UNIVERSITY OF MIAMI

ESSAYS IN ENVIRONMENTAL ECONOMICS AND INTERNATIONAL ECONOMICS

By

Zhuo Tan

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

May 2013

UNIVERSITY OF MIAMI

A dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

ESSAYS IN ENVIRONMENTAL ECONOMICS AND INTERNATIONAL ECONOMICS

Zhuo Tan

Approved:

David L. Kelly, Ph.D. Professor of Economics	M. Brian Blake, Ph.D. Dean of the Graduate School
Michael Connolly, Ph.D. Professor of Economics	Christopher F. Parmeter, Ph.D. Assistant Professor of Economics
Manuel Santos, Ph.D. Professor of Economics	Tie Su, Ph.D. Associate Professor of Finance

TAN, ZHUO (Ph.D., Economics)
Essays in Environmental Economics and International Economics (May 2013)

Abstract of a dissertation at the University of Miami.

Dissertation supervised by Professor David L. Kelly.
No. of pages in text. (90)

This body of work consists of four essays studying topics in environmental economics and international economics. The main interests in this thesis are to evaluate the effects of different government policies. Particularly, we explore the effects the optimal abatement policy for the world economy and the foreign exchange policy in China.

My first essay studies 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 "partial 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. Non-fat tailed uncertainty is largely irrelevant for optimal emissions policy.

My second essay provides a new solution algorithm for discrete time stochastic models of climate and the economy, relying on a nonparametric approximation of the value function. It is known by the dynamic programming theory that the value function is globally increasing and concave, but such information is not exploited by conventional approximation methods. This presents a challenge for solving the integrated assessment models numerically because climate change models have a large state space. The curse of dimensionality limits the size of the grid used in typical solution methods. Without concavity, local maxima can form in areas of the state space where grid points are sparse, which slows convergence. Therefore we establish a general approach to impose shape preserving constraints based on nonparametric econometrics by solving a quadratic programming problem. Then we illustrate stability and accuracy of the algorithm using an optimal growth model and a simple integrated assessment model with analytical solutions.

My third essay evaluates China's neutralization policy by monthly estimations based on the central bank balance sheet from 1999:6 to 2011:6. Our results suggest that China effectively neutralizes 66% of the change of net foreign assets under a pegged currency regime. Consequently, a purchase of one yuan of net foreign assets leads to an effective increase of 1.4 yuan in the money supply, rather than 4 yuan in the absence of neutralization. In the face of rapid growth of foreign reserves, neutralization in China is becoming increasingly difficult, consistent with Mundell's hypothesis that monetary authorities can fix the exchange rate and let the money supply float, or fix the money supply and let the exchange rate float: but it cannot fix both the exchange rate and the money supply.

My fourth essay estimates China's demand for foreign reserves from 1994:1 to 2007:4. Using a monetary model for China's reserve demand, we take into account the People's Bank of China's systematic neutralization policy to reduce inflation. While ultimately inconsistent, this policy has led to a growth in foreign exchange reserves that seems limitless: a neutralization coefficient of 0.57 leading to a magnification effect on the increase in reserves of 2.3. That is, a purchase of foreign reserves leads to a contraction of domestic credit of 57% of the foreign exchange purchase, which in turn magnifies the surplus under a stable exchange rate.

To my parents, Lingui Xie and Youren Tan

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the guidance from my committee members, help from my friends, and support from my family.

My deepest appreciation goes to my advisor, Professor David L. Kelly, for his excellent guidance and extreme patience. His outstanding courses on environmental economics and advanced macroeconomics economics have provoked my interests in optimal climate policy under uncertainty. Whenever I have questions, his office door is always open to me. Many times, his encouragement has helped me keep a positive attitude towards obstacles in research.

I am also indebted to the other members in my doctoral committee, Professor Michael Connolly, Professor Christopher Parmeter, Professor Manuel Santos, and Professor Tie Su, for their insightful comments and continuous support. In addition, I would like to extend my appreciation to the other professors at the Department of Economics at the University of Miami, particularly Nuray Akin, Raphael Boleslavsky, Christopher Cotton, Carlos Flores, Laura Giuliano, and Mehdi Shadmehr for their valuable suggestions and generous help on my thesis as well as my job search.

Moreover, I would like to express gratitude to Penelope Jones and Marbella Santamaria for their professional administrative support at all stages of my doctoral studies. I would not forget to thank my cohort of graduate students, especially Zhigang Feng, Cheng Li, Huong Nguyen, Yongmin Zang, and Lin Zhou for their company.

Finally, I would like to thank my parents for their unconditional love, endless support, and wise guidance throughout my life. This thesis is dedicated to them.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Introduction

This body of work contributes to the studies of government policies in environmental economics and international economics. The first two chapters explores the optimal abatement policy in the economics of climate change: Chapter 2 examines the optiaml abatement policy with Bayesian learning under fat tailed climate uncertainty, and Chapter 3 proposes a novel solution algorithm for the quatitative model in Chapter 2 known as the integrated assessment model based on a nonparametric kernel smoothing method. Chapter 4 and Chapter 5 both evalute the effects of China's neutralization policy.

Learning and Climate Feedbacks

Chapter 2 studies the effect of potentially severe climate change on optimal climate change policy, accounting for learning and uncertainty in the climate system. There is an ongoing debate in this field on whether immediate stringent emissions control for the uncertain low-probability high-impact catastrophes due to extreme climate change is necessary. Chapter 2 provides a quantitative answer to this question.

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 "partial 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. Non-fat tailed uncertainty is largely irrelevant for optimal emissions policy.

Solving Integrated Assessment Models of Climate Change

Chapter 3 provides a new solution algorithm for discrete time stochastic models of climate and the economy, relying on a nonparametric approximation of the value function. It is known by the dynamic programming theory that the value function is globally increasing and concave, but such information is not exploited by conventional approximation methods. This presents a challenge for solving the integrated assessment models numerically because climate change models have a large state space. The curse of dimensionality limits the size of the grid used in typical solution methods. Without concavity, local maxima can form in areas of the state space where grid points are sparse, which slows convergence. Therefore we establish a general approach to impose shape preserving constraints based on nonparametric econometrics by solving a quadratic programming problem. Then we illustrate stability and accuracy of the algorithm using an optimal growth model and a simple integrated assessment model with analytical solutions.

Neutralization in China

Chapter 4 evaluates China's neutralization policy by monthly estimations based on the central bank balance sheet from 1999:6 to 2011:6. Our results suggest that China effectively neutralizes 66% of the change of net foreign assets under a pegged

currency regime. Consequently, a purchase of one yuan of net foreign assets leads to an effective increase of 1.4 yuan in the money supply, rather than 4 yuan in the absence of neutralization. In the face of rapid growth of foreign reserves, neutralization in China is becoming increasingly difficult, consistent with Mundell's hypothesis that monetary authorities can fix the exchange rate and let the money supply float, or fix the money supply and let the exchange rate float: but it cannot fix both the exchange rate and the money supply.

China's Implicit Demand for Foreign Reserves

Chapter 5 estimates China's demand for foreign reserves from 1994:1 to 2007:4. Using a monetary model for China's reserve demand, we take into account the People's Bank of China's systematic neutralization policy to reduce inflation. While ultimately inconsistent, this policy has led to a growth in foreign exchange reserves that seems limitless: a neutralization coefficient of 0.57 leading to a magnification effect on the increase in reserves of 2.3. That is, a purchase of foreign reserves leads to a contraction of domestic credit of 57% of the foreign exchange purchase, which in turn magnifies the surplus under a stable exchange rate.

