Fat-tailed uncertainty and the learning-effect

Hwang, In Chang

14 February 2014

Fat-tailed Uncertainty and the Learning-effect

In Chang Hwang

VU University Amsterdam, Institute for Environmental Studies, Amsterdam, The Netherlands

De Boelelaan 1087, Amsterdam, The Netherlands, 1081 HV, E-mail: i.c.hwang@vu.nl

Abstract

One of the recent findings in the economics of climate change is that emissions control plays a significant role in the reduction of the tail-effect of fat-tailed uncertainty on welfare. The current paper gives another perspective: the learning-effect. The effect of emissions control on welfare is decomposed into the direct effect and the learning-effect. Although this has been known for thin-tailed uncertainty in the literature, this paper takes a different approach: the changes in temperature distributions under fat-tailed uncertainty and learning.

Key words

Climate policy; deep uncertainty; Dismal Theorem; tail-effect; learning-effect

JEL Classification

Q54

1 Introduction

One of the recent findings in the economics of climate change is that optimal carbon tax does not accelerate for many plausible situations as the uncertainty about climate change increases (Karp, 2009; Hwang et al., 2013a, b; Millner, 2013; Horowitz and Lange, 2014), as implied by Weitzman's Dismal Theorem (Weitzman, 2009). This is because emissions control, together with investment, which was absent in the model of Weitzman (2009), plays a significant role in the reduction of the effect of fat-tailed uncertainty on welfare (the tail-effect).

Beside its direct impact, emissions control has an implicit impact on welfare in that carbon emissions produce information on the true state of the world through increased warming. Since learning or the (partial) resolution of uncertainty has value, this should be accounted for when the decision on emissions control is made.

The hypothesis of this paper is that the possibility of learning reduces the marginal benefits of emissions control, compared to the case where there is no learning. As a result, learning reduces the stringency of climate policy compared to the no-learning case. Although this has been known for thin-tailed uncertainty in the literature (e.g., Kolstad, 1996a, b; Ulph and Ulph, 1997; Kelly and Kolstad, 1999; Webster, 2002; Ingham et al., 2007), this paper takes a different approach: the changes in temperature distributions under fat-tailed uncertainty. This approach is also taken by Pindyck (2011, 2012) and Millner (2013), but their models are too stylized and furthermore do not account for learning.

Section 2 presents the model. Sections 3 and 4 investigate the changes in temperature distribution and the rate of tail-slimming (Weitzman, 2013), respectively. The main results are given in Section 5. Section 6 concludes.

2 The Model

The learning-effect is discussed in this paper with a simple dynamic model as in Equation (1). The problem of a decision maker is to choose the rate of emissions control in each period so as to maximize social welfare defined as the discounted sum of expected utility of consumption.

$$W_t = \max_{\mu_t \in [0,1]} U(c_t) + \beta \mathbb{E}_t W_{t+1} = U(c_t) + \beta \int_{\{\lambda\}} W_{t+1} \cdot g_{\lambda_t} d\lambda \tag{1}$$

where W is social welfare, t is time period, μ is the rate of emissions control, U is the utility function, β is the discount factor, $\mathbb{E}$ is the expectation operator, c is consumption per capita, $\lambda \in \{\lambda\}$ is an uncertain parameter such as the equilibrium climate sensitivity, $\{\}$ is the set of any variable, and g_λ is the probability distribution function (PDF) of λ.

A unit increase in carbon emissions induces higher temperature in the future through Equations (2-4).

$$m_{t+1} = (1 - \mu_t) \cdot \sigma \cdot y_t + (1 - \delta) \cdot m_t \tag{2}$$

$$RF_{t+1} = \varphi \cdot ln(m_{t+1}) \tag{3}$$

$$T_{t+1} = (RF_{t+1}/RF_0) \cdot \lambda \tag{4}$$

where m is the carbon stock, σ is the emission-output ratio, y is gross output per capita, δ is the depreciation rate of the carbon stock, RF is radiative forcing from the carbon stock, RF_0 is radiative forcing from a doubling of carbon dioxide, T is air temperature deviations from the initial period, and φ is a constant.

Equation (4) says that a doubling of carbon dioxide induces a temperature increase of λ, the equilibrium climate sensitivity. Without loss of generality it is assumed that $m_1 = 0$, $RF_1 = 0$, and $T_1 = 0$ in this paper.

