Journal of Environmental Economics and Management 72 (2015) 98–122

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

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

Learning and climate feedbacks: Optimal climate insurance and fat tails

David L. Kelly [a,*], Zhuo Tan [b]

[a] Department of Economics, University of Miami, Box 248126, Coral Gables, FL 33146, United States
[b] Department of Economics, Hong Kong Baptist University, Hong Kong

ARTICLE INFO

Article history:
Received 3 March 2014
Available online 10 June 2015

JEL classification:
Q54
Q58
D83
D81
H43

Keywords:
Climate change
Climate insurance
Learning
Fat tails

ABSTRACT

We study the effect of potentially severe climate change on optimal climate change policy, accounting for learning and uncertainty in the climate system. In particular, we test how fat upper tailed uncertainty over the temperature change from a doubling of greenhouse gases (the climate sensitivity), affects economic growth and emissions policy. In addition, we examine whether and how fast uncertainties could be diminished through Bayesian learning. Our results indicate that while overall learning is slow, the mass of the fat tail diminishes quickly, since observations near the mean provide evidence against fat tails. We denote as "tail learning" the case where the planner rejects high values of the climate sensitivity with high confidence, even though significant uncertainty remains. Fat tailed uncertainty without learning reduces current emissions by 38% relative to certainty, indicating significant climate insurance, or paying to limit emissions today to reduce the risk of very high temperature changes, is optimal. However, learning reduces climate insurance by about 50%. The optimal abatement policy is strongly influenced by the current state of knowledge, even though greenhouse gas (GHG) emissions are difficult to reverse. Once the mass of the fat tail diminishes, the remaining uncertainty is largely irrelevant for optimal emissions policy.

Introduction

Uncertainty is a dominant feature of climate change. Recent research highlights a particular aspect of climate change uncertainty: a relatively small chance of severe climate change exists. In particular, a doubling of greenhouse gases (GHGs) above preindustrial levels may cause a very large steady state increase in temperature.[1] The sensitivity of the temperature to GHG concentrations is known as the climate sensitivity. Uncertainty about the climate sensitivity creates an insurance motive for reducing emissions, in that paying to limit GHG emissions today prevents GHG concentrations from rising, which reduces the

* Corresponding author. Fax: +1305 284 2985.
E-mail addresses: dkelly@miami.edu (D.L. Kelly), zhuotan@hkbu.edu.hk (Z. Tan).

[1] Stocker et al. (2013) reviews many studies and finds values higher than 4.5 °C cannot be ruled out, and values greater than 6 °C can only be ruled out with "medium confidence", although a value between 1.5 and 4.5 °C is "likely". Weitzman (2009b) averages 22 studies and finds a 5% chance that a doubling of GHGs will cause temperatures to rise more than 7 °C. Other papers which estimate the current scientific uncertainty regarding the climate sensitivity include: Lemoine (2010), Newbold and Daigneault (2009), Roe and Baker (2007), Schwartz (2007), and Baker and Roe (2009).

http://dx.doi.org/10.1016/j.jeem.2015.05.001
0095-0696/

probability of very high temperature changes.[2] The prior distribution of the climate sensitivity is known to have a fat upper tail, meaning the upper tail of the distribution of temperature changes declines at a rate slower than exponential. The existence of a fat tail significantly increases the insurance value of current GHG abatement, since households are willing to pay more up front abatement costs to eliminate fat tailed risk of severe climate change.[3]

However, climate change uncertainty differs from a standard insurance problem in that learning reduces uncertainty over time. If learning resolves climate uncertainty relatively quickly, then the initial need for climate insurance is small, as the planner still has time to increase abatement if learning quickly indicates the climate sensitivity is large. However, if learning resolves climate uncertainty slowly, then the optimal policy calls for aggressive initial abatement for insurance purposes. The central questions for climate policy are then: how fast will learning resolve fat tailed uncertainty about the climate sensitivity, and what is the optimal climate policy with fat tailed climate uncertainty and learning?

