Ikefuji, M., Laeven, R., Magnus, J., & Muris, C. (2020). Expected utility and catastrophic risk in a stochastic economy-climate model. *Journal of Econometrics*, 214(1), 110-129.
<https://doi.org/10.1016/j.jeconom.2019.05.007>

Peer reviewed version

License (if available):
CC BY-NC-ND

Link to published version (if available):
[10.1016/j.jeconom.2019.05.007](https://doi.org/10.1016/j.jeconom.2019.05.007)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at <https://www.sciencedirect.com/science/article/pii/S0304407619301101> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

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

Masako Ikefuji

Faculty of Humanities and Social Sciences, University of Tsukuba

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, University of Bristol

December 8, 2018

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 (SDICE*) incorporates, possibly heavy-tailed, stochasticity in Nordhaus' deterministic DICE model. 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 SDICE* and examine the effects of light- and heavy-tailed uncertainty. The results indicate that the effects can be substantial, depending on the nature and extent of the uncertainty and the social planner's preferences.

JEL Classification: C1; E2; Q5.

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

*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 addresses: ikefuji.masako.gn@u.tsukuba.ac.jp (Ikefuji), R.J.A.Laeven@uva.nl (Laeven), jan@janmagnus.nl (Magnus), chris.muris@bristol.ac.uk (Muris).

1 Introduction

The current economy-climate debate raises many difficult issues. Only one of these 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' (2017a, 2017b) 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 DICE model is our starting point and the deterministic version of our model reduces to DICE. 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 DICE random shocks featuring potentially heavy-tailed risk. We refer to the resulting stochastic model as SDICE*. We initially focus on uncertainty through the damage-abatement fraction and, later, in an extension of this base stochastic model, we shall also account for uncertainty in the damage parameter, 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, permitting both light- and heavy-tailed risks and stochastic parameters.

In the context of SDICE* 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-abatement fraction of SDICE* leads to a reduction of abatement, while optimal consumption and investment are relatively less affected. Conditional upon the shocks realizing their expected value, namely zero, the pattern remains the same in later periods, and this applies to both power and Pareto utility. Although lower abatement has a negative propagation effect in our economy-climate model, the social planner sacrifices some abatement to maintain consumption and investment, under

uncertainty about the damage-abatement fraction. Compared to a Pareto utility maximizer, the power utility maximizer has a stronger motive to avoid a low damage-abatement reduced output, and therefore consumes less in early periods and invests and abates more in all periods. 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, a power utility maximizer has a stronger motive to abate as a precaution (Kimball, 1990) than a Pareto utility maximizer, and this amplifies with adverse realizations of the shocks. 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 power utility maximizer tends to favor a relatively larger substitution from current consumption to intensified precautionary action when compared to Pareto utility. This effect is distinct from the higher levels of abatement observed for a power utility maximizer in the center of the distribution (and under the deterministic model) which it further amplifies.

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 SDICE* model.

Fourth, in the center of the distribution the impact of uncertainty in the damage-abatement fraction dominates the impact of uncertainty in the damage parameter and an uncertain emissions-to-output ratio and closely resembles the impact of uncertainty through technological efficiency. At the edges, an uncertain damage parameter impacts consumption, but leaves abatement nearly unaffected, while uncertainty in the emissions-to-output ratio significantly impacts abatement but has relatively little effect on consumption. Furthermore, when adverse scenarios for technological efficiency realize, optimal abatement is suppressed compared to the adverse scenarios in which the damage-abatement fraction is relatively small.

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 (2014b) study the effects of climate sensitivity uncertainty, learning, and temperature stochasticity separately, and find pre-

cautionary 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-abatement fraction (and the damage parameter, 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. (2012, 2016), and Berger *et al.* (2017) 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.* (2017) 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.* (2012) 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 SDICE* 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 Stadjé (2014), who solve prob-

lems 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.* (2017), who employ LSMC to analyze model uncertainty in optimal stopping.

The paper is organized as follows. In Section 2 we succinctly summarize DICE. In Section 3 we introduce uncertainty into DICE. 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 SDICE* 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 results of our SDICE* model and discuss their implications, while some extensions are presented in Section 7. Section 8 concludes.

2 Nordhaus' DICE model

Our analysis takes as its starting point the DICE model, more precisely the ‘beta version’ of DICE-2016R, a version with identification DICE-2016R-091916ap.gms; see Nordhaus (2017a, 2017b). (Earlier versions are in Nordhaus and Yang (1996) and Nordhaus (2008, 2013).) We briefly summarize this model in condensed form.

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

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

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

Capital is accumulated through

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

where I_t denotes investment and δ is the depreciation rate of capital.

Carbon dioxide (CO₂) emissions consist of industrial emissions (caused by production) and non-industrial emissions. We denote the latter type by E_t^0 and consider it to be exogenous to the model. Total CO₂ emissions E_t are then given by

$$E_t = \sigma_t(1 - \mu_t)Y_t + E_t^0, \quad (3)$$

where σ_t is the emissions-to-output ratio for CO₂ and μ_t is the abatement fraction for CO₂. The associated CO₂ concentration increase M_t in the atmosphere (GtC from 1750) accumulates through

$$M_{t+1} = (1 - b_0)M_t + b_1 M_t^{(s)} + E_t, \quad (4a)$$

$$M_{t+1}^{(s)} = b_0 M_t + (1 - b_1 - b_3)M_t^{(s)} + b_2 M_t^{(l)}, \quad (4b)$$

$$M_{t+1}^{(l)} = b_3 M_t^{(s)} + (1 - b_2)M_t^{(l)}, \quad (4c)$$

where $M_t^{(s)}$ and $M_t^{(l)}$ are auxiliary variables representing CO₂ concentration increases in shallow and lower oceans, respectively.

Temperature increase H_t (degrees Celsius from 1900) is given by

$$H_{t+1} = (1 - a_0)H_t + a_1 \log(M_{t+1}) + a_2 H_t^{(l)} + F_{t+1}, \quad (5a)$$

$$H_{t+1}^{(l)} = (1 - a_3)H_t^{(l)} + a_3 H_t, \quad (5b)$$

where $H_t^{(l)}$ is an auxiliary variable representing temperature increase of the lower oceans and F_{t+1} is exogenous radiative forcing.