Carbon emissions reduce expected social welfare due to the loss of consumption as a consequence of adverse climate change (Equation 5). Thus the decision maker tries to control the amount of carbon emissions. Emissions control comes at a cost as in Equation (6).

$$c_t = y_t/(1 + \tau {T_t}^{\gamma}) - \Lambda_t \tag{5}$$

$$\Lambda_t = \theta_1 \cdot {\mu_t}^{\theta_2} \tag{6}$$

where Λ is the abatement cost function, $\tau\,(>0)$, $\theta_1\,(>0)$, $\theta_2\,(>1)$ and $\gamma\,(>1)$ are economic parameters. For simplicity y is normalized to be one.

Equations (2-6) and the conditions for the parameter values are generally consistent with the literature (e.g., Gregory and Forster, 2008; Nordhaus, 2008; Weitzman, 2012).

The hyperbolic absolute risk aversion (HARA) utility function is applied in this model. Note that the constant relative risk aversion (CRRA) utility function, usually used in the literature, is a special case of HARA. If $\eta = 0$ (7) becomes CRRA.

$$u(c_t) = \zeta \cdot \{\eta + c_t/\alpha\}^{1-\alpha} \tag{7}$$

where ζ, $\eta(\geq 0)$, and $\alpha(>0)$ are parameters. It is assumed that $\zeta\,(1 - \alpha)/\alpha > 0$ for utility to be increasing and concave in consumption.

The uncertain parameter is assumed to have a fat-tailed distribution in the sense that probability density diminishes slowly than exponentially in the upper tail (Weitzman, 2009; Pindyck, 2011). In this paper the distribution of Roe and Baker (2007) is applied, which is widely discussed in the literature (e.g., Weitzman, 2009; Millner, 2013). Applying the other fat-tailed distributions does not affect the general findings of this paper, as argued in Hwang et al. (2013). The notation for time is dropped for convenience, unless otherwise confused.

$$g_\lambda = \frac{1}{\sigma_f\sqrt{2\pi}} \cdot \frac{\lambda_0}{\lambda^2} \cdot exp\left\{-\frac{1}{2}\left[\frac{\left(1-\bar{f}-\frac{\lambda_0}{\lambda}\right)}{\sigma_f}\right]^2\right\} \tag{8}$$

where $\bar{f}$ and σ_f are the parameters of the climate sensitivity distribution, λ_0 is a constant.

To solve the learning model, the random variable is transformed to derive the PDF of temperature increases as follows.

$$g_T = g_\lambda \cdot \frac{\partial\lambda(T)}{\partial T} = \frac{1}{\sigma_f\sqrt{2\pi}} \cdot \frac{RF}{RF_0} \cdot \frac{\lambda_0}{T^2} exp\left\{-\frac{1}{2}\left[\frac{1-\bar{f}-\frac{RF}{RF_0}\cdot\frac{\lambda_0}{T}}{\sigma_f}\right]^2\right\} \tag{9}$$

where g_T is the PDF of future temperature increases and $\lambda(T) = (RF_0/RF) \cdot T$. Note that g_T is a fat-tailed distribution in that $\lim_{T\to\infty} g_T/exp(-aT) = \infty$ for any $a > 0$.

3 Temperature Distribution and Tail-slimming Rate

The parameters σ_f and $\bar{f}$ are subject to change in the learning model. For instance, the belief of the decision maker on the climate sensitivity can be updated as information (e.g., temperature records) accumulates (Kelly and Kolstad, 1999). In the Bayesian statistics literature, with a normal likelihood function, it is well known that the posterior mean tends to the (pre-specified) true value and the posterior variance approaches zero asymptotically over time (Cyert and DeGroot, 1974). With this in mind and without loss of generality, it is assumed that $\partial \bar{f}/\partial \mu = 0$. The general findings of this paper are not affected by this assumption. In addition, following the literature on learning (e.g., Kelly and Kolstad, 1999; Leach, 2007) it is assumed that $\partial \sigma_f/\partial RF < 0$. For instance, Hwang et al. (2014) find that ${\sigma_f}_{t+1}^2 = {\sigma_f}_t^2/\left(1+\zeta_1^2 T^2\, {\sigma_f}_t^2/\sigma_\varepsilon^2\right)$, where σ_ε^2 is the variance of temperature shocks and ζ_1 is a parameter. Since $\partial T/\partial RF > 0$, it is clear that $\partial \sigma_f/\partial RF = (\partial T/\partial RF)\cdot\left(\partial \sigma_f/\partial T\right) < 0$. Together with the fact that $\partial RF/\partial \mu < 0$ from Equation (2-4), this translates into $\partial \sigma_f/\partial \mu = (\partial RF/\partial \mu)\cdot \left(\partial \sigma_f/\partial RF\right) > 0$.

