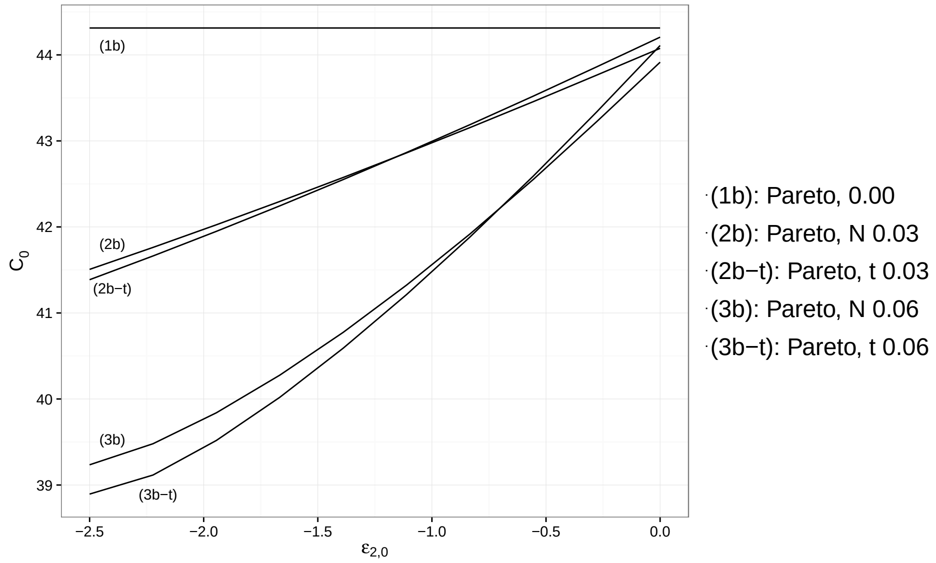

7 Extensions

We generalize the base SSICE model in two directions. First, we allow for an uncertain emissions-to-output ratio. Next, we allow for uncertainty in technological efficiency.

7.1 Emissions-to-output uncertainty

We suppose that $\tau_1 > 0$ and $\tau_2 = 0$, so that uncertainty enters SSICE only through the emissions-to-output ratio σ_t (see (17)), and not through the damage function and the budget constraint (see (18)), as previously. We analyze how our three main findings in Sections 6.3–6.5 are affected under this alternative SSICE model.

We consider first the optimal policies under iid normally distributed errors $\epsilon_{1,t}$, where we restrict our attention to the center of the distribution by considering realizations of $\epsilon_{1,t}$ along the expected trajectory, as in Section 6.3. All three policy variables are now insensitive to the presence of uncertainty along the expected trajectory: the impact on the optimal policies of uncertainty on the emissions-to-output ratio appears to be negligible in the center of the distribution. The reason is that the budget constraint is not affected by uncertainty in the emissions-to-output ratio. Thus, in the center, the effect of emissions-to-output uncertainty is dominated by the effect of uncertainty on the damage function analyzed previously. (Detailed results are available upon request.)

Next, considering the manifestation of tail events analogous to Section 6.4, we find an interesting pattern: while optimal consumption remains insensitive to uncertainty in the emissions-to-output ratio, also under tail scenarios, optimal abatement decreases (increases) when $\epsilon_{1,0}$ takes large negative (positive) values. This is illustrated in Figure 4. This pattern can be explained by the fact that, in the model, abatement directly 'acts upon' the emissions-to-output ratio, while the latter does not appear in the budget constraint, contrary to what happens in Section 6.4. Note also that the scenario in which $\epsilon_{1,0}$ takes large negative values is in fact a very prosperous (rather than adverse) scenario in which emissions are relatively low compared to output, thus facilitating lower abatement.

Finally, we analyze the introduction of heavy-tailed risk attached to the emissions-to-output ratio, using the same distributional assumptions as in Section 6.5. In this setting, the previous two findings are reconfirmed: insensitivity of the optimal policies along the expected trajectory and decreasing (increasing) optimal abatement in the extent of a negative (positive) shock in $\epsilon_{1,0}$.