The prior literature finds that learning is a slow process. Kelly and Kolstad (1999b) consider uncertainty regarding the heat capacity of the ocean. In their integrated assessment model,[4] stochastic weather shocks obscure the climate change signal in the temperature data, which slows Bayesian learning. This result has since been confirmed in models with other types of climate uncertainty and different distributional assumptions. In particular, Leach (2007) considers uncertainty over the climate sensitivity and finds Bayesian learning about the climate sensitivity is extremely slow.[5] Roe and Baker (2007) argue that resolving uncertainty regarding the climate sensitivity is difficult, because small uncertainties in climate feedbacks magnify the uncertainty about the climate sensitivity.[6] Keller et al. (2004) show that slow learning about the climate sensitivity combined with an uncertain climate threshold, implies significant near term abatement is optimal, to avoid accidentally exceeding the threshold. Lemoine and Traeger (2014) study alternative uncertain thresholds with learning and find somewhat smaller effects on near term abatement. The above literature on learning with thin tails finds that learning is too slow to have much impact initially: near term abatement policy is similar with or without learning.

However, it is possible that the planner learns enough to reject severe climate change with a high degree of confidence quickly, even though the climate sensitivity is difficult to pin down precisely.[7] We define this case as “tail learning.”

To investigate tail learning, we develop a quantitative integrated assessment model in which the planner faces stochastic weather shocks and uncertainty over the first order autoregressive coefficient in the equation governing the evolution of temperature, the climate feedback parameter. Because the climate feedback parameter is uncertain, the climate sensitivity is also uncertain. If the climate feedback parameter is close to one, then GHG “shocks” to temperature are long lived, and an increase in GHG emissions causes high steady state temperature changes. Hence, thin tailed uncertainty in the feedback parameter causes fat tailed uncertainty in the climate sensitivity. Roe and Baker (2007) show that normally distributed uncertainty in the feedback parameter results in climate sensitivity uncertainty which approximates the current uncertainty in the scientific literature. The social planner learns the feedback parameter, and therefore the climate sensitivity, using Bayes rule.

We define the lower bound of the fat tail as when a doubling of GHGs implies steady state temperatures increase by 1.5 °C more than the mean of the current prior distribution. For example, we calibrate that the mean of the initial prior results in a temperature increase of 2.76 °C, so the tail of the distribution equals values greater than or equal to 4.26 °C.[8] When the planner rejects the hypothesis that the climate sensitivity implies a steady state temperature increase greater than or equal to the lower bound at the 1% or 0.1% level, we say that tail learning is complete.[9] Such learning is partial in the sense that significant uncertainty typically remains even after a high climate sensitivity is rejected.

Our results show that the social planner rejects that the climate sensitivity is in the upper tail of the prior distribution very quickly. That is, although we confirm results in the previous literature that learning the actual true value *precisely* is a relatively slow process, the planner rejects values of the climate sensitivity in the upper tail of the prior distribution quickly. In fact, tail learning is complete in less than a decade, if the true climate sensitivity is moderate. First, observations near the moderate true value provide evidence against the tail of the distribution. In addition, the density of even a fat tail is not

[2] We are referring here to the definition of insurance as protection against a possible adverse outcome, not the purchase of a contract which provides compensation in the event of a loss.

[3] Indeed, fat tailed uncertainty is now arguably the most important current issue in climate change policy, since it is the dominant motivation (either directly or indirectly through the effect of uncertainty on the discount rate) for stringent immediate limits on GHG emissions (see, for example, Weitzman, 2007).

[4] An integrated assessment model is broadly defined as a model which combines scientific and socio-economic aspects of climate change to assess policy options for climate control (Kelly and Kolstad, 1999a).

[5] Most of the literature and our paper consider observational learning in the sense that the planner learns from the data on temperature and GHG concentrations. An alternative is to allow learning where the planner pays for R&D. Nonetheless, to fully resolve uncertainty, all R&D must eventually be confirmed in the data.

[6] Climate feedbacks are changes in the climate system brought on by higher temperatures which amplify or diminish the relationship between GHGs and temperature (climate forcing). For example, higher temperatures melt ice, which in turn implies less heat is radiated back into space, which amplifies climate forcing. The magnitude of many climate feedbacks are uncertain (Forest et al., 2006).

[7] Strictly speaking, the planner takes the entire current and expected future distribution of uncertainty into consideration when calculating the optimal policy. That is, the planner does not statistically test the hypothesis that climate change is severe. Instead, we show the optimal abatement policy can be intuitively understood as one in which the test statistic plays a dominant role.

[8] No generally agreed upon value for what constitutes the tail of the prior distribution exists. Nonetheless, much of the literature uses higher values (e.g. the Weitzman (2009b) discusses the values above 7 °C). A larger lower bound would only strengthen our results.

[9] See Kelly and Kolstad (1999b) for a justification for using hypothesis tests to measure learning.