The PDFs of temperature increases for three hypothetical scenarios are illustrated in Figure 1. Distribution 1 refers to the case where radiative forcing is doubled relative to the pre-industrial level. Distributions 2 and 3 refer to the cases where there is learning (a 50% reduction in σ_f^2) and there is a 50% reduction in carbon emissions, respectively compared to the case for Distribution 1. Since this paper focuses on the tail-effect, only the probability density of the upper tail is considered below. This is reasonable in that the upper tail dominates the others in a usual cost-benefit analysis under deep uncertainty (Weitzman, 2009). As shown in the figure, for any T in the upper tail, probability density increases in radiative forcing ($\partial g_T/\partial RF > 0$). Likewise, for any T in the upper tail, probability density increases in uncertainty ($\partial g_T/\partial \sigma_f > 0$).

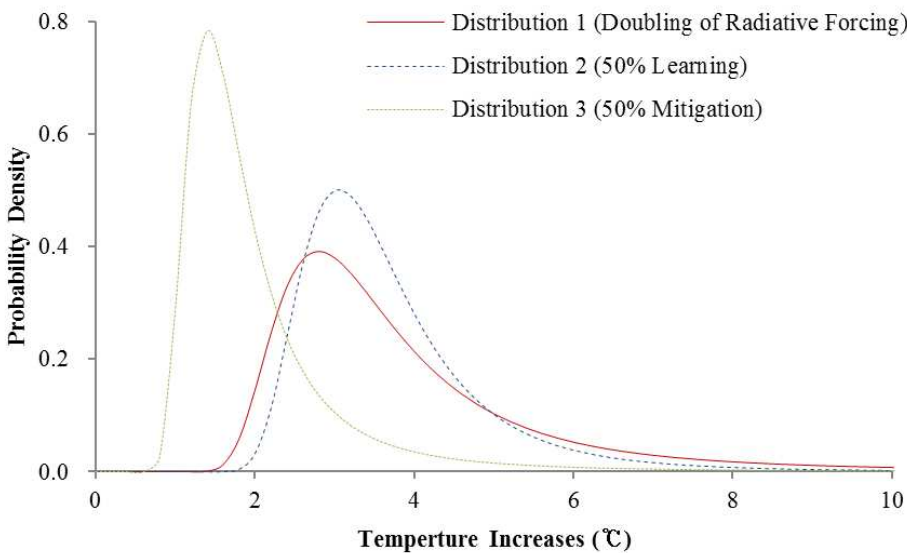

Figure 1 PDFs of temperature increases

One of the important things to be considered in the economics of catastrophic climate change is the rate of tail-slimming in the bad tail (Weitzman, 2013). If the tail-slimming rate of the upper tail is slower than the one for objective function, willingness to pay to avoid catastrophic climate change becomes arbitrarily large.

The tail-sliming rate at temperature T_u in the upper tail can be defined as in Equation (10). Since the fatness of the PDF decreases as temperature increases in the upper tail, the negative sign is attached in Equation (10).

$$-\frac{\partial g_T}{\partial T}\bigg|_{T=T_u} = \frac{\lambda_0 RF}{\sigma_f\sqrt{2\pi}RF_0T^3} exp\left\{-\frac{1}{2}\left[\frac{1-\bar{f}-\frac{RF}{RF_0}\frac{\lambda_0}{T}}{\sigma_f}\right]^2\right\}\left\{\frac{RF}{RF_0}\frac{\lambda_0}{T}\frac{\left(1-\bar{f}-\frac{RF}{RF_0}\frac{\lambda_0}{T}\right)}{\sigma_f^2}+2\right\}\Bigg|_{T=T_u}$$
$$\xrightarrow[T\to\infty]{} \sqrt{\frac{2}{\pi}}\frac{RF}{RF_0}\frac{1}{T^3}\frac{1}{\sigma_f} exp\left\{-\frac{1}{2}\left[\frac{1-\bar{f}}{\sigma_f}\right]^2\right\}\Bigg|_{T=T_u} \tag{10}$$