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

$$(1 - \omega_t - \xi H_t^2) Y_t = C_t + I_t. \quad (6)$$

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

$$\omega_t = \psi_t \mu_t^\theta \quad (\theta > 1). \quad (7)$$

When μ_t increases then so does ω_t , and a larger fraction of GDP will be spent on abatement. Damage is represented by a fraction ξH_t^2 of Y_t , and it depends only on temperature. The optimal temperature is $H_t = 0$, the temperature in 1900. Deviations from the optimal temperature cause damage.

Equations (1)–(7) imply a condensed system consisting of six dynamic equations (in the state variables K , M , H , $M^{(s)}$, $M^{(l)}$, and $H^{(l)}$) in terms of the three policy variables I , μ , and (through the budget constraint) C .

As in DICE-2016R one period is five years. Period 1 refers to the time interval 2015–2019, period 2 to 2020–2024, and so on. Stock variables are measured at the beginning of the period. 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

$$C_t \geq 0, \quad I_t \geq 0, \quad \mu_t \geq 0. \quad (8)$$

Nordhaus restricts $\mu_t \leq 1$ in the early periods, but he allows an upper bound of 1.2 from period 30 onwards (year 2160). We do the same.

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

$$W_t = L_t U(C_t/L_t). \quad (9)$$

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

$$W = \sum_{t=1}^T \frac{W_t}{(1+\rho)^t} \quad (0 < \rho < 1), \quad (10)$$

where ρ is the discount rate and $T = 100$ (500 years). 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

$$U(x) = \frac{x^{1-\alpha} - 1}{1-\alpha} \quad (\alpha > 0), \quad (11)$$

where α is the elasticity of marginal utility of consumption. This is the so-called *power* function. Many authors, including Nordhaus, select this function. All parameter- and initial values are the same as in DICE-2016R.

3 Introducing stochasticity

We now introduce uncertainty in the DICE model, focussing on the uncertainty about the economic impact of future climate change. Uncertainties can arise in many ways. Gillingham *et al.* (2015) distinguish between seven types of uncertainties including most notably parametric uncertainty, such as uncertainty about the climate sensitivity parameter; model or specification uncertainty, such as the specification of the aggregate production function or the damage function; measurement error, for example related to the level and trend of global temperature; and random errors in structural equations, for example due to weather shocks.

We shall study the impact of stochasticity in four scenarios, defined by adding four random shocks $u_{j,t}$ ($j = 1, \dots, 4$) to the condensed system (1)–(7), as follows:

- (i) Uncertainty through the damage-abatement fraction $d_t = 1 - \omega_t - \xi H_t^2$ in (6) by modifying d_t to $\bar{d}_t = d_t u_{1,t}$;
- (ii) Uncertainty in the damage parameter by modifying ξ in (6) to $\bar{\xi}_t = \xi u_{2,t}$;
- (iii) Uncertainty in the emissions-to-output ratio by modifying σ_t in (3) to $\bar{\sigma}_t = \sigma_t u_{3,t}$; and
- (iv) Uncertainty through technological efficiency by modifying A_t in (1) to $\bar{A}_t = A_t u_{4,t}$.

We specify the shocks $u_{j,t}$ as

$$u_{j,t} = e^{-\tau_j^2/2} e^{\tau_j \epsilon_{j,t}} \quad (j = 1, 2, 3, 4), \quad (12)$$

where the normalizing constants $e^{-\tau_j^2/2}$ are chosen such that $\mathbb{E}(u_{j,t}) = 1$ when $\epsilon_{j,t} \sim N(0, 1)$. We never change the parameters of a given probability distribution of $\epsilon_{j,t}$, we only vary the level of τ_j . Also, note that under scenario (ii), when the temperature change is zero, there is no uncertainty.

The general model formulation resulting from these four scenarios 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. Thus we obtain a stylized stochastic IAM of climate economics, to which we refer as SDICE* (= stochastic DICE).

Uncertainty in the damage parameter ξ (scenario (ii)) has been considered by Keller *et al.* (2004), Ackerman *et al.* (2010), Dietz (2011), Hwang *et al.* (2013), Berger *et al.* (2017), and Howard and Sterner (2017). The impact of uncertain technological efficiency (scenario (iv)) is studied by Cai *et al.* (2012) and Jensen and Traeger (2014a), and its importance is emphasized by Gillingham *et al.* (2015). Uncertainty in the emissions-to-output ratio (scenario (iii)) has not received much attention. Uncertainty in the abatement cost has been studied in Hwang *et al.* (2013). We first focus on the combined uncertainty of damage and abatement, as specified in scenario (i), but later we shall discuss all scenarios.

We shall consider both light- and heavy-tailed risk. If the $\epsilon_{j,t}$ are independent and identically distributed (iid) and follow a normal distribution $N(0, 1)$, then the moments of $u_{j,t}$ exist, and we have $\mathbb{E}(u_{j,t}) = 1$ and $\text{var}(u_{j,t}) = e^{\tau_j^2} - 1$. Since the distribution of $u_{j,t}$ is heavily skewed, more uncertainty (higher τ_j) implies more probability mass of $u_{j,t}$ close to zero. If, however, we move

away from the normal distribution and assume, e.g., that $\epsilon_{j,t}$ follows a Student distribution with *any* (finite) degrees of freedom, then the expectation is infinite (Geweke, 2001). The analysis in, among others, Weitzman (2009), 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 $\tau_j > 0$ then the assumption of iid distributed errors $\epsilon_{j,t}$ is sufficient to generate the possibility of incompatibility between preferences and distributional assumptions, as discussed and proved in Section 5. Provided the compatibility conditions are satisfied, 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 SDICE* model in Section 3 belongs as a special case, and develop a regression-based method to solve such problems. We first introduce some notation, define our general class of optimization problems, and show how it encompasses SDICE* 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 cap-

ture tail-risk behavior far away from a rapidly evolving steady state. We shall indicate how our solution algorithm deals with each of these challenges. In the final subsection we compare our method with a standard solution approach.

4.1 SDICE* as a stochastic dynamic finite-horizon optimization problem

The social planner in SDICE* faces a discrete-time stochastic dynamic finite-horizon programming problem, which consists of maximizing expected total discounted welfare W given in (10) subject to the SDICE* model specified in the condensed system of equations (1)–(7) and including one (or several) of the four shocks specified in Section 3.

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 SDICE*, the prime state variables are the capital stock K_t , temperature H_t , and carbon dioxide concentrations M_t , and the auxiliary variables $(M_t^{(s)}, M_t^{(l)}, H_t^{(l)})$. The prime control variables are consumption C_t and the abatement fraction μ_t . Stochasticity enters SDICE* through $\epsilon_t = (\epsilon_{1,t}, \epsilon_{2,t}, \epsilon_{3,t}, \epsilon_{4,t})$, where $j = 1, \dots, 4$ refers to the scenario through which we introduce the stochasticity.