CHAPTER 2
Learning and Climate Feedbacks

2.1 Background

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

However, climate change uncertainty differs from a standard insurance problem in that learning reduces climate change uncertainty over time. If learning resolves climate uncertainty relatively quickly, then the initial insurance premium might be small, as the planner would still have time to increase abatement if learning quickly indicated the climate sensitivity was large. However, if learning is slow to resolve climate uncertainty, then the optimal policy calls for aggressive initial abatement for

[1]The Intergovernmental Panel on Climate Change (2007) reviews many studies and finds values higher than 4.5°C cannot be ruled out, although the best estimate is closer to 3°C. Weitzman (2009) 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: Gregory et al. (2002), Lemoine (2010), Newbold and Daigneault (2009), Roe and Baker (2007), Schwartz (2007), and Baker and Roe (2009).

insurance purposes. The central question for climate policy is then: how fast will learning resolve uncertainty about potential large steady state temperature changes caused by GHG emissions, 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 (1999a) consider uncertainty regarding the heat capacity of the ocean. In their integrated assessment model, 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, which is the steady state temperature change per unit of radiative forcing,[2] a measure of the elasticity of the climate to changes in GHG concentrations, and finds Bayesian learning about the climate sensitivity is extremely slow.[3] 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.[4] Keller, Bolker, and Bradford (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 (2013) study alternative uncertain thresholds with learning and find somewhat smaller effects on near term abatement.

[2]The more familiar $\Delta T_{2\times}$, or the steady state temperature change from a sustained doubling of GHG concentrations, is proportional to the climate sensitivity.

[3]Both of these papers 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.

[4]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, Stone, and Sokolov, 2005).

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. We define this case as "partial learning." To investigate partial learning, we develop a quantitative integrated assessment model of the world economy.[5] In the model, the planner faces stochastic weather shocks and uncertainty over the first order autoregressive coefficient in the equation governing evolution of temperature, which is the climate feedback parameter.

Because the climate feedback parameter is uncertain, the climate sensitivity is also uncertain. Further, if the climate feedback parameter is close to one, then GHG "shocks" to temperature are long lived, and therefore an increase in GHG emissions causes very high steady state temperature changes. Hence, although uncertainty in the feedback parameter is normally distributed, uncertainty in the climate sensitivity has a fat upper tail (see for example, Roe and Baker, 2007). The social planner learns the feedback parameter, and therefore the climate sensitivity, by updating prior beliefs given stochastic temperature data 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 the mean of the current prior equal to 2.76°C, so the tail of the distribution equals values greater than or equal to 4.26°C.[6] When the planner can reject the hypothesis that the climate sensitivity implies a steady state temperature increase greater than or equal to the critical temperature at the 1% or 0.1% level, we say that partial learning is complete.[7] Such learning is partial

[5] 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, 1999b).

[6] 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. Weitzman, 2009 discusses values above 7°C). A larger lower bound would only strengthen our results.

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

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 is able to reject values of the climate sensitivity in the upper tail of the prior distribution quickly. In fact, the planner can reject very high values of the climate sensitivity (e.g. 1.5°C or more above the mean estimate) at 99% confidence interval in less than a decade, if the true climate sensitivity is moderate. First, observations near the moderate true value provide evidence against the tails of the distribution. In addition, the density of even a fat tail is not large, so Bayes rule requires relatively few observations to reduce the mass of the fat tail below the critical confidence level. This result is surprising given the common intuition in the literature that reducing uncertainty in the tail of the climate sensitivity prior distribution must be a slow process since climate disasters are rare (see for example, Weitzman, 2009, page 12).

If the true climate sensitivity turns out to be relatively high, learning progresses more slowly. First, Bayes rule requires more observations to move the mean estimate from the prior of a moderate climate sensitivity to the true high value. Second, Bayes rule requires more observations to resolve the difference between a climate sensitivity that is relatively high and a climate sensitivity which is very high. Nonetheless, because a high climate sensitivity is relatively unlikely according to the prior, the possibility that learning is slower due to a high climate sensitivity receives relatively little weight when computing the expected time until partial learning is complete.

For example, we find that the expected learning time when the lower bound of the fat tail is 4.26°C is only about 8 years at the 0.1% level. However, if the lower bound

is 10°C, partial learning is not complete for almost 50 years. Integrating over the prior distribution, we find that the expected time, conditional on prior information, until partial learning is complete is only about 17 years at the 0.1% level.

Like Weitzman (2009), our model considers a high climate sensitivity as a possible scenario with high damages from climate change. Other potentially high damage scenarios exist, including high sea level rise (Nicholls, Tol, and Vafeidis, 2008), thermohaline circulation collapse (Keller, Bolker, and Bradford, 2004), and a reduction in the decay rate of carbon into the ocean (Lemoine and Traeger, 2013).[8] Learning in each of these contexts may differ from our results. For example, if variation in the size of ice sheets unrelated to temperature was significant, learning about sea level rise would proceed more slowly. Nonetheless, fat tailed uncertainty in the climate sensitivity is most commonly analyzed high damage scenario in the literature, presumably because the fat tail is clearly evident in the scientific priors.

In terms of optimal policy, we quantify the effect of uncertainty on near term emissions and abatement policy. With uncertainty but without learning, in the initial period emissions are about 38% lower, and the carbon tax is $22.94 higher, than under certainty. The planner insures by reducing emissions, paying for more abatement to reduce the probability of high damages that occur if the climate sensitivity is high. However, in the current period, emissions with uncertainty and learning are only about 19% lower than under certainty. The optimal carbon tax with uncertainty and learning is only $8.84 per ton higher than under certainty. Therefore, learning reduces emissions abatement for insurance purposes by about 50%. Further, optimal emissions with uncertainty and learning converge quickly to emissions given perfect information, typically in about 16 years. Uncertainties remain after 16 years, but the remaining uncertainty is not relevant for the optimal emissions policy. The fat tail

[8] Tol (2009) reviews the findings of the literature on various damages caused by climate change.

drives policy, and learning shrinks the mass of the fat tail quickly.[9]

Fat tailed uncertainty arises naturally when multiple uncertainties exist,[10] or, in the approach taken here, when an uncertain parameter has a multiplicative effect through feedbacks in the climate system. Climate sensitivity uncertainty is then fat tailed, even if the prior distribution for the uncertain parameter is normal (Roe and Baker, 2007). In turn, Weitzman (2009) shows that, given a coefficient of risk aversion greater than or equal to one, the risk premium required to accept an uncertain climate sensitivity with fat tails is infinite (the dismal theorem).[11]

However, Costello, Neubert, Polasky, and Solow (2010) shows that the dismal theorem is an asymptotic result. They find that truncating the distribution of uncertainty at a high level invalidates the dismal theorem. They argue that temperature changes should be truncated, since infinite temperatures are not physically possible, which we adopt here.[12]

Our results do not contradict the dismal theorem. Fat tails do in fact remain in our climate sensitivity distribution for any finite number of observations, and without truncation, the risk premium is infinite.

Our results therefore show the importance of results by Costello, Neubert, Polasky, and Solow (2010) for policy. We show that the fat tail is important for near term policy, even with truncation. Nonetheless, we show that with a small number of observations, learning reduces the upper tail of the distribution to close to an exponential. Therefore, although learning has no effect on the infinite risk pre-

[9]Optimal policy under learning converges quickly to the perfect information case even if the true climate sensitivity is very high. This is because with a higher mean estimate, deviations from certainty are driven by the mass of the tail of the new distribution with a higher mean, and the mass of the tail of the new distribution still shrinks quickly.

[10]For example, if the prior distribution for the climate sensitivity conditional on the variance of the weather shocks is the thin tailed normal distribution and the variance of the weather shocks is also unknown, with a gamma prior, then the unconditional prior distribution for the climate sensitivity is the fat tailed t distribution.

[11]See also Geweke (2001).

[12]Millner (2011) discusses other key assumptions of the dismal theorem.

mium without truncation, with truncation learning does significantly reduce the risk premium.

Put differently, our results highlight the difference between *variance* and *fat tails*. It is possible, through learning, to reduce the variance so that the mass in the upper tail is arbitrarily small. For example, we show that in many cases the planner can reduce the mass in the tail to 1% or 0.1% in just 10-16 years. The planner then easily rejects the hypothesis of a high climate sensitivity. However, tails are still fat in that the rate of decline in the upper tail is eventually slower than exponential. Therefore, with truncation, abatement policy is sensitive to the mass of the fat tail, which learning reduces by reducing the variance.

Considerable debate exists in the literature on the importance of fat tailed uncertainty for near term climate policy when learning is possible. For example, Weitzman (2011) argues that strong inertia in the climate should mean learning is less relevant for near term policy. Given that the stock of GHG emissions is difficult to reduce quickly, it will be difficult to reduce GHG concentrations if we learn climate change is more severe than expected. Conversely, Nordhaus (2011) argues that severe climate change should be evident in the data within the next 50 years, and so time exists to reduce GHG concentrations. Pindyck (2011) points out that the question is inherently quantitative and depends on the cost of insurance, the probability of severe climate change, etc.