As expected, the tail-sliming rate is increasing in radiative forcing (see also Figure 2). This means that carbon emissions play a role in reducing the fatness of the upper tail, other things being equal. In addition, the tail-sliming rate is increasing in uncertainty as in Equation (11). This implies that learning is faster for larger uncertainty.

$$\frac{\partial}{\partial \sigma_f}\left(-\frac{\partial g_T}{\partial T}\bigg|_{T=T_1}\right) \propto \frac{1}{\sigma_f^2}\cdot\frac{RF}{T^3}\cdot\left[\left(\frac{1-\bar{f}}{\sigma_f}\right)^2-1\right]\cdot exp\left\{-\frac{1}{2}\left[\frac{1-\bar{f}}{\sigma_f}\right]^2\right\}\Bigg|_{T=T_1} > 0 \qquad (11)$$

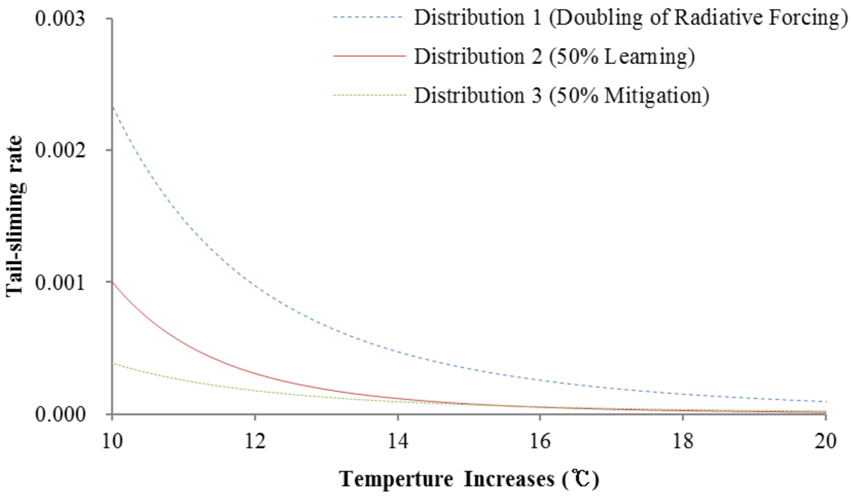

Figure 2 Tail-sliming rate of each distribution in Figure 1

4 The Learning-effect

From Equation (1), optimal climate policy should satisfy the first order condition as follows. For simplicity arguments of each function are dropped.

$$-\frac{\partial \Lambda_t}{\partial \mu_t}\cdot\frac{\partial u}{\partial \Lambda_t} = \beta\int_{\{T\}}\frac{\partial\left(W_{t+1}\cdot g_{T_{t+1}}\right)}{\partial \mu_t}dT \qquad (12)$$

The left hand side (LHS) of Equation (12) is the marginal abatement costs, whereas the right hand side (RHS) is the expected marginal benefits of emissions control (or the expected marginal avoided damage costs).

Without loss of generality we consider a three-period problem below. For the last period it is assumed that $\mu_3 = 0$ and $W_3 = u(c_3)$. The problem is recursively solved by backward induction (Bellman and Dreyfus, 1962). If μ_2^* is assumed to be a solution for the second period, given μ_1, the maximized social welfare W_2^* is calculated as in Equation (13).[1]

$$W_2^* = u(c_2^*) + \beta \int_{\{T\}} u(c_3^*) \cdot g_{T_3}^* \left(RF_3^*, \overline{f_3}^*, \sigma_{f_3}^* \right) dT \tag{13}$$

where $c_2^* = c_2^*(\mu_2^*|\mu_1)$, $c_3^* = c_3^*(\mu_2^*|\mu_1)$, $RF_3^* = RF_3^*(\mu_2^*|\mu_1)$, $\overline{f_3}^* = \overline{f_3}^*(\mu_2^*|\mu_1)$, $\sigma_{f_3}^* = \sigma_{f_3}^*(\mu_2^*|\mu_1)$.

The first order condition for the first period reads:

$$-\frac{\partial \Lambda_1}{\partial \mu_1} \cdot \frac{\partial u(c_1)}{\partial(\Lambda_1)} = \beta \int_{\{T\}} \left(\frac{\partial W_2^*}{\partial \mu_1} \cdot g_{T_2} + W_2^* \cdot \frac{\partial g_{T_2}}{\partial \mu_1} \right) dT \tag{14}$$

where $g_{T_2} = g_{T_2}\left(RF_2, \overline{f_2}, \sigma_{f_2} \right)$, $RF_2 = RF_2(\mu_1)$, $\overline{f_2} = \overline{f_2}(\mu_1)$, $\sigma_{f_2} = \sigma_{f_2}(\mu_1)$.