Figure 4: Abatement μ_0: SSICE-extension-I with normal errors and $\tau_1 = 0.00$, 0.03, and 0.06 — power versus Pareto utility

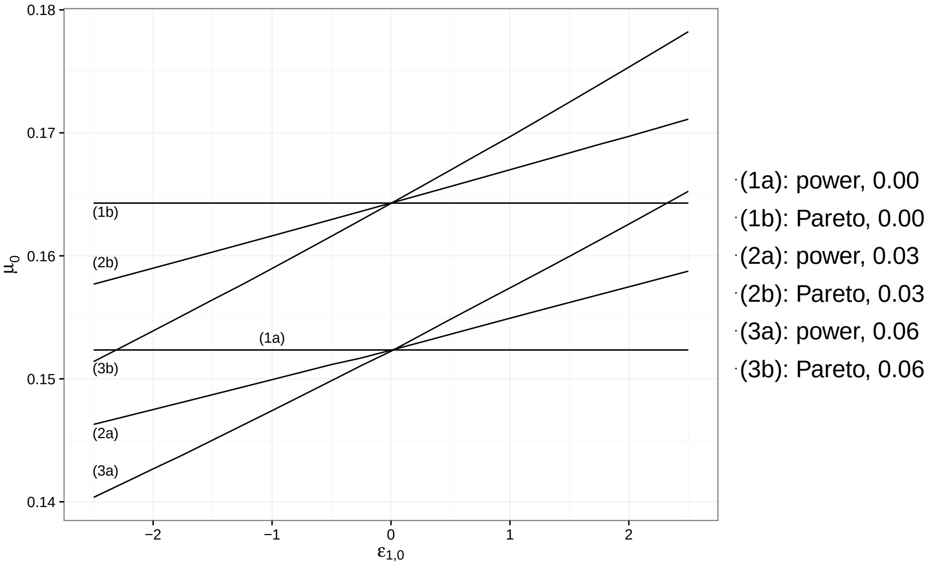

7.2 Uncertainty in technological efficiency

We next suppose that $\tau_1 = \tau_2 \equiv \tau > 0$ and $\epsilon_{1,t} = \epsilon_{2,t} \equiv \epsilon_t$, which means that uncertainty enters SSICE through technological efficiency A_t (see (17)–(18)). This implies in particular that uncertainty now appears again in the budget constraint just like in Section 6, and the current extension can technically be viewed as a marriage between the settings of Sections 6 and 7.1. We analyze again the impact of this alternative specification in the spectrum of uncertainties that our model formulation accommodates on the three main findings in Sections 6.3–6.5.

With iid normally distributed errors ϵ_t taking values along their expected trajectory, i.e., in a setting analogous to Section 6.3, all three optimal policies under the current extension closely resemble those observed under the base SSICE model in Section 6.3. Intuitively, this follows from the insensitivities of the optimal policies along the expected trajectory observed in Section 7.1 and the fact that the current extension is technically a marriage between the base model and the first extension.

Next, when catastrophic risk realizes, that is, when ϵ_0 takes large negative values analogous to the analysis at the edges in Section 6.4, optimal consumption responds exactly as in Figure 1. However, optimal abatement, while increasing with more uncertainty (that is, behaving monotonically in τ as in Section 6.4), decreases with the extent of the negative shock ϵ_0; see

Figure 5: Abatement μ_0: SSICE-extension-II with normal errors and τ = 0.00, 0.03, and 0.06 — power versus Pareto utility