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 SDICE*, the exogenous variables and parameters are: the initial values of the state variables, $(K_1, H_1, M_1, M_1^{(s)}, M_1^{(l)}, H_1^{(l)})$, the time-varying exogenous stock variables and parameters $(A_t, L_t, \psi_t, \sigma_t, E_t^0, F_t)$, the time-invariant parameters $(\gamma, \delta, \rho, \phi, \xi, \theta, a_0, a_1, a_2, a_3, b_0, b_1, b_2, b_3)$, and the stochasticity parameters $(\tau_1, \tau_2, \tau_3, \tau_4)$. The remaining endogenous variables in the model — investment I_t , emissions E_t , the abatement cost fraction ω_t , welfare W_t , and output Y_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 (6) 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 SDICE*. 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.

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

$$z_t \in \mathcal{Z}_t(x_t, \epsilon_t). \quad (13)$$

For the first twenty periods of the SDICE* model with uncertainty as in scenario (i), this set of constraints specializes to:

$$0 \leq C_t \leq (1 - \omega_t - \xi H_t^2) A_t K_t^\gamma L_t^{1-\gamma} e^{-\tau_1^2/2} e^{\tau_1 \epsilon_{1,t}}$$

and

$$0 \leq \mu_t \leq 1;$$

see also (17). For scenarios (ii) and (iii), the first constraint will incorporate uncertainty in a different manner. For later time periods, the second constraint is modified to $0 \leq \mu_t \leq 1.2$.

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 ($1 \leq s \leq T$) 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 (13):

$$V_s(x_s) = \max_{z_s, \dots, z_T} \mathbb{E}_{s-1} \left[\sum_{t=s}^T g_t(z_t) \right] \quad (14)$$

subject to

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

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 . We note that, given the information at time $s - 1$, the vector function f_{s-1} dictates the prime state variables at time s (but not the control variables). Therefore, optimization for period s can be based upon the conditional expectation at time $s - 1$. The function V_s is referred to as the value function. The corresponding Bellman equation is given by

$$V_t(x_t) = \max_{z_t \in \mathcal{Z}_t(x_t, \epsilon_t)} \mathbb{E}_{t-1} [g_t(z_t) + \beta V_{t+1}(f_t(x_t, z_t, \epsilon_t))], \quad (15)$$

where β is a discount factor ($0 < \beta < 1$). The time- t objective function $g_t(z_t)$ in the SDICE* 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 SDICE* is equal to $\beta = 1/(1 + \rho)$.

If the value function in period $t + 1$ is known, then Equation (15) 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. A potentially even more efficient approach may be to select fewer but better nodes for approximation and thus economize on the number of generated paths, as in Cai *et al.* (2012) and Traeger (2014). We don't pursue this. Methods to handle high-dimensional problems, for example with adaptive sparse grids, are discussed in Brumm and Scheidegger (2017), but the dimensionality of our setting is only six so we don't pursue this either. Moreover, a variety of numerical approaches for infinite-horizon problems is available, in particular in the DSGE literature. Examples of such approaches and problems are King and Rebelo (1999) who accommodate stochastic trends via an equivalent representation without trending variables in the context of technology shocks and real business cycles, and Schmitt-Grohé and Uribe (2004) who consider a stochastic production economy with sticky product prices. Our setting differs from these approaches and problems due to at least one of the following three features: (i) our problem is a finite-horizon problem with general preferences and beliefs; (ii) our problem is highly non-stationary; (iii) we are interested in a solution at (a) specific time period(s), not necessarily in the steady state, and not only in the center of the distribution, but also at the edges.

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

$$V_T(x_T) = \max_{z_T \in \mathcal{Z}_T(x_T, \epsilon_T)} \mathbb{E}_{T-1}[g_T(z_T)]. \quad (16)$$

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 (15) simplifies to the Bellman equation in (16) 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 $r = 1, \dots, R$. For SDICE*, 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 (DICE). 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 SDICE* 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 $\beta_{j,T}$ ($j = 0, 1, 2, \dots$) 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 $\hat{\beta}_{j,T}$ ($j = 0, \dots, J$), and we define

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

as our approximation to the value function at time T .

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

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

The algorithm then proceeds as above in four steps: (a) generate draws $(x_{T-1}^r, \epsilon_{T-1}^r)$ for $r = 1, \dots, R$; (b) given the r th draw $(x_{T-1}^r, \epsilon_{T-1}^r)$, com-

pute the deterministic maximum v_{T-1}^r , using the approximation $\widehat{V}_T$ obtained above; (c) obtain estimates $\hat{\beta}_{j,T-1}$ for the coefficients $\beta_{j,T-1}$ in

$$v_{T-1}^r = \sum_{j=0}^J \beta_{j,T-1} \phi_j(x_{T-1}^r) + \nu_{r,T-1},$$

using least squares; and (d) define the 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 $t = 1, \dots, T$ 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 SDICE* 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 (DICE). We draw x_t^r from a uniform distribution with support $[0.5x_t^*, 2.5x_t^*]$. Taking the support even wider does not help our solution: the optimization routine never evaluates values of the state variables

outside those intervals. 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 1, capital is equal to $K_1 = 223$. The specified support for K_2 is now given by $[134.2, 671.2]$. Even under very large shocks in period 1, the lower bound of the support is never binding for the optimal choice of K_2 .

Number of time periods and draws. We must also specify the number of time periods for which we solve the model and the number of simulation draws R to be drawn in every time period (each time period consists of five years). We solve the model for $T = 100$ time periods. We ignore uncertainty in periods 21 through 100. We report results at the center for periods 1, 6, 11, 16, and at the edges in periods 1 and 2. We experimentally determine that considering stochasticity for more than 20 periods does not affect our reported results: the reported results do not change if we increase the number of periods with stochasticity to 30 or 40.

An elegant alternative way to formulate the economy-climate problem would be to set it in an infinite-horizon framework, and to solve it on small time-steps to approximate the continuous-time solution, as in Cai *et al.* (2012) and Traeger (2014). However, because we wish to stay close to our starting point given by the benchmark DICE model, we adopt the same time-steps and finite-horizon formulation as in DICE.