This paper provides a quantitative answer. It is indeed important to reduce emissions initially due to the fat tailed uncertainty and the difficulty in reversing GHG stocks. Nonetheless, consistent with Nordhaus' idea, learning is indeed fast enough so that mid course corrections are possible. For example, we show that if the true climate sensitivity implies a steady state temperature change of 5°C for a doubling of GHG concentrations, then the planner reduces emissions to within 1% of the certainty

level in only 14 years. Emissions under learning remain slightly below (0.6-0.7%) the certainty level for 53 years as the planner corrects for initially over-emitting, to bring GHG concentrations to the perfect information trajectory.

Note that in all of above results, convergence to the perfect information level of emissions requires only partial learning. Indeed we show that learning is not complete (in the sense that the planner can reject values plus or minus 5% away from the true value with 95% confidence, see for example, Leach, 2007) for 50-100 years depending on the true value of the climate sensitivity. [13] Learning is slow, but the remaining uncertainty after the fat tail is statistically rejected is not important for policy. Only the fat tail provides an insurance motivation for near term abatement.

The structure of this paper is as follows. Section 2 presents our integrated assessment model with learning. The model is solved and calibrated in Section 3, and Sections 4 and 5 presents simulation results. Section 6 concludes.

2.2 Model

The model is similar in spirit to a simplified DICE model (Nordhaus, 2007b). Economic growth generates GHG emissions, which in turn cause temperatures to rise, reducing productivity. However, we use the abatement cost function from Bartz and Kelly (2008), and use different assumptions about future improvements in the emissions intensity of output. The learning model and stochastic temperature change follows Kelly and Kolstad (1999a) and Leach (2007) except that we look at learning the fat tailed climate sensitivity rather than learning the primitives of the temperature model. The damage function is from Weitzman (2009).

[13] In this sense, our results are consistent with Kelly and Kolstad (1999a), who estimate learning about the heat capacity was complete after about 90 years.

2.2.1 Economic system

The population of L_t identical households have preferences over consumption C_t/L_t given by the period constant relative risk aversion utility function:

$$U\left(\frac{C_t}{L_t}\right) = \frac{\left(\frac{C_t}{L_t}\right)^{1-\sigma} - 1}{1-\sigma}. \tag{2.1}$$

A constant returns to scale technology exists that produces output Q_t, from capital K_t, and productivity-augmented labor $A_t L_t$. Here A_t is labor productivity, which grows exogenously at rate ϕ. Population grows at exogenous rate η. The production technology is such that:

$$Q_t = F\left(K_t, A_t L_t\right) = K_t^{\psi}\left(A_t L_t\right)^{1-\psi}. \tag{2.2}$$

Unabated GHG emissions are an exogenous proportion $1/B_t$ of output. Let u_t denote the fraction of emissions abated, then $(1-u_t)/B_t$ is the emissions intensity of output, and emissions, E_t, are:

$$E_t = \left(1-u_t\right)\frac{Q_t}{B_t}. \tag{2.3}$$

The cost of abatement is $\Lambda\left(u_t\right)Q_t$. Hence, output net of abatement costs, Y_t, is

$$Y_t = \left(1-\Lambda\left(u_t\right)\right)Q_t. \tag{2.4}$$

We assume a convex cost function:

$$\Lambda\left(u_t\right) = 1-\left(1-u_t\right)^{\epsilon}. \tag{2.5}$$

The abatement cost function (2.5), differs from standard cost functions in the literature. For example, Nordhaus (2008) uses a two parameter function: $\Lambda(u_t) = \epsilon_1 u_t^{\epsilon_2}$. The abatement cost function (2.5) has a particular advantage in that it is consistent with a balanced growth path (Bartz and Kelly, 2008), which simplifies the computations considerably. Further, without a balanced growth path, either emissions goes to zero or infinity, or the growth rate in labor productivity must be forced to go to zero.

Using equations (2.2)-(2.5) to substitute out for Q_t and u_t implies output net of abatement costs is Cobb-Douglas:

$$Y_t = F(K_t, B_t E_t, L_t A_t) = K_t^{\theta} (B_t E_t)^{\epsilon} (A_t L_t)^{1-\theta-\epsilon}. \tag{2.6}$$

Here $\theta = \psi(1-\epsilon)$ is the capital share and ϵ can thus be interpreted as the emissions share.[14]

A balanced growth path is a steady state where aggregate capital, output, and consumption all grow at the same constant rate $(1+\eta)(1+\phi) - 1$. A balanced growth path exists with constant emissions if the exogenous growth rate of B_t equals the growth rate of output:[15]

$$B_{t+1} = (1+\eta)(1+\phi) B_t. \tag{2.7}$$

Exogenous growth of B_t captures both technological change in abatement and compositional changes in output.

Let capital depreciate at rate δ_k. The resource constraint then sets consumption

[14]Bartz and Kelly (2008) calibrate the emissions share for four air pollutants.

[15]Note that if B_t grew slower than the rate of output then the returns to emissions savings innovation would approach infinity, while the returns to labor productivity would go to zero, and the reverse if B_t grew faster than the rate of output.

plus net investment equal to production net of abatement costs after damages, $D(T)$, due to climate change:

$$C_t = (1 - D(T_t)) Y_t + (1 - \delta_k) K_t - K_{t+1}, \tag{2.8}$$

where the damage function is:

$$D(T_t) = 1 - e^{-b_1 T_t^{b_2}}. \tag{2.9}$$

Here b_1 and b_2 are damage parameters.

2.2.2 Climate system

Let M_t represent the current accumulation of carbon-equivalent GHG in the atmosphere, and MB is the preindustrial stock. We assume the ocean and biosphere absorb atmospheric carbon at a constant rate δ_m. Let $\gamma = A_0 L_0 E_0 / Q_0$ be the initial emission intensity coefficient. The stock of pollution accumulates according to:

$$M_{t+1} - MB = (1 - \delta_m)(M_t - MB) + \gamma E_t. \tag{2.10}$$

We use a one equation physical model for temperature:

$$\hat{T}_t = \hat{T}_{t-1} + \frac{1}{\alpha}\left(F_t - \frac{\hat{T}_{t-1} - \Gamma}{\lambda}\right) + \nu_t. \tag{2.11}$$

Here $\hat{T}_t$ is the annual global temperature (difference in °C between year t and the 1961-1990 average temperature); Γ is preindustrial temperature difference from the 1961-1990 average temperature; α is the heat capacity of the upper ocean; λ is the climate sensitivity; $\nu \sim N(0, \sigma_\nu^2)$ is the stochastic weather shock; and F_t is radiative

forcing of GHGs:

$$F_t = \Omega \log_2 \left(\frac{M_t}{MB} \right). \tag{2.12}$$

Here Ω is the radiative forcing parameter. Equations (2.11) and (2.12) vastly simplify large physical models of climate, known as general circulation models (GCMs). Nonetheless, models similar to (2.11) and (2.12) are frequently estimated and used for policy analysis.[16]

The uncertain parameter is the climate sensitivity λ, which measures how responsive temperature is to GHG concentrations. The climate sensitivity amplifies the effect of radiative forcing on temperature. To see this, rewrite equations (2.11) and (2.12) as:

$$T_t = \beta_1 T_{t-1} + \beta_2 \log_2 \left(\frac{M_t}{MB} \right) + \nu_t. \tag{2.13}$$

Here $T_t = \hat{T}_t - \Gamma$ is the annual global temperature deviation from preindustrial level, $\beta_1 = (1 - \frac{1}{\lambda \alpha})$ is the climate feedback parameter, and $\beta_2 = \frac{\Omega}{\alpha}$. The climate feedback parameter is positively related to the climate sensitivity. Since the climate sensitivity is uncertain, the climate feedback parameter β_1 is also uncertain.