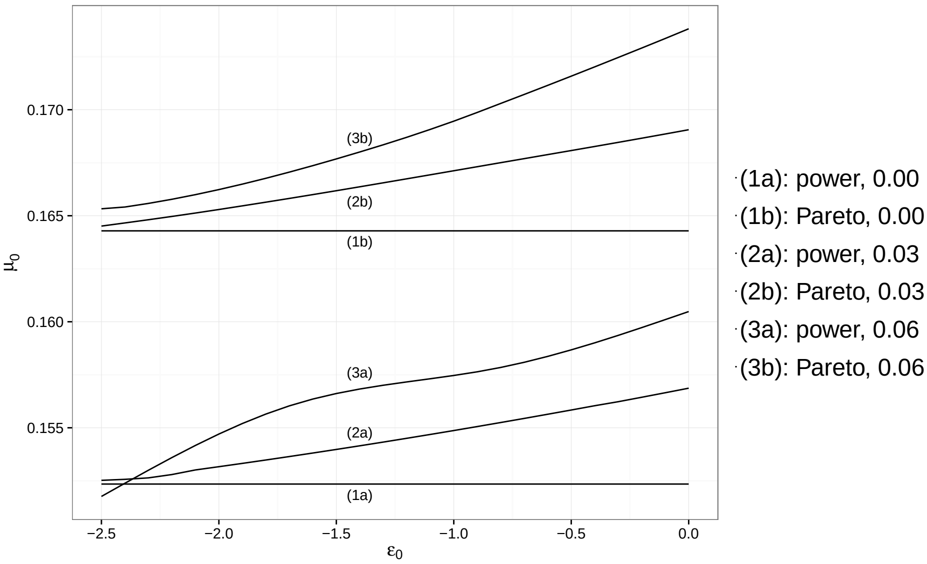

Figure 5. The latter effect is in part induced by the abatement results in Section 7.1, illustrated in Figure 4.

Finally, analyzing heavy-tailed uncertainty in technological efficiency, employing the same distributional assumptions as in Section 6.5, we recover the exact same pattern as in Figure 3.

Apparently, the impact of uncertainty is similar whether we model it through its impact in the damage function or through technological efficiency. The key difference between the base SSICE model and our second extension is that, while uncertainty in technological efficiency increases optimal abatement just like uncertainty in the damage function, this effect is suppressed in adverse technology scenarios in which budgets are lower but also emissions will automatically be lower, thus facilitating lower abatement.

8 Conclusions

We have developed a stochastic dynamic finite-horizon economic framework with climate change and a regression-based method for numerically solving the associated optimization problem. Our framework (SSICE) provides a parsimonious representation of Nordhaus' deterministic DICE model, but it incorporates, possibly heavy-tailed, stochasticity. Upon applying our solution method to SSICE our analysis reveals that the introduction of uncertainty into a deterministic integrated assessment model can have a substantial impact on the optimal policies of abatement, consumption, and investment. This cannot just be explained by a regular precautionary savings motive. It is also due to our framework's recognition (under Pareto utility) of the benefits of precautionary action in adverse circumstances which can in part get lost in more conventional economy-climate models (with power utility) in the mire of a strong desire to keep up consumption. These findings remain intact, and get reinforced, under heavy tails.

Therefore, precautionary action under conventional integrated assessment models can be too low if there is no account of uncertainty; if a substitution from precaution to consumption tends to occur in adverse scenarios; and if there is no consideration of heavy tails in identifying the optimal abatement, consumption, and investment policies.

Appendix A: SICE versus DICE

Implementation of the SICE model of Section 2 requires specification of parameter values and initial levels. These are presented in Table 4.

Table 4: Parameter values in SICE (simplified DICE) model

Parameter	Value	Description
Endogenous stocks: initial levels		
K_0	135	Capital stock, beginning of period 0
M_0	819	CO_2 concentration, beginning of period 0
H_0	0.83	Temperature, beginning of period 0
Technology		
γ	0.30	Elasticity of capital in production function
δ	0.4095	Depreciation rate on capital, per five years
Pollution, damage, and abatement		
ϕ	0.0529	Depreciation rate on CO_2 concentration, per five years
ξ	0.00267	Quadratic term, temperature-impact function
θ	2.80	Exponent in abatement function
Temperature		
η_0	−3.1761	Constant term, temperature equation
η_1	0.9042	Previous period impact, temperature equation
η_2	0.4995	CO_2 concentration impact, temperature equation
Discount rate		
ρ	0.0773	Welfare discount rate, per five years

The parameter values are the same as in DICE-2013 in all cases where we use the same equations. The initial values of the endogenous stock variables are also the same as in DICE-2013. In cases where we simplified the equations, notably the carbon equation (4) and the temperature equation (5), we had to adjust the parameters.