The value of R used for the first 20 periods has been determined by trial and error. We started with $R = 1,000$ simulations per period, and then assessed whether the solution is sensitive to increases in R using steps of 1,000. After $R = 5,000$, the change in optimal consumption was 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. Similarly, for periods 21 through 100, we find that results are not sensitive to increasing the number of draws beyond 200. We conservatively set $R = 1,000$ for those periods.

Basis functions. The selection of the basis functions in the approximation of the value function is an important ingredient of our procedure. We use polynomials as basis functions. To select the order of the polynomials, and to select which interactions terms to include in our model, we proceed in three steps. First, we use a model selection approach to choose the order of the polynomial. Second, we use model selection to choose a set of interaction terms. Both steps are performed using power utility. In a third step, we repeat the first two steps under Pareto utility.

Let us now describe each of the three steps in more detail. We start by solving the DICE model for all time periods under power utility using as basis functions only a constant term and one linear function for each state variable. For that solution, we measure the fit of the value function approximation in each time period by computing the adjusted R-squared, Akaike's Information Criterion, and the leave-one-out cross-validation criterion. Next, we solve the model using as basis functions a constant term, and a linear and quadratic function in each state variable, and record the same criteria in each time period. We repeat this for increasing sets of basis functions, up to polynomial order 8.

The following conclusions hold for each of the three criteria: (i) inclusion

of the auxiliary state variables $(M_t^{(s)}, M_t^{(l)}, H_t^{(l)})$ does not reduce the criteria, and occasionally leads to numerical instability; (ii) a fourth-order polynomial in (K_t, M_t, H_t) outperforms lower-order polynomials in all time periods; (iii) a fifth-order polynomial is competitive with the fourth-order polynomial in later time periods ($t > 15$) but not in the early time periods. (iv) numerical instability is more serious for polynomials of order greater than 5, while those specifications do not improve the criteria relative to the fourth-order polynomial. This is visualized in Figure 1, where we plot Akaike's Information Criterion (AIC) for the sets of basis functions up to polynomial order 5, against the time period. (Plots for the other criteria are similar.) Because the criteria are computed over discrete time periods, the curves are somewhat granular rather than smooth. On the basis of these findings, we proceed to the next step of our model selection procedure using the fourth-order polynomial in (K_t, M_t, H_t) .

In the second step, we repeat this procedure by gradually adding interaction terms to the fourth-order specification selected in step 1. We solve the model using a fourth-order polynomial augmented by interaction terms with a combined degree of 2, i.e., $K_t M_t$, $K_t H_t$, and $M_t H_t$. This leads to improved criteria in all time periods. For example, the adjusted R-squared, averaged across time periods, is improved from approximately 0.9995 to 0.9999. We then solve the model again after adding further interaction terms with a combined degree of 3, e.g., $K_t^2 M_t$. This improves the criteria in almost all time periods, although the performance gain is much smaller. Adding further interaction terms leads to numerical instability in the early periods, without improving the criteria in the later periods. We illustrate this in Figure 2,

where we plot the AIC for the sets of interaction terms up to a combined degree of 3, against the time period. (Plots for the other criteria are similar.)

As a third step, we apply the procedure outlined above to the model with Pareto utility. For both steps, we reach the same conclusion as in the power utility case. The adjusted R-squared exceeds 0.9999 in all time periods for the selected sets of basis functions. In view of the fact that we simulate the state variables in the regression and their realizations lie in a very large interval with a substantial fraction far away from the center (see above sub ‘The support’), this suggests in particular that our approximation to the value function is good also at the edges. As a result, we choose a fourth-order polynomial with interaction terms up to a combined degree of three for all of the numerical results in this paper.

Code and testing details.

The code is written in `Julia` and is available from the authors upon request. It was tested with `Julia` version 0.6.2, on a desktop computer with Core i5-6300U architecture running Windows 10.

4.4 Comparison to standard algorithms

We briefly compare our algorithm with a more standard solution approach. We shall do so in a simplified setting without climate. More specifically, we shall use the same model and identical values of all time-varying and time-invariant parameters, except that we set $\omega_t = 0$ for all t and also $\xi = 0$. These conditions imply that climate plays no role. In this experiment, utility is always power utility and welfare is defined as before; see (9)–(11). Uncer-

tainty enters through technological efficiency as described in scenario (iv) of Section 3. Thus, we modify A_t to $\bar{A}_t = A_t u_{4,t}$, where $u_{4,t}$ is defined in (12).

In order to compare our algorithm (described in Sections 4.2 and 4.3) to a ‘standard’ algorithm, we need to define what we mean by a ‘standard’ algorithm. We first consider a deterministic version of the model with T periods and, for any t , J values of the discretized state variable $K_j \in \{K_1, \dots, K_J\}$. In this simplified setting without climate, the state variable capital fully dictates the control variable investment, and hence consumption, via the budget constraint.

We start in the final period T . For any K_j , the value function is given by

$$v_T(K_j) = L_T U(A_T K_j^\gamma L_T^{1-\gamma} / L_T).$$

Given a solution in period $t + 1$, the standard algorithm moves one period back to period t and solves for the value function $v_t(\cdot)$ at each gridpoint K_j . This involves selecting the maximum over a finite set:

$$v_t(K_j) = \max_{K_m \in \{K_1, \dots, K_J\}} L_t U(A_t K_j^\gamma L_t^{1-\gamma} / L_t) + \beta v_{t+1}(K_m).$$

We repeat this backward induction until we arrive at period 1. To introduce stochasticity, we draw, for each period and each K_j , S error terms $u_{4,t}^{(s)}$ ($s = 1, \dots, S$) and solve the corresponding step. With the exception of the final period, this requires for each (s, j, t) to take the maximum over J points on the grid (i.e., over the next period’s value function for each possible value of K_m on the grid). For each (j, t) , we compute the estimate of the value

function by averaging across the S draws.

We assume normal errors $\epsilon_{4,t}$, and we let $T = 5$ and $\tau = 0.2$. To implement the standard grid approach we let $200 \leq K_j \leq 2000$ with a stepsize of 20, so that $K_j = 200, 220, \dots, 2000$. For our algorithm we take six polynomial basis functions, like previously. For our algorithm we use 50,000 simulations. For the standard algorithm we distinguish between two cases: coarse (1,000 simulations) and fine (20,000 simulations).

The results are summarized in Figure 3, where we present the estimated value functions of period 4 for four different cases: the standard algorithm for the deterministic model (6.1 seconds), our algorithm for the stochastic model (3.8 minutes), and the standard algorithm for the stochastic model either coarse (4.7 minutes) or fine (1 hour 36 minutes). We find that for about the same amount of computation time, our algorithm is much less volatile for the estimated value function. Using a larger number of simulations leads to a less volatile solution, but at the cost of additional computation time. The computation time is linear in the number of simulations.

