Economics Letters 120 (2013) 552–558

Contents lists available at SciVerse ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

# Optimal climate policy: Uncertainty versus Monte Carlo[☆]

Benjamin Crost[a], Christian P. Traeger[b,*]

[a] *Department of Economics, UC Denver, United States*

[b] *Department of Agricultural & Resource Economics, UC Berkeley, United States*

## HIGHLIGHTS

- How well can Monte-Carlo averaging of deterministic scenarios replicate optimal climate policy under uncertainty?
- Answer 1: Quantitatively off.
- Answer 2: Can imply the wrong sign of the uncertainty effect.
- Answer 3: Can imply contradictory recommendations that depend on the depicted policy variable.
- Results hold for standard preferences as well as comprehensive risk preferences.

## ARTICLE INFO

*Article history:*
Received 17 January 2013
Received in revised form
3 May 2013
Accepted 17 May 2013
Available online 29 May 2013

*JEL classification:*
Q54
Q00
D61
D90
C61
C63

*Keywords:*
Climate change
Uncertainty
Integrated assessment
Monte Carlo
Risk aversion
DICE

## ABSTRACT

The integrated assessment literature frequently replicates uncertainty by averaging Monte Carlo runs of deterministic models. This Monte Carlo analysis is, in essence, an averaged sensitivity analyses. The approach resolves all uncertainty before the first time period, drawing parameters from a distribution before initiating a given model run. This paper analyzes how closely a Monte Carlo based derivation of optimal policies is to the truly optimal policy, in which the decision maker acknowledges the full set of possible future trajectories in every period. Our analysis uses a stochastic dynamic programming version of the widespread integrated assessment model DICE, and focuses on damage uncertainty. We show that the optimizing Monte Carlo approach is not only off in magnitude, but can even lead to a wrong sign of the uncertainty effect. Moreover, it can lead to contradictory policy advice, suggesting a more stringent climate policy in terms of the abatement rate and a less stringent one in terms of the expenditure on abatement.



## 1. Introduction

The precise nature and extent of climate change and its socio-economic consequences remain uncertain. Interest groups as well as political leaders frequently cherry pick information. Economists should help to avoid polarization by explicitly incorporating uncertainty into their models. In the integrated assessment community, uncertainty is usually emulated by averaging deterministic sensitivity runs. The community refers to such weighted averaging as a Monte Carlo approach. The approach is convenient because it employs deterministic models. However, these models are not seeking optimal policies, because decisions are not actually taken under uncertainty. Seeking an optimal policy, the decision maker has to find one optimal decision at a given information set anticipating all possible futures.

We compare the optimal policy under damage uncertainty to the probability weighted average of deterministic runs under different damage draws. We use an integrated assessment model that recently shaped US climate policy. We show that deterministic averaging can result in contradictory policy recommendations:

☆ We thank Larry Karp for helpful advice. Financial support from the Giannini Foundation and the National Science Foundation is gratefully acknowledged.

* Correspondence to: Department of Agricultural & Resource Economics, 207 Giannini Hall #3310, University of California, Berkeley, CA 94720-3310, United States.

*E-mail address:* traeger@berkeley.edu (C.P. Traeger).

0165-1765/$ – see front matter © 2013 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.econlet.2013.05.019

depending on whether the modeler averages the abatement rate or the abatement expenditure, he/she finds that uncertainty significantly decreases or increases the optimal mitigation effort. Moreover, under damage uncertainty, an optimizing Monte Carlo approach underestimates the optimal carbon tax and significantly overestimates the optimal peak carbon concentration. We find examples where even the sign of the policy change implied by introducing uncertainty can be wrong under an averaging of deterministic paths.

The recent literature addressing uncertainty in climate change comprises three strands. First, highly abstract models discuss whether uncertainty is important for climate change evaluation (Weitzman, 2010; Millner, 2011; Traeger, 2012b). While these models yield important insights into the fundamentals of the problem, they are too abstract to yield quantitative information on how uncertainty affects optimal climate policy. Second, Monte Carlo studies simulate deterministic integrated assessment models for different parameter draws and average the results to resemble uncertainty (Richels et al., 2004; Hope, 2006; Nordhaus, 2008; Dietz, 2009; Anthoff et al., 2009; Ackerman et al., 2010; Interagency Working Group on Social Cost of Carbon, 2010; Pycroft et al., 2011; Kopp et al., 2012). The most popular approach does not try to optimize policies, but merely simulates the effect of uncertainty under an exogenously given policy. Often, these assessments cite the social cost of carbon (SCC). This social cost of carbon is merely the damage of a ton of carbon along a non-optimal path. It is not the optimal carbon tax. The derivation of an optimal mitigation policy takes into account that a change in the policy changes the emissions, the carbon stock, the temperature, the damages, the economic capital, and, thus, the value of and damages from an additional ton of carbon. Few approaches, for example in DICE-2007 Ackerman et al. (2010), have tried to get at optimal policies using Monte Carlo analysis. However, the Monte Carlo approach resolves all uncertainty by drawing a parameter before a given model run is initiated. Then, an optimizing decision maker can set his/her policy optimally for each state of the world. He/she is only uncertain right before running his/her model: he/she is only uncertain ex ante. In contrast, a truly uncertain policy maker has to identify a single optimal policy not knowing the true state of the world. A third strand of literature properly handles uncertainty in complex integrated assessment models (Kelly and Kolstad, 1999; Keller et al., 2004; Leach, 2007). We follow this literature, using a recursive approach similar to Kelly and Kolstad (1999), but model uncertainty over the damage function and compare the optimal trajectories to those of an optimizing Monte Carlo approach.