Let $\Delta T_{2\times}$ be the steady state temperature change that results from a steady state doubling of the GHG concentrations, then from equations (2.11) and (2.13):

$$\Delta T_{2\times} = \Omega \lambda = \frac{\beta_2}{1 - \beta_1}. \tag{2.14}$$

Since λ is uncertain $\Delta T_{2\times}$ is also uncertain. The steady state temperature change from a doubling of CO_2 is most straightforward to understand, so we will report most results in terms of the equivalent $\Delta T_{2\times}$.

[16] See for example, Andronova and Schlesinger (2001), Gregory, Stouffer, Raper, Stott, and Rayner (2002), and Schwartz (2007). Kelly, Kolstad, Schlesinger, and Andronova (1998) discuss some weaknesses of one equation physical models. Traeger (2012) calibrates a one equation model to match near term temperature changes predicted by the Nordhaus DICE model.

Emissions, E_t, temperature, T_t, and GHG concentrations, M_t, are constant along the balanced growth path.

2.2.3 Learning

Each period the social planner observes new statistical records of the climate system and updates beliefs on the uncertain feedback parameter. Bayesian learning characterizes this process.

Assume the planner has prior beliefs that the true β_1 is drawn from a normal distribution, $N(\mu_0, S_0)$, where S_0 is the variance of the prior distribution. Let:

$$H_t = T_t - \beta_2 \log_2 \left(\frac{M_t}{MB} \right) = \beta_1 T_{t-1} + \nu_t. \tag{2.15}$$

Then $H_{t+1} \sim N\left(\mu_{H,t}, \sigma^2_{H,t}\right)$ combines the stochastic weather shocks and feedback uncertainty into a single random variable, where $\mu_{H,t} = \mu_t T_t$ and $\sigma^2_{H,t} = T_t^2 S_t + \sigma^2_\nu$.

The $t+1$ weather shock occurs at the start of the period, before decisions are made. The social planner thus observes H_{t+1}, and T_t and updates the prior on β_1. Let $p_\nu = 1/\sigma^2_\nu$ be the precision of ν. Then Bayes rule implies that the posterior distribution of β_1 is also normally distributed with:

$$\mu_{t+1} = \frac{\mu_t + S_t p_\nu T_t H_{t+1}}{1 + S_t p_\nu T_t^2}, \tag{2.16}$$

$$S_{t+1} = \frac{S_t}{1 + S_t p_\nu T_t^2}. \tag{2.17}$$

Note that from Equation (2.17) that the variance estimate on β_1 is monotonically non-increasing with time. We use variance instead of the usual precision since the variance is bounded above by the prior, while the precision is unbounded above.[17] Perfect information implies $S = 0$ and $\mu = \beta_1$.

[17] The computational solution method requires a bounded state space.

Roe and Baker (2007) compute the probability density function (PDF) for the climate sensitivity from a Jacobian transformation. Let, $\beta_1 = 1 - \frac{1}{\alpha\lambda} \sim N(\mu, S)$, then the density for the climate sensitivity is:

$$h_\lambda(\lambda) = h_{\beta_1}(\beta_1(\lambda)) \left(\frac{\partial \beta_1}{\partial \lambda}\right) = \frac{1}{\sqrt{2\pi S}} \frac{\frac{1}{\alpha}}{\lambda^2} \exp\left[-\frac{1}{2S}\left(1 - \frac{1}{\alpha\lambda} - \mu\right)^2\right]. \quad (2.18)$$

The prior distribution of the climate sensitivity has a fat upper tail if and only if the tail probability declines to zero at a rate slower than exponential:

$$\lim_{\lambda \to \infty} \frac{h_\lambda(\lambda)}{\exp(-a\lambda)} > 0, \quad a > 0. \quad (2.19)$$

It is straightforward to verify that $h_\lambda(\lambda)$ satisfies condition (2.19). The fat upper tail is clearly visible in Figure D.1, which plots $h_\lambda(\lambda)$.

2.2.4 The Recursive Problem

The social planner maximizes the expected present discounted value of the stream of period utilities weighted by the population size:

$$W = \max \sum_{t=0}^{\infty} \beta^t L_t U\left(\frac{C_t}{L_t}\right). \quad (2.20)$$

We normalize the problem so that variables are in per labor productivity unit terms. Let $f(K, E) = F(K, E, 1)$, $k_t = K_t/(A_t L_t)$ and similarly for c_t and y_t, $m_t = \frac{M_t}{MB}$ and assume $\hat{\beta} \equiv \beta(1+\eta)(1+\phi)^{1-\sigma} < 1$. Then we can write the recursive version of the social planning problem as:

$$\begin{aligned} v(k, m, T, \mu, S) &= \max_{k', E} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \right. \\ &\left. \hat{\beta} \int_{H'} v(k', m', T', \mu', S') \, \mathrm{N}\left(H', \mu T, T^2 S + \sigma_\nu^2\right) dH' \right\}, \quad (2.21) \end{aligned}$$

subject to:

$$c = (1 - D(T))\, f\,(k, E) + (1 - \delta_k)\, k - (1 + \eta)\,(1 + \phi)\, k', \tag{2.22}$$

$$m' = 1 + (1 - \delta_m)\,(m - 1) + \left(\frac{\gamma}{MB}\right) E, \tag{2.23}$$

$$T' = H' + \beta_2 \log\,(m')\,, \tag{2.24}$$

$$\mu' = \frac{\mu + Sp_\nu T H'}{1 + Sp_v T^2}, \tag{2.25}$$

$$S' = \frac{S}{1 + Sp_v T^2}. \tag{2.26}$$

Equation (2.21) condenses the double expectation over unknown variables β_1 and ν' into an expectation over a single unknown variable H'. Table A.1 gives the variable definitions for the above problem.

2.3 Calibration and Solution Method

Table A.2 gives the parameter values. For the economic parameters, we chose a risk aversion coefficient of $\sigma = 1.5$. The discount factor is consistent with a 5% rate of time discount, which implies the model economy matches the US capital to output ratio. The depreciation rate of capital is 6%, which implies the model economy matches the US investment to output ratio. The population growth rate is 1% and the growth rate of per capita GDP is 1.8%, which match US data. These values are broadly consistent with the business cycle literature (Bartz and Kelly, 2008; Cooley and Prescott, 1995).

The emissions share parameter is also the parameter of the abatement cost function. In general, the abatement cost function has only one parameter and therefore cannot match both the low cost of abatement when u_t is low, and the convexity of the Nordhaus (2008) cost function. We chose the calibrated value to match well for

low values of u_t, which results in an abatement cost function which is more convex. Therefore, initial control rates, emissions, and carbon taxes will tend to be similar to Nordhaus, but in later years, when larger control rates are optimal, our results will generally show lower abatement costs and carbon taxes.

The damage parameters are taken from Weitzman (2009). The damage parameters and discount factor are set relatively conservatively. A discount factor closer to one or a high damage scenario (for example, Weitzman also considers $b_2 = 4$) would increase both the value of abatement as climate insurance, and the value of learning. We consider the high discount factor case in Section 2.6.

One issue with the calibration is that the model sets the growth rate in the inverse of the emissions intensity of output (B_t) equal to the growth rate of output. This is done so that the balanced growth path exists with constant emissions. In the long run, the growth rate of B_t should equal the growth rate of output otherwise the returns to labor productivity innovations will either be infinite or zero. That is, any other value is inconsistent with a more general model in which R&D expenditures flow towards the sector (emissions or labor) with the highest marginal returns to innovation. Nonetheless, in the short run the rate of growth in B_t can differ from that of output (for example, if random innovations in one sector are more successful than the other). In fact, the growth rate of B_t has exceeded the growth rate of output in recent years. Allowing for different short run growth rates (as is done in Nordhaus, 2008, for example), would considerably complicate the analysis by adding additional state variables, however.

The variance of the weather shock is taken from Leach (2007). The remaining climate and forcing parameters are set consistent with Nordhaus (2008).