This experiment has been conducted in a stylized setting. For a satisfactory solution of the full model as used in this paper, the relative performance of the standard algorithm would deteriorate much further, because we then need to consider one hundred rather than five periods and three rather than two control variables. Furthermore, we would need a much finer grid to accurately capture what happens to utility when consumption becomes small (in the tails) and marginal utility large, which is the interest in our paper.

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 SDICE*, using backward induction.

We know from Sections 2 and 3 that consumption is bounded by

$$\begin{aligned} 0 \leq C_t &\leq C_t + I_t = (1 - \omega_t - \xi H_t^2) A_t K_t^\gamma L_t^{1-\gamma} u_{1,t} \\ &\leq A_t K_t^\gamma L_t^{1-\gamma} u_{1,t} = A_t K_t^\gamma L_t^{1-\gamma} e^{-\tau_1^2/2} e^{\tau_1 \epsilon_{1,t}}, \end{aligned} \tag{17}$$

since $I_t \geq 0$, $\omega_t \geq 0$, and $\xi > 0$. Note that $u_{1,t} > 0$ (with probability one) and that $A_t K_t^\gamma L_t^{1-\gamma}$ is positive for all t . The last statement follows because

the exogenous variables A_t and L_t are assumed to be positive for all t ; the parameter δ satisfies $0 < \delta < 1$; the initial condition $K_1 > 0$ holds; and $K_t \geq (1 - \delta)^{t-1} K_1 > 0$ since $I_t \geq 0$. Now, because A_t and L_t are exogenous, and K_t is deterministic given all information at time $t-1$, since K_t depends on K_{t-1} and I_{t-1} , $A_t K_t^\gamma L_t^{1-\gamma}$ is deterministic given all information at time $t-1$.

Since the social planner in our setup has time-additive expected utility preferences, the inequality (17) 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 $\epsilon_{1,T}$ 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 $\epsilon_{1,T}$ 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 instead of expected marginal utility, but we shall not expand on this. Note that what is relevant here is whether $\epsilon_{1,T}$ (not $\exp(\epsilon_{1,T})$) is heavy-tailed to the left or not. If $\epsilon_{1,T}$ is heavy-tailed rather than light-tailed to the left, then $\exp(\epsilon_{1,T})$ has, loosely speaking, more probability mass near zero.

The fragility of expected power utility to heavy-tailed distributional as-

sumptions 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 SDICE*. The Pareto family, introduced by Ikefuji *et al.* (2013) and given by

$$U(x) = 1 - \left(1 + \frac{x}{\lambda}\right)^{-k} \quad (k > 0, \lambda > 0), \quad (18)$$

enjoys a combination of appealing properties especially relevant in heavy-tailed risk analysis. These properties were primarily normatively motivated but also have some empirical support (see Ikefuji *et al.*, 2013, p. 45). Let

$$\text{ARA}(x) = -\frac{U''(x)}{U'(x)}, \quad \text{RRA}(x) = -\frac{xU''(x)}{U'(x)} \quad (19)$$

denote the local indexes of absolute and relative risk aversion. Under Pareto utility,

$$\text{ARA}(x) = \frac{k+1}{x+\lambda}, \quad \text{RRA}(x) = \frac{x(k+1)}{x+\lambda}, \quad (20)$$

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

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 SDICE*.

6 Main findings

We now have 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 SDICE* model given by scenario (i) 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 Nordhaus' solution to DICE, our solution to DICE, referred to as DICE*, and DICE* under Pareto utility instead of power utility; all in a deterministic setting.

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 SDICE* model, allowing in particular for an uncertain damage parameter, 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 $\epsilon_{j,t}$, $j = 1, \dots, 4$. Following our discussion in Section 3, we consider both a normal distribution (light tails) and a Student distribution (heavy tails). Under normality, the damage-abatement fraction $\bar{d}_t = d_t u_{1,t}$ has a finite expectation. Under a Student distribution, its expectation is infinite. Various routes lead to heavy-tailed distributional assumptions in economy-climate models. For instance, light-tailed distributions for input variables may generate heavy-tailed distributions for output variables via feedback loops (see e.g., Roe and Baker, 2007; and Mahadevan and Deutch, 2010). The heavy-tailed Student distribution we employ can also be interpreted as the posterior predictive distribution of a normal distribution with uncertain standard deviation as suggested in Geweke (2001), Weitzman (2009), and Pindyck (2011).

We need to specify adequate values for the uncertainty parameters τ_j and for the number of degrees of freedom of the Student distribution. Let us consider τ_4 first. The stochasticity generated by $\epsilon_{4,t}$ captures uncertainty about technological efficiency affecting GDP. Historical variation in GDP may therefore serve as a sensible proxy for τ_4 . 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_4 in the set $\{0, 0.02, 0.04, 0.06\}$. Since scenario (iv) is related to scenarios (i), (ii), and (iii) through (6) and (3), and also to make it possible to compare the impact of the four scenarios, we choose the same range for τ_1 – τ_3 . Throughout this section we focus on uncertainty through the damage-abatement fraction ($\tau_1 > 0$, $\tau_2 = \tau_3 = \tau_4 = 0$) and assume the errors to be iid. Other assumptions on τ_j , in particular ($\tau_2 > 0$, $\tau_1 = \tau_3 = \tau_4 = 0$), ($\tau_3 > 0$, $\tau_1 = \tau_2 = \tau_4 = 0$), and ($\tau_4 > 0$, $\tau_1 = \tau_2 = \tau_3 = 0$), 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 2016 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 $\epsilon_{1,t}$, 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 DICE versus DICE*

In a nonstochastic world we find that the optimal policy and state variables under DICE closely match their counterparts under DICE* with power utility. In fact, the maximum absolute difference is 0.1 over the period that we consider; see Tables 1–2. The differences between DICE and DICE* are mainly due to differences in the solution method, the absence of ad hoc bounds ‘for stability’ on some state and control variables in DICE*, and the absence of a heuristic (rather than an optimal) solution for the last ten time periods in DICE*. We also note that abatement at time $t = 1$ is fixed at 0.03 in DICE*, as in DICE.