These adjusted parameters are chosen to mimic DICE-2013. The exogenous variables L_t, A_t, σ_t, and ψ_t are the same as in DICE-2013. We let $E_t^0 = 42.2733 - 0.3128\, E_{land,t}$, where $E_{land,t}$ is equal to exogenous 'land-use emissions' as specified in DICE-2013, and the coefficients 42.2733 and −0.3128 have been calibrated through linear regression and DICE-2013 output so as to bring our climate model close to DICE-2013.

To supplement Table 1 we present some further results for the deterministic SICE model; see also Section 6.2.

Table 5 presents results for the control variables consumption, investment, and abatement. Rather than consumption C_t and investment I_t we present

Table 5: Control variables: DICE versus SICE

		2010	2035	2060	2085
	DICE	0.7395	0.7464	0.7442	0.7358
Consumption	power	0.7378	0.7457	0.7448	0.7400
(C_t/Y_t)	Pareto	0.6973	0.7408	0.7516	0.7526
	DICE	0.2587	0.2467	0.2405	0.2373
Investment	power	0.2601	0.2481	0.2420	0.2376
(I_t/Y_t)	Pareto	0.3005	0.2530	0.2353	0.2248
	DICE	0.0000	0.0016	0.0040	0.0076
Abatement spending	power	0.0003	0.0008	0.0013	0.0012
($\omega_t D_t$)	Pareto	0.0004	0.0007	0.0010	0.0009

C_t/Y_t and I_t/Y_t, in relative terms. Regarding abatement, we write, using (6) and (8),

$$D_t = \frac{C_t + I_t}{Y_t} + \omega_t D_t, \qquad \omega_t = \psi_t \mu_t^{\theta},$$

and present $\omega_t D_t$ rather than μ_t itself in the third panel of Table 5, so that the three variables become comparable.

Table 6: State variables: DICE versus SICE

		2010	2035	2060	2085
	DICE	2.1246	2.2017	2.2887	2.3627
Capital	power	2.1244	2.2124	2.3031	2.3759
(K_t/Y_t)	Pareto	2.1244	2.3319	2.3009	2.2875
	DICE	12.8889	6.8168	4.2907	2.8942
Concentration	power	12.8876	6.8446	4.4201	3.1746
(M_t/Y_t)	Pareto	12.8876	6.7100	4.4581	3.2698
	DICE	0.8300	1.4037	2.0673	2.6864
Temperature	power	0.8300	1.4279	2.1219	2.8467
(H_t)	Pareto	0.8300	1.4302	2.1359	2.8790

Table 6 presents results for the state variables and is an extended version of Table 1.

We compare the DICE model (under power utility) with two versions of the SICE model: power utility and Pareto utility. Over periods 0–15 both the control and state variables related to the economy part (C, I, and K)

of the SICE model under power utility closely follow the DICE model. In particular, C/Y, I/Y and K/Y under SICE with power utility deviate from DICE by no more that 0.6% over the 16 periods that we consider. The main differences between DICE and SICE are related to the climate modules. The deviation between the control and state variables pertaining to the climate part of SICE (μ, M, and H) and the climate part of DICE is somewhat more pronounced, but still rather small.

Furthermore, the control and state variables under the SICE model with Pareto utility closely resemble those of the SICE model under power utility. This is true for both the economy and climate parts of the SICE model. This, perhaps, is not surprising because Pareto utility was calibrated to resemble power utility along the expected trajectory of $\epsilon_{2,t}$ over the reported period.