We use a multidimensional spline approximation of the value function and value function iteration to solve the problem. After we solve the model, we approximate

the policy function using another spline approximation and do the simulations. Our numerical integration truncates the temperature distribution at 26.8°C, for computational reasons.[18]

2.4 Results: Learning

2.4.1 Partial Learning

The solution of the dynamic program is a value function $v(k, m, T, \mu, S)$ and the corresponding policy functions giving optimal investment and emissions $k'^*(k, m, T, \mu, S)$ and $E^*(k, m, T, \mu, S)$. The evolution of the state and control variables over time follows:

$$\left[\begin{array}{cccccc} k' & m' & T' & \mu' & S' & E \end{array} \right] \equiv G\left([k, m, T, \mu, S], \lambda, \nu'\right). \tag{2.27}$$

Here the function G is comprised of the policy functions and the laws of motion for the state variables, equations (2.23)-(2.26). A simulated time path thus requires a set of random draws for ν over time and a true value of λ. Each experiment consists of 150 simulations, over 6,500 years, for a given true value of λ. We report simulations for various hypothetical true values for λ, and then take expectations over the results using the prior distribution to get expected results given the current state of information.

Table A.3 gives the initial conditions, which are set to the year 2005. The initial conditions for the learning parameters are from Roe and Baker (2007), who survey various global circulation models (GCMs). Initial conditions for capital, temperature, and GHG concentrations are from Nordhaus (2008).

[18]See Costello, Neubert, Polasky, and Solow (2010) for a justification. Further, the numerical integration routine also truncates so that values of the climate feedback parameter outside the unit interval have zero probability. In this sense, a prior distribution bounded by zero/one, such as the beta distribution may be more appropriate. We use the normal distribution since any computational solution method requires a conjugate family of distributions, but also to follow the literature (see for example, Roe and Baker, 2007).

We say partial learning is complete if the planner can reject values of the climate sensitivity which are in the tail of the distribution at a given confidence level. That is, partial learning is complete if the planner rejects the hypothesis $\lambda < T^L/\Omega$ or equivalently $\Delta T_{2\times} < T^L$, where T^L is the lower border of the tail of the distribution (°C).

The lower border of the climate sensitivity distribution which constitutes the "tail" has no precisely agreed upon definition. Further, from Equation (2.21) the optimal decision takes into account the entire distribution in a continuous way. Our hypothesis is that the tail of the distribution drives current optimal abatement policy. If the mass of the tail shrinks quickly, the remaining uncertainty is irrelevant and the optimal policy converges to the certainty case. The first step is therefore to show that learning about the tail of the distribution differs from learning about the mean, which requires a definition of what constitutes the tail of the distribution.

Therefore, we set the lower border of the tail of the distribution conservatively as $T^L = \Delta T_{2\times,t} + 1.5°\text{C}$, where $T_{2\times,t}$ is the current mean estimate of the climate sensitivity.[19] The initial prior mean equals 2.76°C, so the initial lower border of the tail is $T^L = 4.26°\text{C}$, which constitutes the upper 16.7% of the mass of the distribution. If the true value is large, say 5°C, then learning will move the mean of the prior higher, and the uncertainty can eventually be partitioned into uncertainty about the exact value of the climate sensitivity (near 5) and uncertainty about the probability of still higher values (even more disastrous values greater than 6.5 are possible).[20]

The planner's optimization problem (2.21) does not specify such a test as part

[19] Values above 1.5°C will speed up the learning, values less than 1.5°C mean the tail is relatively likely to occur, which is inconsistent with the idea that the true value is in the tail with relatively low probability.

[20] Alternatives are less attractive. If the lower border of the tail is fixed for all t (say $T^L = 5°\text{C}$), then the planner must essentially learn the exact value of the climate sensitivity for true values near the lower border, since in that case the tail of the distribution converges to the upper half of the distribution. As $\Delta T_{2\times,t} \to T^L$, this definition of partial learning requires an arbitrary large number of periods to complete, which is not consistent with the optimal policy in Section 2.5.

of the optimal policy. Instead, we are developing a hypothesis that the effect of uncertainty on optimal policy can be approximately partitioned into two parts: the effect of uncertainty over the exact value of the climate sensitivity, and the effect of uncertainty over the mass of the tail of the distribution.

We consider two confidence levels, 99% and 99.9%. Given the high damages associated with severe climate change, we assume the planner requires a relatively high level of confidence that the climate sensitivity is not large before rejecting the hypothesis that the true climate sensitivity is in the tail of the distribution.

We consider 50 possible true values, $\Delta T_{2\times}^{i*}$, indexed by i. For each $\Delta T_{2\times}^{i*}$, for each simulation j, and associated random vector of weather shocks ν_j, we record the first period, $n^*\left(\Delta T_{2\times}^{i*}, \nu_j\right)$, for which the hypothesis that $\Delta T_{2\times,jt} > T^L$ is rejected and not subsequently not-rejected. We then say that the planner achieves partial learning in period n^*, for true value of $\Delta T_{2\times}^{i*}$ and simulation j.

We then average over all simulations, and weight all n^* using the prior distribution. Mathematically:

$$\mathrm{E}\left[n^* | \, \mu_0, S_0\right] = \int_{1/\alpha}^{\infty} \int_{\nu} n^*\left(\lambda, \nu\right) N\left(0, \sigma_\nu^2\right) h_\lambda\left(\lambda\right) d\nu d\lambda. \tag{2.28}$$

Table A.4 shows the results. The expected time to complete partial learning varies from 8.99-14.58 periods, depending on the required confidence level. Partial learning is relatively quick for three reasons. First, by definition the initial mass of the upper tail is not large. Therefore it takes relatively few observations away from the tail to reduce the mass to 1% or 0.1%. Second, although the global temperature is subject to random weather shocks which makes the exact value of the climate sensitivity difficult to pin down, the calibrated standard deviation of the shocks is only 0.11°C, which makes it relatively easy to reject high values of the climate sensitivity if the true value

is not too large. Third, most of the mass of the prior distribution is relatively close to the mean. Therefore, true values with smaller learning times receive more weight when calculating the expected learning time conditional on the prior distribution. As expected, partial learning at the 99.9% confidence level requires more periods than at the 99% confidence level.

Figure D.2 shows the learning time as a function of $\Delta T^*_{2\times}$ for both confidence levels. For true values near the mean of the prior distribution (2.76°C), learning takes less than 10 periods, but rises to 90 years or more as $\Delta T^*_{2\times}$ increases. Partial learning becomes increasingly difficult as $\Delta T^*_{2\times}$ increases. Since $\lambda = 1/(\alpha(1-\beta_1))$, for values of β_1 near one, small differences in β_1 generate large differences in λ. Therefore, the planner must learn β_1 with increased precision to reject values of λ in the tail of the distribution when β_1 is near one. For example, with β_1 equal to the prior, the lower bound of the tail ($T^L = 4.26$) corresponds to $\beta_1^L = 0.77$, which is rejected in reasonable time given a true value of $\beta_1^* = 0.65$ and a calibrated standard deviation of the weather shocks equal to 0.11. Conversely, if $\Delta T^*_{2\times} = 5$ and $T^L = 6.5$ then $\beta_1^* = 0.81$ and $\beta_1^L = 0.85$, which is much harder to reject.

Put differently, for β_1 large, precise learning is much more important, because small differences in β_1 imply large differences in the decay rate of GHG "shocks" to temperature. Small differences in β_1 therefore imply large differences in steady state temperature for a given concentration of GHGs, making a precise estimate of β_1 more important.

Figures D.1 and D.2 together show that, for most of the prior distribution, partial learning requires less than 20 periods to complete. The expected number of periods until partial learning is complete conditional on prior information is relatively small. Longer learning times are possible but unlikely.[21]

[21] Note that the priors are based on physical, rather than statistical models of temperature change.

2.4.2 Alternative Assumptions

We next report learning times given alternative assumptions. First, we simulate the model with an ocean layer.[22] Let $\hat{O}_t$ be the ocean temperature in deviations from the recent average and ζ and τ be heat transfer coefficients, then the climate system (2.11) now becomes:

$$\hat{T}_t = \hat{T}_{t-1} + \frac{1}{\alpha}\left(F_t - \frac{\hat{T}_{t-1} - \Gamma}{\lambda}\right) - \zeta\left(\hat{T}_{t-1} - \hat{O}_{t-1}\right) + \nu_t, \tag{2.29}$$

$$\hat{O}_t = \hat{O}_{t-1} + \tau\left(\hat{T}_{t-1} - \hat{O}_{t-1}\right). \tag{2.30}$$

Because heat transfer to the ocean is a slow process, the model now takes much longer to reach the steady state. However, Table A.4 shows that partial learning is complete in 15.84-24.39 periods, which is only slightly longer than with the one equation model. To learn the true climate sensitivity, the planner must statistically isolate the upward trend in the climate from the stochastic weather shocks. The ability of learning to do so is only marginally affected by the ocean and other climate processes with long lags, since the ocean temperature is essentially constant on a year to year basis. Thus our results are robust to the addition of a more complicated climate model.