When we move from power utility to Pareto utility, we see that the optimal policy and state variables under DICE* with Pareto utility match their counterparts under DICE* with power utility quite closely, and that this applies to both the economy and climate parts of the DICE* model. A power utility maximizer consumes less and invests and abates more compared to a Pareto utility maximizer. Indeed, a power utility maximizer has a stronger motive to avoid building up a climate-economy with low damage-abatement reduced output, and therefore uniformly consumes less and invests and abates more.

6.3 Light tails in the center

We now introduce stochasticity and consider the SDICE* model with iid normally distributed errors $\epsilon_{1,t}$ (i.e., light tails), for different values of the degree of uncertainty τ_1 , under both power and Pareto utility. (Recall that

we assume $\tau_j = 0$ for $j = 2, 3, 4$ in this section.)

We focus on the ‘center’ of the distribution by considering shocks along the expected trajectory of $\epsilon_{1,t}$. Specifically, the results reported here are derived by solving for the optimal initial ($t = 1$) policies under uncertainty, and then computing the optimal policies over the following periods 2 to 16 under uncertainty, by assuming that the realized shocks in the previous periods are equal to zero.

The three panels in Table 3 present the results for optimal consumption, investment, and abatement, respectively. Our benchmark is $\tau_1 = 0$, which is the case without uncertainty, that is, DICE*. The introduction of light-tailed uncertainty in the damage-abatement fraction of DICE* leads to a reduction of abatement for both power and Pareto utility. Conditional upon the shocks realizing their expected value, that is zero, we find a reduction in abatement in all periods. Consumption and investment are relatively less affected.

Lower levels of abatement correspond to choosing higher levels of concentration, and this has a negative propagation effect in our model. Conversely, investment has a positive propagation effect. When faced with uncertainty about the damage-abatement fraction, the social planner sacrifices some abatement to maintain investment (first) and consumption (later). In the presence of uncertainty, the power utility maximizer continues to consume less in early periods and invest and abate more in all periods compared to the Pareto utility maximizer.

Overall, the effect of uncertainty on the optimal policies is 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 ‘monotonic’ in the variance of the shock as represented by τ_1 .

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). In particular, we analyze how a realized shock impacts the optimal policies, under future uncertainty. (We don’t consider specific forms of learning whereby the social planner is more ‘on-edge’ because of a recent shock.)

Figures 4 and 5 present optimal consumption C_t^* and optimal abatement μ_t^* as a function of $\epsilon_{1,t}$ at time $t = 1$ and at time $t = 2$, respectively (recall that abatement at time $t = 1$ is fixed), for both the power and Pareto SDICE* models. Considering SDICE* against the benchmark given by DICE* but now allowing the light-tailed shocks to take large negative values, we find that the optimal policy variables are more affected at the edges than in the center. In fact, towards the edges we observe pronounced differences in the optimal policy variables, both within and between the SDICE* models.

As expected, optimal policy derived under certainty — lines with label DICE* — does not respond to negative shocks. A power utility maximizer has a stronger motive to abate as a precaution (cf. Kimball, 1990) compared to a Pareto utility maximizer, and this amplifies with adverse realizations of

$\epsilon_{1,t}$. Under such adverse circumstances, the power utility maximizer keeps abatement at a substantial level, but this comes at the cost of lower consumption. While the presence of uncertainty reduces abatement (as found in Section 6.3), a power utility maximizer puts larger emphasis on keeping up abatement in adverse circumstances compared to a Pareto utility maximizer.

6.5 Heavy tails

Heavy-tailed risk is represented by a Student- t distribution. The random shock $\epsilon_{1,t}$ is not $N(0, 1)$ anymore but rather follows a t -distribution with 10 degrees of freedom, so that $\text{var}(\epsilon_{1,t}) = 1.25$. 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 4 report optimal values in the SDICE* model under Pareto utility for consumption, investment, and abatement, both for light- and heavy-tailed uncertainty, and for different values of τ_1 . 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 a reduction in abatement under damage-abatement fraction uncertainty, 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 ‘monotonic’ in the variance of the shock.

We also report results at the ‘edges’. The changes in the optimal policy

variables between the Pareto utility models with light and heavy tails are virtually identical; see Figure 6. This means that under heavy tails and towards the edges, pronounced differences occur both within and between the power and Pareto models. Contrary to the ‘discontinuity’ that occurs under power utility when we move from light to heavy tails, the Pareto utility maximizer only very slightly adapts his/her optimal policies. In an influential paper, Weitzman (2009) indicated that power utility is fragile with respect to heavy-tailed consumption risk, in the sense that expected marginal utility may become infinite. Our results confirm that this is remedied when preferences are compatible with statistical assumptions, that is, by avoiding an *ex ante* incompatible model specification.

In summary, under heavy tails the main findings of Sections 6.3 and 6.4 broadly remain valid, reinforcing their robustness.

7 Extensions

We generalize the base SDICE* model in three directions. First, we allow for uncertainty in the damage parameter; next, for an uncertain emissions-to-output ratio; and finally, we allow for uncertainty in technological efficiency. The discussion of these three extensions is brief to save space. Details are available from the authors upon request.

7.1 Uncertainty in the damage parameter

In this extension we consider parametric uncertainty by supposing that stochasticity enters SDICE* only through the damage parameter. That is, we

replace ξ in (6) by $\bar{\xi}_t = \xi u_{2,t}$ and suppose $\tau_2 > 0$ while $\tau_1 = \tau_3 = \tau_4 = 0$. We explore how this change in stochasticity impacts our three main findings in Sections 6.3–6.5.

In a setting that mimics Section 6.3, i.e., with $\epsilon_{2,t}$ realizing its expected trajectory, we find that all the optimal policy variables are now nearly insensitive to the presence of uncertainty. The reason is that, under the parameters and levels of uncertainty that we consider, the term $\bar{\xi}_t H_t^2$ contributes relatively little to the budget constraint. Hence, the impact of uncertainty about the damage-abatement fraction analyzed in Section 6.3 dominates the impact of uncertainty about the damage parameter. The power utility maximizer continues to consume less in early periods and to invest and abate more in all periods in comparison to the Pareto utility maximizer.

When tail events manifest themselves analogous to Section 6.4, optimal consumption reduces (increases) when $\epsilon_{2,1}$ takes large positive (negative) values, both under power and Pareto utility. Note that under the present scenario, $\epsilon_{2,t}$ taking large negative values is a prosperous event: it means that the realized damage parameter is low. This is illustrated in Figure 7. Furthermore, optimal abatement is still nearly insensitive even under large realizations of $\epsilon_{2,2}$ (whether positive or negative).