Appendix B: SSICE with normal errors

In Table 2 of Section 6.3 we presented a comparison between power and Pareto utility for the control variables consumption, investment and abatement. We complement this table by presenting the corresponding results for the state variables.

Table 7: SSICE model with normal errors, under power and Pareto utility

	power			Pareto		
$t \backslash \tau_2$	0.00	0.03	0.06	0.00	0.03	0.06
Capital K_t						
2010	135.00	135.00	135.00	135.00	135.00	135.00
2035	309.80	304.62	302.03	333.99	336.20	338.55
2060	586.23	574.59	560.89	585.42	575.22	562.93
2085	968.42	967.45	968.20	917.34	921.61	925.49
Concentration M_t						
2010	818.99	818.99	818.99	818.99	818.99	818.99
2035	958.44	957.35	956.56	961.03	960.72	960.25
2060	1125.06	1123.59	1122.48	1134.26	1134.11	1133.85
2085	1293.97	1293.32	1292.71	1311.26	1311.60	1311.55
Temperature H_t						
2010	0.8300	0.8300	0.8300	0.8300	0.8300	0.8300
2035	1.4279	1.4264	1.4252	1.4302	1.4294	1.4284
2060	2.1219	2.1187	2.1165	2.1359	2.1357	2.1351
2085	2.8467	2.8423	2.8384	2.8790	2.8780	2.8762

The three panels in Table 7 present the results for capital, concentration, and temperature, respectively.

References

Ackerman, F., Stanton, E. A., Bueno, R., 2010. Fat tails, exponents, extreme uncertainty: Simulating catastrophe in DICE. *Ecological Economics* 69, 1657–1665.

Barro, R. J., 2009. Rare disasters, asset prices, and welfare costs. *American Economic Review* 99, 243–264.

Berger, L., Emmerling, J., Tavoni, M., 2016. Managing catastrophic climate risks under model uncertainty aversion. *Management Science*, published online.

Bernoulli, D., 1738. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5, 175–192.

Bertsekas, D. P., 2005. *Dynamic Programming and Optimal Control*, Third Edition. Athena Scientific, Belmont, MA.

Brandt, M. W., Goyal, A., Santa-Clara, P., Stroud, J. R., 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies* 18, 831–873.

Buchholz, W., Schymura, M., 2012. Expected utility theory and the tyranny of catastrophic risks. *Ecological Economics* 77, 234–239.

Cai, Y., Judd, K. L., Lontzek, T. S., 2015. DSICE: A dynamic stochastic integrated model of climate and economy. RDCEP Working Paper No. 12-02.

Cai, Y., Lenton, T. M., Lontzek, T. S., 2016. Risk of multiple interacting tipping points should encourage rapid CO_2 emission reduction. *Nature Climate Change* 6, 520–525.

Carriere, J. F., 1996. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19, 19–30.

Cerreia-Vioglio, S., Dillenberger, D., Ortoleva, P., 2015. Cautious expected utility and the certainty effect. *Econometrica* 83, 693–728.

Clément, E., Lamberton, D., Protter, P., 2002. An analysis of a least squares regression method for American option pricing. *Finance and Stochastics* 6, 449–471.

Dietz, S., 2011. High impact, low probability? An empirical analysis of risk in the economics of climate change. *Climatic Change* 108, 519–541.

Egloff, D., 2005. Monte Carlo algorithms for optimal stopping and statistical learning. *The Annals of Applied Probability* 15, 1396–1432.

Egloff, D., Kohler, M., Todorovic, N., 2007. A dynamic look-ahead Monte Carlo algorithm for pricing Bermudan options. *The Annals of Applied Probability* 17, 1138–1171.

Geweke, J., 2001. A note on some limitations of CRRA utility. *Economics Letters* 71, 341–345.

Gillingham K., Nordhaus, W. D., Anthoff, D., Blanford, G., Bosetti, V., Christensen, P., McJeon, H., Reilly, J., Sztorc, P., 2015. Modeling uncertainty in climate change: A multi-model comparison. NBER Working Paper No. 21637.