We use a recursive dynamic programming implementation of Nordhaus's (2008) DICE model as described in detail in Traeger (2012a). In contrast to the original DICE model, our implementation features persistent uncertainty, an annual time step, and an infinite planning horizon. In order to limit the run-time and to improve the control rule approximations, our model calibrates an exogenous carbon decay rate to the carbon cycle and collects different warming delays into a simplified delay equation. The literature has criticized various details of the DICE model, including its damage function (Hanemann, 2009). Nevertheless, the DICE model remains a benchmark; it is the most widespread model in the literature, and it was recently used as one of three models that informed the official estimate of the SCC to be used for evaluating policy in the US (Interagency Working Group on Social Cost of Carbon, 2010). DICE is open source, easily accessible, and has received continuous updates and improvements over the last 20 years. We focus on damage uncertainty, paying particular attention to the coefficient uncertainty suggested in Nordhaus's (2008) Monte Carlo analysis. The latter damage coefficient states the damages for a 1 °C warming over 1900 levels as a fraction of world output. We show that similar results hold for the case of uncertainty over the damage exponent, which determines how quickly damages increase with rising temperatures. Moreover, Crost and Traeger (2010) show that a comprehensive risk attitude model is crucial for evaluating climate change. The economic standard model assumes that risk aversion equals the propensity to smooth consumption over time. In consequence, the calibration of DICE to market interest can only calibrate either the risk premia or the risk-free discount rate correctly. An additional scenario therefore follows Crost and Traeger (2010) in using Epstein–Zin preferences and parameter estimates from the finance literature resolving the risk-free rate and the equity premium puzzles. We show that similar differences between the Monte Carlo approach and the fully stochastic model also emerge in the more comprehensive model of risk and risk attitude.

## 2. The model

The first part of this section discusses those parts of the DICE model that are most relevant to our subsequent analysis. Then, we introduce uncertainty, and briefly describe its numerical implementation. Subsequently, we state the welfare side of the model and the resulting Bellman equation. We close with a generalized DICE Bellman equation matching market interest and discount rates better than the standard model.

### 2.1. DICE and climate damages

Nordhaus's (1994,2008) DICE model couples a Ramsey–Cass–Koopmans exogenous growth economy to a model of the climate system. The economy produces emissions that accumulate in the atmosphere, change the radiative forcing in the atmosphere, and warm the planet's surface. This warming feeds back into economic production and consumption. Except for emission accumulation, all of these interactions are nonlinear. In addition, warming is subject to a significant delay because of the atmosphere's and the ocean's heat capacity as well as various feedback processes. The original DICE model has a time step of 10 years. In our recursive dynamic programming problem, an annual time step comes at no additional cost.[1] We adopt an annual time step not only because it gives a better resolution of optimal policy, but also because we focus on the effects of risk aversion and intertemporal substitution. A ten-year time step changes the intertemporal and the risk fluctuations significantly, making questionable the validity of attitude parameters that are usually derived on an annual basis.[2] An accompanying paper explains the details of the recursive dynamic programming implementation of the DICE model and its scaling to a one-year time step (Traeger, 2012a).[3]

The main difference between the economic side of DICE and the standard Ramsey growth model is a wedge between gross production and net production. We use a per effective unit of labor

[1] This finding holds because we make time a state variable and fit the value function continuously on the full state space. In contrast, in the usual forward optimization, decreasing a ten-year time step to one year increases the time-indexed control and state variables by the factor 10.

[2] A ten-year time step implies fewer, but much larger jumps. The risk and consumption smoothing parameters are not easily adjusted to the changed dynamics when the model is outside of a steady state, which is clearly the case for the relevant centuries of the DICE model.

[3] Simply scaling the DICE model to an annual time step would imply two changes affecting optimal policies. First, a finer policy resolution can reduce the SCC, which is a small effect. Second, a simple downscaling of the carbon cycle would result in a faster effective decay of atmospheric carbon. This second effect would be physically wrong and requires a recalibration of the carbon cycle. We decided to calibrate the model to reproduce exactly the optimal policies of the deterministic DICE baseline in order to focus exclusively on uncertainty effects.