Finally, we explore the impact of heavy-tailed risk associated to the damage parameter, adopting the same distributional assumptions as in Section 6.5. We find that the two previous results are reconfirmed: the near insensitivity of the optimal policy variables in the center and the fact that optimal current consumption is decreasing (increasing) in the extent of the positive (negative) shock in $\epsilon_{2,1}$.

7.2 Emissions-to-output uncertainty

We next suppose that $\tau_3 > 0$ and $\tau_1 = \tau_2 = \tau_4 = 0$, so that uncertainty enters SDICE* only through the emissions-to-output ratio σ_t in (3), and not through the damage-abatement fraction and the budget constraint in (6), as previously. We analyze how our three main findings in Sections 6.3–6.5 are affected under this alternative SDICE* model.

We consider first the optimal policies under iid normally distributed errors $\epsilon_{3,t}$, where we restrict our attention to the center of the distribution by considering realizations of $\epsilon_{3,t}$ 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 prime 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-abatement fraction analyzed previously.

Next, considering the manifestation of tail events analogous to Section 6.4, we find an interesting pattern: while optimal consumption remains relatively insensitive to uncertainty in the emissions-to-output ratio, also under tail scenarios, optimal abatement decreases (increases) pronouncedly when $\epsilon_{3,2}$ takes large negative (positive) values under both power and Pareto utility. This is illustrated in Figure 8.

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 $\epsilon_{3,2}$ 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 again reconfirmed: insensitivity of the optimal policies along the expected trajectory and decreasing (increasing) optimal abatement in the extent of a negative (positive) shock in $\epsilon_{3,2}$.

7.3 Uncertainty in technological efficiency

We finally suppose that $\tau_4 > 0$ and $\tau_1 = \tau_2 = \tau_3 = 0$, which means that uncertainty enters SDICE* through technological efficiency A_t in (1) and hence (6). This implies in particular that uncertainty 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 7.2. 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 $\epsilon_{4,t}$ 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 SDICE* model of Section 6.3. Intuitively, this follows from the insensitivities

of the optimal policies along the expected trajectory observed in Section 7.2 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 $\epsilon_{4,t}$ takes large negative values analogous to the analysis at the edges in Section 6.4, optimal consumption responds exactly as in Figure 4. However, optimal abatement, decreases with the extent of the negative shock $\epsilon_{4,2}$ for both power and Pareto utility; see Figure 9. The latter effect is in part induced by the abatement results in Section 7.2, illustrated in Figure 8.

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 scenario (i).

Apparently, the impact of uncertainty is similar whether we model it through its impact in the damage-abatement fraction or through technological efficiency. The key difference between the base SDICE* model and our third extension is that, while uncertainty in the damage-abatement fraction increases optimal abatement for the power utility maximizer, this effect is suppressed in adverse technology scenarios in which budgets and hence emissions are 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 (SDICE*) is based on

Nordhaus’ deterministic DICE model, but it incorporates, possibly heavy-tailed, stochasticity. Upon applying our solution method to SDICE* 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, depending on the nature and extent of the uncertainty and the social planner’s preferences.

A general criticism about integrated assessment models, whether deterministic or stochastic, is that some of the inputs such as functional forms and parameters are to some extent arbitrary, and yet quite relevant for the outputs they predict. This also applies to canonical benchmark integrated assessment models such as Nordhaus’ DICE model. SDICE* provides a framework to assess the effects of a rich spectrum of uncertainties related to the complex interplay between climate and the economy on the optimal consumption, investment and abatement policies. While we have made a significant effort to calibrate our parameters to resemble e.g., output growth rates generated by our model, we prefer to interpret our results not as direct literal policy implications, but rather as implications about the extent and sensitivity of the interactions between the main variables of interest and about the roles played by the various variables, dynamic equations, functional forms, and parameters.

The regression-based method we develop for solving a general class of stochastic dynamic finite-horizon optimization problems relies on forward-simulated paths and avoids nested simulation. This makes our method relatively efficient and we find that it competes favorably with standard solution approaches. To further improve efficiency, future work may consider econo-

mizing on simulated paths by selecting fewer but better nodes.

If we combine heavy-tailed uncertainty about climate change and its economic damage with an arbitrary utility function, such as the conventional power utility function, we may be confronted with infinite expected marginal utility. In that case, the model would predict that the social planner should reduce current consumption to ultra-low levels, in order to limit the possibility of an economy-climate catastrophe. The resolution to this unacceptable conclusion in cost-benefit analysis is to impose compatibility conditions on beliefs and preferences.

Our results show that introducing light-tailed uncertainty to a conventional economy-climate model with power utility can yield reduced levels of current consumption, driven by a strong desire to save us from an economy-climate catastrophe. This can be especially so under very adverse scenarios, in which a relatively strong substitution from current consumption to precaution occurs. This effect is more limited under Pareto utility. These findings remain intact, and get exacerbated, under heavy tails.

Conventional integrated assessment models can thus overemphasize precautionary action when they are confronted with uncertainty and heavy tails.

Acknowledgements

We are very grateful to the editor of the Journal of Econometrics, to the editors of this special issue, and to three referees for their constructive comments and suggestions that have significantly improved our paper. We are also 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 feedback. 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).

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., 2017. Managing catastrophic climate risks under model uncertainty aversion. *Management Science* 63, 749–765.
- 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 sim-

- ulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies* 18, 831–873.
- Brumm, J., Scheidegger, S., 2017. Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica* 85, 1575–1612.
- 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., 2012. DSICE: A dynamic stochastic integrated model of climate and economy. RDCEP Working Paper No. 12-02. Available at SSRN: <https://ssrn.com/abstract=1992674>.
- Cai, Y., Lenton, T. M., Lontzek, T. S., 2016. Risk of multiple interacting tipping points should encourage rapid CO₂ 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.
- Howard, P. H., Sterner, T., 2017. Few and not so far between: a meta-analysis of climate damage estimates. *Environmental Resource Economics* 68, 197–225.
- 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., 2014a. Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review* 69, 104–125.
- Jensen, S., Traeger, C. P., 2014b. 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.
- King, R. G., Rebelo, S. T., 1999. Resuscitating real business cycles, in: *Handbook of Macroeconomics*, Vol. 1 (Eds. J. B. Taylor and M. Woodford), Elsevier, Chapter 14.