Gordon, R. B., Koopmans, T., Nordhaus, W., Skinner, B., 1987. *Toward a New Iron Age? Quantitative Modeling of Resource Exhaustion.* Harvard University Press, Cambridge, Mass.

Hwang, I. C., Reynès, F., Tol, R. S. J., 2013. Climate policy under fat-tailed risk: An application of DICE. *Environmental and Resource Economics* 56, 415–436.

Ikefuji, M., Laeven, R. J. A., Magnus, J. R., Muris, C., 2013. Pareto utility. *Theory and Decision* 75, 43–57.

Ikefuji, M., Laeven, R. J. A., Magnus, J. R., Muris, C., 2015. Expected utility and catastrophic consumption risk. *Insurance: Mathematics and Economics* 64, 306–312.

Jensen, S., Traeger, C. P., 2014. Optimally climate sensitive policy under uncertainty and learning. Working paper, presented at the Conference on Sustainable Resource Use and Economic Dynamics (SURED), Ascona, Switzerland.

Keller, K., Bolker, B. M., Bradford, D. F., 2004. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management* 48, 723–741.

Kelly, D. L., Kolstad, C. D., 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics & Control* 23, 491–518.

Kimball, M., 1990. Precautionary saving in the small and in the large. *Econometrica* 58, 53–73.

Krätschmer, V., Ladkau, M., Laeven, R. J. A., Schoenmakers, J. G. M., Stadje, M., 2015. Robust optimal stopping. Working Paper, University of Amsterdam, University of Ulm, and WIAS Berlin.

Laeven, R. J. A., Stadje, M., 2014. Robust portfolio choice and indifference valuation. *Mathematics of Operations Research* 39, 1109–1141.

Leach, A. J., 2007. The climate change learning curve. *Journal of Economic Dynamics & Control* 31, 1728–1752.

Lemoine, D., Traeger, C. P., 2014. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy* 6, 137–166.

Longstaff, F. A., Schwartz, E. S., 2001. Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies* 14, 113–147.

Lontzek, T. S., Cai, Y., Judd, K. L., Lenton, T. M., 2015. Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nature Climate Change* 5, 441–444.

Manne, A. S., Richels, R. G., 1992. *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emission Limits.* MIT Press, Cambridge, MA.

Mastrandrea, M. D., Schneider, S. H., 2004. Probabilistic integrated assessment of 'dangerous' climate change. *Science* 304, 571–575.

Menger, K., 1934. Das Unsicherheitsmoment in der Wertlehre: Betrachtungen im Anschluß an das sogenannte Petersburger Spiel. *Zeitschrift für Nationalökonomie (Journal of Economics)* 5, 459–485.

Nordhaus, W. D., 1994. *Managing the Global Commons: The Economics of Climate Change.* MIT Press, Cambridge, MA.

Nordhaus, W. D., 2008. *A Question of Balance: Weighing the Options on Global Warming Policies.* Yale University Press, New Haven, CT.

Nordhaus W. D., 2013. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World.* Yale University Press, New Haven, CT.

Nordhaus, W. D., Yang, Z., 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. *American Economic Review* 86, 741–765.

Pindyck, R. S., 2011. Fat tails, thin tails, and climate change policy. *Review of Environmental Economics and Policy* 5, 258–274.

Powell, W. B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Second Edition. John Wiley & Sons, Hoboken, NJ.

Roughgarden, T., Schneider, S. H., 1999. Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* 27, 415–429.

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

Traeger, C. P., 2014. A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics* 59, 1–37.

Tsitsiklis, J., Van Roy, B., 1999. Regression methods for pricing complex American style options. *IEEE Transactions on Neural Networks* 12, 694–703.

Ursúa, J. F., 2010. Long-run volatility. Mimeo, Harvard University.

Weitzman, M. L., 2009. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics* 91, 1–19.