A very large weather shock can cause a hypothesis which was rejected in period n to be no longer rejected in $n+1$. We also considered an alternative criterion in which n^* is the first period for which the hypothesis is first rejected, even if the hypothesis is no longer rejected in a subsequent period. Table A.4 gives the results, which are only slightly smaller than the base case given in Table A.4. As expected, learning is faster in this case, but still of a similar magnitude.

Therefore, the results may be interpreted as saying that to confirm or refute physical models which predict a relatively high climate sensitivity will require relatively little additional data.

[22] For computational reasons, we use the emissions and investment policy functions given in Equation (2.27), which assumes the ocean layer is constant. Since the ocean layer changes slowly over time, this assumption is reasonable for the short term.

2.4.3 Other Criteria

The above results constitute partial learning in that the only goal is to reject values of the climate sensitivity in the tail of the prior distribution, not to fully learn the true value. Similar to previous work (Leach, 2007), we define two hypothesis tests, $\Delta T_{2\times} \leq 0.95 \cdot \Delta T^*_{2\times}$ and $\Delta T_{2\times} \geq 1.05 \cdot \Delta T^*_{2\times}$, where $\Delta T^*_{2\times}$ is the true value, and a desired confidence level of 95%. If the planner rejects both hypothesis, we say that the planner has fully learned the uncertain climate sensitivity, or that full learning is complete. Figure D.3 plots the mean number of periods required to fully learn the climate sensitivity. The learning time is increasing in the true value. First, as the true value gets farther from the prior, Bayes rule must completely reject the prior information in favor of the new observations. Second, as noted above, $\Delta T_{2\times}$ is a nonlinear function of β_1. When $\Delta T_{2\times}$ is large, the range of values of β_1 that constitute the plus or minus 5% of the true value becomes narrower, which increases the expected learning time. Integrating over the prior distribution, we find an average learning time equal to 79.67 years, This is roughly consistent with the previous literature (Kelly and Kolstad, 1999a; Leach, 2007).

Figure D.4 plots the learning dynamics for the mean of the feedback parameter μ for three simulations, when the true value is 0.65 ($\Delta T_{2\times} = 2.76$). In all simulations, learning converges rather slowly to the true value (about 55-150 years). However, even for simulation 5, where the first few weather shocks are positive, the mean estimate of $\Delta T_{2\times}$ is well below the border of the tail, $T^L = 4.26$. The hypothesis $T^L \geq 4.26$ is first rejected at the 99.9% level in periods 7, 10, and 5, respectively. Figure D.5 plots the density for simulation 5 after 10 periods. Considerable uncertainty remains. Values of $\Delta T_{2\times}$ between 2.5 and 4 are plausible. However, values above 4 are very unlikely. Partial learning is complete, although full learning is not complete.

The last 3 rows of Table A.5 presents the same information for the mean of 150

simulations. The initial mass of the tail is 16.5%, and the 95% confidence interval admits a wide range of possible values for the climate sensitivity, from the benign 0.97 to the disastrous 7.09°C temperature change for a doubling of GHGs. By 2015, however, partial learning is complete at the 99.9% confidence level, whereas full learning is not complete: the 95% confidence interval is still fairly wide at almost 1°C. By 2050, the 95% confidence is fairly narrow, although full learning is still not quite complete.

Figure D.6 reports the learning dynamics for the standard deviation of the prior distribution of the feedback parameter, $\sqrt{S}$, for simulation 5 when the true climate sensitivity equals the prior. The standard deviation is monotonically decreasing as expected, and converges to zero.

Learning the climate sensitivity precisely is a slow process. According to the prior literature with thin tails, slow learning indicates the optimal policy under learning is unlikely to be much different than the optimal policy without learning. Since learning is slow, the planner acts using current information. However, the above analysis shows that with learning the planner rejects extreme values relatively quickly, unless the true value is large. With fat tails, the extreme values drive current policy. Therefore, the learning is potentially much more policy relevant with fat tails. The next section makes these ideas precise.

2.5 Results: Optimal Policy

2.5.1 Optimal Insurance

In this section, we examine how learning and fat tailed uncertainty affect the optimal emissions policy. Figure D.7 plots optimal emissions for the case where the true climate sensitivity is equal to the initial prior value. The circle line corresponds to perfect information, where the initial variance of the prior is set to zero. In this case, optimal emissions increases for a short period of time and is then decreasing. The

initial world capital stock is only 68% of its steady state level. Therefore, the planner postpones most emissions control until the capital stock has converged and more resources are available. Both damages and costs are a fraction of world GDP, so an increase in GDP affects both damages and costs equally. However, more wealth implies more consumption, which decreases the marginal utility of consumption. Therefore, emissions control becomes more attractive. In addition, each year's emissions have only a small effect on the GHG concentration, so the planner does not incur much additional damage by waiting.

The line with squares shows the optimal policy under uncertainty with learning. As the planner learns the true value, the emissions under learning and uncertainty converges to emissions given perfect information. Notice that emissions under learning are initially below the perfect information case. The planner insures, emitting less than under perfect information just in case climate change turns out to be severe. Emissions are initially 19.3% lower under learning than under perfect information (Table A.5), but are only 1.1% lower by 2021. Uncertainty matters for a relatively short period of time. In Figure D.7, the true value is the initial guess. Therefore, the planner quickly rejects values of the climate sensitivity in the fat tail, and the policy approximately converges to the perfect information case by 2025.[23]

The plus line corresponds to the no learning case. In this case, the learning parameters μ and S are not updated, despite the new information. In addition, the model is solved so that the planner knows μ and S will not be updated. Therefore, the no learning case differs from the learning case, even in the initial period when the state vector is identical for both policies. Emissions are lowest under no learning.

[23]Note that emissions are slightly above the perfect information case for a short period of time. This is because the planner has under-emitted relative to perfect information during the learning process. The planner can therefore over-emit after rejecting the fat tail to reach the same steady state stock of GHGs.

The planner must insure more by reducing emissions without learning, because the planner knows that she cannot adjust later as more information is revealed. Therefore, learning reduces the need for climate insurance. Emissions for the no learning case are initially 38% below emissions given perfect information, whereas emissions under learning are 19.3% below perfect information. Therefore, learning reduces climate insurance by about 50% (Table A.5).

Emissions under no learning are below the true optimal emissions for the entire time path. The planner under no learning must continue to insure, whereas under learning climate insurance is required for only a short time. Emissions for the no learning case are 48% below perfect information in 2020, whereas emissions under learning are only 1.42% below certainty at this point.

Given a true value equal to the prior ($\Delta T_{2\times} = 2.76$), the planner rejects the tail at the 99.9% confidence level after 7.2 periods (see Figure D.2 and Table A.5). Emissions under learning is only 4.64% lower than emissions under certainty after 8 periods, and is only 1.07% lower after 16 periods (Figure D.7 and Table A.5). In contrast, full learning is not complete for 63.8 years if $\Delta T_{2\times} = 2.76$ (see Figure D.3 and Table A.5). Therefore, optimal emissions policy is more sensitive to partial learning. Although uncertainty is present and the climate sensitivity is difficult to pin down precisely, the planner rejects values in the tail quickly and thereafter proceeds as if the planner was certain that $\Delta T_{2\times}$ equals the mean of the current prior.

Figure D.8 gives the emissions control rate, which is increasing over time under certainty. This is similar to the "ramp up" strategies found for example by Nordhaus (2008). The control rate under learning is initially more stringent, but converges to the certainty case at approximately the same rate as emissions. Without learning, the control rate remains elevated as the planner continues to insure. The initial control rate under learning is 19.3%, which is similar to values found in the literature. In

contrast, without learning the initial control rate is 38.3%, approximately twice as high (Table A.5). Figure D.9 shows the carbon tax. The initial carbon tax is $46.1 per ton, also within the range of typical estimates.