Substituting Equation (13) into Equation (14) and rearranging lead to Equation (15).

---

[1] Hwang et al. (2013a) investigates the conditions for the convergence of RHS in Equation (12) under the no-learning case.

$$-\frac{\partial \Lambda_1}{\partial \mu_1} \cdot \frac{\partial u(c_1)}{\partial(\Lambda_1)} = \tag{15}$$

$$\beta \cdot \frac{\partial}{\partial \mu_1} \int_{\{T\}} \left(u(c_2^*) \cdot g_{T_2}\right) dT + \beta \cdot \frac{\partial}{\partial \mu_1} \int_{\{T\}} \left\{ \beta \int_{\{T\}} \left(u(c_3^*) \cdot g_{T_3}^*\right) dT \right\} \cdot g_{T_2} dT$$

Equation (15) says that the marginal benefits of emissions control are the discounted sum of the expected marginal social welfare. The last term of RHS can be decomposed applying a chain rule:

$$\beta^2 \int_{\{T\}} \int_{\{T\}} \left\{ \frac{\partial u(c_3^*)}{\partial \mu_1} \cdot g_{T_2} \cdot g_{T_3}^* + u(c_3^*) \cdot \frac{\partial g_{T_2}}{\partial \mu_1} \cdot g_{T_3}^* + u(c_3^*) \cdot g_{T_2} \cdot \frac{\partial g_{T_3}^*}{\partial \mu_1} \right\} (dT)^2 \tag{16}$$

The last term in the bracket of RHS can be further decomposed into Equation (17). Note that it has been assumed that $\partial \bar{f}_3 / \partial \mu_1 = 0$.

$$\frac{\partial g_T}{\partial \mu} = \frac{\partial RF}{\partial \mu} \cdot \frac{\partial g_T}{\partial RF} + \frac{\partial \sigma_f}{\partial \mu} \cdot \frac{\partial g_T}{\partial \sigma_f} \tag{17}$$

The first term of RHS in Equation (17) reflects the effect of emissions control on the PDF of temperature increases through the changes in radiative forcing, whereas the second term is added because the parameters of the distribution change as learning takes place. Thus the second term is named the 'learning-effect' hereafter.[2]

[2] Note that the PDF of temperature increases for the second period g_{T_2} changes only by radiative forcing. After observing information in the second period the decision maker updates her PDF for the upcoming period.

The conditions for RHS in Equation (15) to converge do not change from the no-learning case of Hwang et al. (2013), since the first term in RHS dominates the other terms regarding the existence of solutions.

It is clear that $(\partial RF/\partial\mu)\cdot(\partial g_T/\partial RF)<0$ and $(\partial\sigma_f/\partial\mu)\cdot(\partial g_T/\partial\sigma_f)>0$ in the upper tail from Figure 1. These relations imply that the learning-effect offsets to some extents the effect of deep uncertainty on welfare. The marginal benefits of emissions control falls as the decision maker learns, and thus the optimal level of emissions control or carbon tax decreases when there is a possibility of learning, compared to the no-learning case.

The offsetting ratio is calculated as in Equation (18). The ratio grows in three conditions: 1) the quality of information that carbon emissions produce $(\partial\sigma_f/\partial\mu>0)$ increases, 2) carbon emissions are larger $(m\cdot RF)$, and 3) the level of learning is larger $(1/\sigma_f^3)$.

$$\left|\frac{(\partial\sigma_f/\partial\mu)\cdot(\partial g_T/\partial\sigma_f)}{(\partial g_T/\partial\mu)}\right|\propto\frac{\partial\sigma_f}{\partial\mu}\cdot\frac{m\cdot RF}{\sigma_f}\cdot\frac{\left(1-\bar{f}-\frac{RF}{RF_0}\cdot\frac{\lambda_0}{T}\right)^2/\sigma_f^2-1}{1+\left(1-\bar{f}-\frac{RF}{RF_0}\cdot\frac{\lambda_0}{T}\right)\cdot\frac{\lambda_0 RF}{\sigma_f^2 RF_0 T}}\xrightarrow[T\to\infty]{}\frac{\partial\sigma_f}{\partial\mu}\cdot\frac{m\cdot RF}{\sigma_f^3}\qquad(18)$$