- Krätschmer, V., Ladkau, M., Laeven, R. J. A., Schoenmakers, J. G. M., Stadjé, M., 2017. Optimal stopping under uncertainty in drift and jump intensity. *Mathematics of Operations Research*, forthcoming.
- Laeven, R. J. A., Stadjé, 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.
- Mahadevan, L., Deutch, J. M., 2010. Influence of feedback on the stochastic evolution of simple climate systems. *Proceedings of the Royal Society A* 466, 993–1003.
- 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., 2017a. Revisiting the social cost of carbon. *Proceedings of the U. S. National Academy of Sciences* 114, 1518–1523, February 14.
- Nordhaus, W. D., 2017b. Projections and uncertainties about climate change in an era of minimal climate policies. No. 22933, National Bureau of Economic Research.
- 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.
- Roe, G. H., Baker, M. B., 2007. Why is climate sensitivity so unpredictable? *Science* 318, 629–632.
- Roughgarden, T., Schneider, S. H., 1999. Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* 27, 415–429.
- Schmitt-Grohé, S., Uribe, M., 2004. Optimal fiscal and monetary policy under sticky prices. *Journal of Economic Theory* 114, 198–230.
- 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.

Figures and Tables

Figure 1: Akaike's Information Criterion (AIC) for the sets of basis functions up to polynomial order 5 against the time period

Figure 2: Akaike's Information Criterion (AIC) for the sets of cross terms against the time period

Figure 3: Comparison of standard algorithm to our algorithm

Figure 4: Consumption C_1 : SDICE* with normal errors and $\tau_1 = 0.00, 0.02, 0.04$, and 0.06 — power vs Pareto, scenario (i)

Figure 5: Abatement μ_2 : SDICE* with normal errors and $\tau_1 = 0.02, 0.04$, and 0.06 — power vs Pareto, scenario (i)

Figure 6: Abatement μ_2 : SDICE* under Pareto utility and $\tau_1 = 0.02, 0.04$, and 0.06 — normal (solid) vs Student (dotted), scenario (i)

Figure 7: Consumption C_1 : SDICE* with normal errors and $\tau_2 = 0.02, 0.04$, and 0.06 — power vs Pareto, scenario (ii)

Figure 8: Abatement μ_2 : SDICE* with normal errors and $\tau_3 = 0.02, 0.04$, and 0.06 — power vs Pareto, scenario (iii)

Figure 9: Abatement μ_2 : SDICE* with normal errors and $\tau_4 = 0.02, 0.04$, and 0.06 — power vs Pareto, scenario (iv)

Table 1: State variables — DICE vs DICE*, power vs Pareto						
Control variable	Model	Utility	2015	2040	2065	2090
Capital (K_t/Y_t)	DICE	power	2.120	2.429	2.457	2.488
	DICE*	power	2.120	2.328	2.421	2.491
	DICE*	Pareto	2.120	2.037	2.128	2.228
Concentration (M_t/Y_t)	DICE	power	8.091	4.938	3.108	2.084
	DICE*	power	8.091	4.943	3.035	2.018
	DICE*	Pareto	8.091	5.342	3.378	2.331
Temperature (H_t)	DICE	power	0.850	1.694	2.525	3.247
	DICE*	power	0.850	1.673	2.460	3.148
	DICE*	Pareto	0.850	1.698	2.555	3.361

Table 2: Control variables — DICE vs DICE*, power vs Pareto						
Control variable	Model	Utility	2015	2040	2065	2090
Consumption (C_t/Y_t)	DICE	power	0.738	0.819	0.798	0.778
	DICE*	power	0.762	0.824	0.798	0.777
	DICE*	Pareto	0.799	0.862	0.830	0.803
Investment (I_t/Y_t)	DICE	power	0.317	0.315	0.291	0.277
	DICE*	power	0.277	0.311	0.293	0.280
	DICE*	Pareto	0.217	0.275	0.265	0.256
Abatement cost fraction (ω_t)	DICE	power	0.000	0.002	0.004	0.007
	DICE*	power	0.000	0.003	0.005	0.007
	DICE*	Pareto	0.000	0.001	0.001	0.002

Table 3: SDICE* with normal errors — power vs Pareto, scenario (i)

$t \backslash \tau_1$	power				Pareto			
	0.00	0.02	0.04	0.06	0.00	0.02	0.04	0.06
<i>Consumption C_t</i>								
2015	80.18	80.26	80.31	80.34	84.02	84.06	84.07	84.06
2040	169.00	168.86	168.72	168.57	166.86	166.89	166.94	167.00
2065	306.44	307.13	307.75	308.38	301.58	301.88	302.29	302.73
2090	491.41	491.00	490.40	489.83	484.42	484.29	484.24	484.25
<i>Investment I_t</i>								
2015	29.15	29.02	28.93	28.88	22.78	22.70	22.69	22.70
2040	63.82	63.93	64.14	64.39	53.22	53.18	53.21	53.31
2065	112.48	111.89	111.57	111.33	96.19	95.97	95.84	95.80
2090	177.25	177.46	178.03	178.69	154.34	153.94	153.91	154.02
<i>Abatement μ_t</i>								
2015	0.030	0.030	0.030	0.030	0.030	0.030	0.030	0.030
2040	0.361	0.357	0.354	0.351	0.205	0.204	0.202	0.200
2065	0.523	0.517	0.511	0.506	0.323	0.318	0.312	0.307
2090	0.719	0.717	0.714	0.711	0.474	0.469	0.464	0.460

Table 4: SDICE* under Pareto utility — light- vs heavy-tailed, scenario (i)

	$\tau_1 = 0.02$		$\tau_1 = 0.04$		$\tau_1 = 0.06$	
	light	heavy	light	heavy	light	heavy
<i>Consumption C_t</i>						
2015	84.06	84.06	84.07	84.07	84.06	84.06
2040	166.89	166.89	166.94	166.94	167.00	167.00
2065	301.88	301.88	302.29	302.29	302.73	302.73
2090	484.29	484.29	484.24	484.24	484.25	484.25
<i>Investment I_t</i>						
2015	22.70	22.70	22.69	22.69	22.70	22.70
2040	53.18	53.18	53.21	53.21	53.31	53.31
2065	95.97	95.97	95.84	95.84	95.80	95.80
2090	153.94	153.94	153.91	153.92	154.02	154.02
<i>Abatement μ_t</i>						
2015	0.030	0.030	0.030	0.030	0.030	0.030
2040	0.204	0.204	0.202	0.202	0.200	0.200
2065	0.318	0.317	0.312	0.312	0.307	0.306
2090	0.469	0.469	0.464	0.464	0.460	0.459