Figure D.10 gives the path of temperature increases. The policies with learning and perfect information differ very little in terms of the GHG stock, since the planner adjusts emissions to keep the economy on the same GHG stock trajectory after learning takes place. Emissions and the GHG stock are lower under no learning due to the insurance, resulting in a smaller temperature increase.

Figures D.11-D.13 repeat the experiment with the true $\Delta T_{2\times} = 2$ rather than 2.76. In this case, the planner learns over time that the climate has only small feedback effects, and therefore that GHG concentrations cause smaller steady state temperature increases. The planner then emits considerably more under perfect information than when $\Delta T_{2\times} = 2.76$, 3.55 gigatons more in the steady state. Figure D.12 indicates the planner under perfect information reduces emissions by less than one percent initially, which increases to a maximum reduction of only 1.5%. Under learning, the planner begins with the same information set as when the true $\Delta T_{2\times} = 2.76$, and therefore chooses the same initial policy as in Figure D.7. As new information arrives which decreases the prior, the planner under learning begins to increase emissions.

The planner rejects the fat tail sooner here, since a relatively small λ implies β_1^L is not close to β_1^*. Emissions under learning converge to within 0.4% of emissions under perfect information after 4 periods. Similarly, partial learning is complete after 3.65-5.07 periods at the 99% and 99.9% confidence levels, respectively (Figure D.2). In contrast, full learning is not complete until after 57.0 periods (Figure D.3). After ruling out the fat tail, the planner proceeds along a path very close to perfect information, even though uncertainty remains.

Without learning, the initial conditions are identical as when the true value is 2.76.

Therefore, the initial policies without learning are identical in Figures D.7 and D.11. Since the planner does not update the prior, emissions are below perfect information indefinitely. In fact, emissions without learning are quite similar in Figures D.7 and D.11. The insurance motivation is the main determinant of emissions policy without learning, and differences in emissions policy caused by the different temperature trajectories are minor.

Finally, Figures D.14-D.16 repeat the experiment with the true $\Delta T_{2\times} = 5$. Given that GHG concentrations are projected to more than double,[24] this represents a high damage case. Under certainty, the planner responds to the high $\Delta T_{2\times}$ by severely limiting emissions. Figure D.15 indicates the planner reduces emissions initial by 52.4% and by 72.2% in the steady state under certainty.

Initial optimal emissions policy with learning is unchanged from the previous cases, since the initial beliefs are unchanged. Emissions fall over time as the planner increases the mean belief of the climate sensitivity over time in response to higher than expected temperatures. Emissions policy under learning converges to within 1.1% of emissions policy under certainty in 17 periods. From Figure D.2, partial learning is complete after 16.09 periods at the 99.9% confidence level. In contrast, learning is not complete for about 105.6 years (see Figure D.3) for a true value of $\Delta T_{2\times} = 5$. The planner learns that the climate sensitivity is much higher than the prior, but also learns that still higher values are unlikely, even if the exact value is difficult to pin down. Optimal emissions converges to the certainty case, and the remaining uncertainty is not relevant for policy.

Emissions under learning is about 0.6%-0.7% less than emissions under certainty from 2022 to 2270. Under learning the planner has over-emitted relative to perfect

[24] However, the optimal emissions path under certainty limits GHG concentrations to only 55% above preindustrial levels when $\Delta T_{2\times} = 5$.

information in the first 16 years, and must therefore under emit relative to certainty to bring GHG concentrations to the same steady state trajectory. The planner smoothes out the error correcting over a considerable period because emissions control costs are convex and because of consumption smoothing.

For all three true values of $\Delta T_{2\times}$, the upper tail of the distribution is initially policy relevant. Emissions under learning are significantly below emissions if the prior is known with certainty (Figure D.11). However, learning quickly reduces the mass of the upper tail of the distribution, and emissions quickly converge to the certainty case.

2.5.2 Fat Tails and Variance

Figures D.1, D.17, and D.18 shows the evolution of the posterior PDF for $\Delta T_{2\times}$ when the true value is 2.76 (equal to the current prior), after 0, 10, and 50 periods, respectively. We also plot the normal distribution with identical parameters on each graph,[25] to emphasize the fat tail of the distribution of the climate sensitivity. Contrasting Figures D.1 and D.17, we see that after only 10 periods the mass of the fat tail shrinks considerably. Values of $\Delta T_{2\times}$ above 3.5 are very unlikely.

However, as Weitzman (2009) points out for the case of the t distribution, the tail remains fatter than normal for any finite number of observations. Here our distribution is not t, but the same result applies. Condition (2.19) holds for any $S > 0$, so the tails remain fat for any finite number of observations. After 50 years, the mass of the tail barely visible on Figure D.18, but nonetheless, the tail is still fatter than normal.

Section 2.5.1 shows fat tails are relevant initially for optimal emissions policy. Emissions with learning are 19.3% below the certainty case. Further, when a climate

[25]Note that the variance of λ is infinite. Figures D.1, D.17, and D.18 plot $h(\lambda; \mu, S)$ along with a normal distribution with mean μ and variance S.

sensitivity above 4.26 is rejected at the 99.9% level, emissions are within 1% of certainty, even though fat tails remain for all finite observations. Therefore, what is important for emissions policy is not the existence of fat tails, but the mass of the fat tail, which is a function of the variance. Learning is relevant since learning affects the variance.

2.6 Sensitivity Analysis

The climate sensitivity is considered a major source of uncertainty in integrated assessment models (Kelly and Kolstad, 1999a). Nonetheless, other sources of uncertainty also exist, especially the level and convexity of damages (Weitzman, 2009) and the discount factor (Nordhaus, 2008). Higher or more convex damages or a discount factor closer to one all increase the benefits of abatement. We therefore consider a representative sensitivity analysis, decreasing the pure rate of time preference from 0.05 to 0.03, or alternatively increasing β to 0.97.

Figures D.19-D.21 graph the mean optimal emissions for the high discount factor case. Contrasting Figures D.7 and D.19, we see that optimal emissions are lower with the higher discount factor, as expected. Optimal emissions under certainty in the initial period fall by 24.33% when the discount factor increases from 0.95 to 0.97. Emissions under learning and no learning fall by 39.63% and 35.37%, respectively, when the discount factor increases.

Since damages are potentially greater, optimal insurance also increases. Table A.6 indicates current emissions under no learning are 47.31% lower than emissions under certainty for the high discount factor case, versus 38.31% for the low discount factor case. Similarly, emissions under learning are 35.62% lower under learning than emissions under certainty for the high β case, versus 19.3% for the low β case. Although damages are potentially greater, learning becomes somewhat less important,

since the high discount factor means large emissions reductions are optimal for a wide range of outcomes of the learning process. Learning reduces the insurance premium by only 24.71% in the high β case, versus about 50% in the low β case.

Both partial and full learning are slightly slower for the high β case. Kelly and Kolstad (1999a) prove that more restrictive climate policies slow learning because the climate change signal is less pronounced amidst the noisy weather. Table A.6 shows that this effect is small, however. The mass of the fat tail is 1.4% in 2010 when $\beta = 0.97$, versus 1.3% when $\beta = 0.95$. The 95% confidence interval is $[2.37, 3.31]$ in 2015 when $\beta = 0.97$ versus $[2.39, 3.30]$ when $\beta = 0.95$. Although learning is slightly slower, emissions under learning converges faster to perfect information when $\beta = 0.97$. Table A.6 and Figure D.19 indicate emissions under learning converges to within 1% of emissions under certainty in 7 years, versus about 17 period when $\beta = 0.95$.

Overall, sensitivity analysis indicates that if the discount factor is closer to one or damages are greater or more convex, the main results remain. Fat tails are initially policy relevant, with a large insurance premium, learning significantly reduces the insurance premium, and learning is fast in that the fat tail is rejected quickly, although significant uncertainty remains for decades, that uncertainty is not relevant for optimal policy.

2.7 Concluding Remarks

In this paper, we study the effect of a possible high climate sensitivity on near term optimal climate change policy, accounting for learning and uncertainty in the climate system. We find three major results. First, fat tails are initially policy relevant in that near term GHG emissions policy is much more restrictive when the planner accounts for fat tailed uncertainty in the climate sensitivity (a 38% reduction in emissions).