Finally let us consider the case for active learning. Active learning refers to the case where the decision maker explicitly affects the rate of learning (see Hwang et al., 2013c). If the decision maker invests on climate science to raise the speed of learning ($R>0$, where R is the amount of investment), and if research is effective in reducing uncertainty ($\partial\sigma_f/\partial R<0$), then the offset ratio becomes far larger (see Equation 18). Consequently, there is an additional reduction of the marginal benefits of emissions control, hence the less stringent climate policy.

5 Conclusion

The effect of learning under fat-tailed uncertainty has been investigated in this paper using a simple dynamic model. The effect of emissions control on welfare is decomposed into the direct effect and the learning-effect. Main findings of this paper are that the possibility of learning reduces the marginal benefits of emissions control, compared to the case where there is no learning. As a result, learning reduces the stringency of climate policy. The fatter is the tail of the distribution and the faster is learning the larger is the learning-effect. Hwang et al. (2013c, 2014) investigate this issue with numerical models.

Acknowledgement

I am very grateful to Richard S. J. Tol and Marjan Hofkes for their helpful comments and suggestions. The remaining errors are my own.

References

Bellman, R., Dreyfus, S.E., 1962. Applied dynamic programming. The RAND Corporations.

Cyert, R.M., DeGroot, M.H., 1974. Rational expectations and Bayesian analysis. The Journal of Political Economy 82, 521-536.

Gregory, J.M., Foster, P.M., 2008. Transient climate response estimated from radiative forcing and observed temperature change. Journal of Geophysical Research 113, D23105.

Horowitz, J., Lange, A., 2014. Cost-benefit analysis under uncertainty: A note on Weitzman's dismal theorem. Energy Economics 42, 201-203.

Hwang, I.C., Reynès, F., Tol, R.S.J., 2013a. Climate policy under fat-tailed risk: An application of DICE. Environmental and Resource Economics 56(3), 415-436.

Hwang, I.C., Tol, R.S.J., Hofkes, M.W., 2013b. Tail effects and the role of emissions control. University of Sussex Working Paper Series No. 66-2013.

Hwang, I.C., Tol, R.S.J., Hofkes, M.W., 2013c. Active learning about climate change. University of Sussex Working Paper Series No. 65-2013.

Hwang, I.C., Tol, R.S.J., Hofkes, M.W., 2014. The effect of learning on climate policy under fat-tailed uncertainty. Unpublished manuscript. The earlier version was presented at the 20th EAERE Conference, June 2013, Toulouse, France.

Ingham, A., Ma, J., Ulph, A., 2007. Climate change, mitigation and adaptation with uncertainty and learning. Energy Policy 35, 5354-5369.

Karp, L., 2009. Sacrifice, discounting and climate policy: five questions. CESifo Working Paper No. 2761.

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

Kolstad, C.D., 1996a. Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. Journal of Environmental Economics and Management 31, 1-18.

Kolstad, C.D., 1996b. Fundamental irreversibilities in stock externalities. Journal of Public Economics 60, 221-233.

Leach, A.J., 2007. The climate change learning curve. Journal of Economic Dynamics and Control 31, 1728-1752.

Millner, A., 2013. On welfare frameworks and catastrophic climate risks. Journal of Environmental Economics and Management 65, 310-325.

Nordhaus, W.D., 2008. A question of balance: Weighing the options on global warming policies. New Haven and London: Yale University Press.

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

Pindyck, R.S., 2012. Uncertain outcomes and climate change policy. Journal of Environmental Economics and Management 63, 289-303.

Roe, G.H., Baker, M.B., 2007. Why is climate sensitivity so unpredictable? Science 318, 629-632.

Ulph, A., Ulph, D., 1997. Global warming, irreversibility and learning. The Economic Journal 107, 636-650.

Webster, M., 2002. The curious role of learning in climate policy: Should we wait for more data? The Energy Journal 23(2), 97-119.

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

Weitzman, M.L., 2012. GHG targets as insurance against catastrophic climate damages. Journal of Public Economic Theory 14(2), 221-244.

Weitzman, M.L., 2013. A Precautionary Tale of Uncertain Tail Fattening. Environmental and Resource Economics 55**,** 159-173.