Second, when the planner accounts for learning, the near term emissions reduction falls by half to only 19.3%. Third, although full learning is slow, learning quickly reduces the mass of the fat tail. Optimal emissions policy is much more sensitive to the mass of the fat tail than the uncertainty in the prior around the mean. Therefore, optimal emissions policy converges quickly to perfect information, even though some uncertainty remains for decades.

The planner knows values of the climate sensitivity in the tail of the prior distribution will be rejected quickly at a high level of confidence if the true climate sensitivity is moderate. If the climate sensitivity is high the planner can quickly reject still higher values, and quickly adjusts emissions to get back on the optimal temperature trajectory. The planner has an option to essentially purchase climate insurance: by paying to limit GHG emissions today, the planner prevents GHG concentrations from rising, which in turn prevents the possibility of very high temperature changes. Without learning, the planner takes out a significant amount of insurance. However, with learning the planner insures about 50% less. First, learning quickly rejects values of the climate sensitivity in the fat tailed part of the prior distribution, if the true climate change is moderate. Second, the planner has time to adjust emissions to keep the economy on the same GHG stock trajectory. Climate insurance under learning in most cases falls to less than 1% after about 17 years as the planner reduces the mass of the tail end of the distribution and the remaining uncertainty is not important for emissions policy.

Several caveats are in order. First, for computational reasons, our model of the climate system is highly simplified. For example, we do not include a separate ocean layer. Nonetheless, we computed an optimal policy assuming the ocean temperature is constant, but simulated the model and learning with an ocean layer. The results are not much different since learning here is about isolating the upward trend in the

atmospheric temperature from the stochastic weather shocks. The ability of learning to do so is only marginally affected by the ocean and other climate processes with long lags, since the ocean temperature is essentially constant on a year to year basis. Thus our results are robust to the addition of longer lags to the climate model.

Second, we consider only a single uncertainty, the climate sensitivity. Climate change has many uncertainties, including the parameters of the damage function, the heat capacity of the ocean, etc. In general, multiple simultaneous uncertainties slows learning. It is unclear, however, how much the partial learning we consider here would slow.

Third, our model has no irreversibilities, tipping points, etc. The existence of irreversibilities makes the planner much more cautious, which increases insurance with or without learning. Learning would certainly still reduce climate insurance in this case, but by less as the planner may not be able to correct a mistake of initially over-emitting. We leave this interesting extension to future research.

Fourth, in our model the planner estimates climate feedbacks using current data. If the climate is subject to regime shifts which occur in the far future, then it might be difficult to learn about the existence of regime shifts today. However, if the process which causes the regime shift is observable today, then our model still applies. Suppose for example, the climate sensitivity is different if the polar ice caps melt as sunlight no longer reflects back into space as efficiently (the albedo effect). If one can estimate the albedo effect by observing changes in ice cover and changes in temperature, then the planner can learn the albedo effect before a regime shift to a world without polar ice caps occurs.

Regardless, our main results are likely robust to any of these extensions. Fat tails matter for climate policy, even if the distribution has a truncation point. Nonetheless, we show that learning significantly reduces the influence of fat tails, especially over

the near term. Given these results it is important for policy makers to maintain policy flexibility, and to stand ready to quickly adjust the emissions policy as new information arrives.

Finally, our results have interesting potential implications for recent research which finds fat tailed uncertainty in other contexts (equity market returns, banking crises, etc.). Fat tails in other contexts is typically modeled as an exogenous property of the return distribution rather than an endogenous implication of parameter uncertainty. Our results show that if the exogenously imposed fat tails are in fact the result of parameter uncertainty, then learning has the potential to reduce fat tailed uncertainty over time, which limits the risk premium of fat tailed uncertainty. We leave this interesting possibility to future research.

CHAPTER 3
Solving Integrated Assessment Models of Climate Change

3.1 Background

Global warming caused by anthropogenic carbon emissions is the one of the most formidable challenges facing mankind today. The optimal abatement policy in principle should balance the cost of abatement and the expected future damages caused by climate change, which is usually determined by solving complex quantitative models. In the literature, models which combine "scientific and socio-economic aspects of climate change for the purpose of assessing policy options for climate change control" are referred to as "integrated assessment models (IAMs)" (Kelly and Kolstad, 1999b). IAMs are the workhorse for researchers and policy makers to understand the linkages and interactions between the complex climate system and the human society, and set the optimal abatement and mitigation policies. The widely used IAMs include Nordhaus (2007a) DICE, Nordhaus (2010) RICE, and Manne, Mendelsohn, and Richels (1993) MERGE.

One method to build an IAM is to construct a recursive dynamic programming problem with a built-in climate sector, such as Kelly and Kolstad (1999a) and Leach (2007). Such method allows an endogenous economic response, i.e., an agent makes a decision optimally given a climate policy (Kelly and Kolstad, 1999b). However, the solution of these models presents special challenges to the existing numerical methods. For example, although climate change is a long term issue, the policy maker is more interested in the near term abatement policy. Since the current state of the world, or the initial condition, normally is not close to the steady state, numerical methods which use local approximations around the steady state, such as LQ approximations,

are particularly inaccurate (Kelly and Kolstad, 2001). Smooth approximation methods such as polynomial or spline interpolations often violate the shape properties of the value function and therefore lead to unstable fluctuations and errors in the value function iterations (Cai and Judd, 2012). Furthermore, most of these models have a large number of states and controls, especially when incorporating uncertainties, and thus suffer from the curse of dimensionality. For one thing, computational time increases exponentially in the number of states. For another, numerical approximations in a higher dimensional space are generally more difficult. The curse of dimensionality limits the size of the grid used in typical solution methods. Without shape constraints such as monotonicity and concavity, local maxima can form in areas of the state space where grid points are sparse, which leads to approximation errors and slows convergence.

It is known by the dynamic programming theory that the value function is globally increasing and concave (see for example, Santos, 1999), but such information is not exploited by conventional approximation methods. Recently, Cai and Judd (2012) introduce a shape-preserving dynamic programming algorithm with value function iteration based on shape-preserving Chebyshev interpolations. They show that imposing basic shape restrictions on value function approximations will stabilize value function iteration with their numerical examples.

In this chapter we develop a new solution algorithm for discrete time stochastic models of climate and the economy, relying on a nonparametric approximation of the value function. The nonparametric approximation method is based on Du, Parmeter, and Racine (2013). In the paper, the authors propose a kernel smoothing estimation method that can handle multiple general shape constraints for multivariate functions. Here we use the method to approximate the value function numerically. We show that with the new shape-preserving approximation method, value function iteration

in dynamic programming problems will be more stable.

This chapter is organized as follows. Section 3.2 summarizes the structure and types of IAMs. Section 3.3 presents the standard dynamic programming solution algorithm and introduces the nonparametric numerical approximation method. Section 3.4 gives an optimal growth model and a simple IAM, both with analytical solutions, to show the stability and accuracy of our algorithm.

3.2 Integrated Assessment Models

Kelly and Kolstad (1999b) provide a comprehensive review on IAMs. Based on the policy options available to the regulator, or the behavior of the choice variables, IAMs can be broadly divided into *policy evaluation models* and *policy optimization models*. Policy evaluation IAMs consider the effect of a single policy option on the biosphere, climate, and sometimes economic systems. Policy evaluation models are also known as simulation models. In contrast, policy optimization IAMs have a dual purpose: to seek to find the optimal policy which trades off expected costs and benefits of climate change control or the policy which minimizes costs of achieving a particular goal, and to simulate the effect of an efficient level of carbon abatement on the world economy (Kelly and Kolstad, 1999b).

Kelly and Kolstad (2001) provides a summary of the structure of IAMs. Let S_t be a vector of the state variables, C_t be a vector of control variables, ξ_t be random variables which represents either stochastic shocks to the economics and climate systems or parametric uncertainties, g be vectors of transition equations, the laws of motion in the economic and climate systems can be summarized as:

$$S_{t+1} = g(S_t, C_t, \xi_{t+1}). \tag{3.1}$$

Suppose there exists a set of all feasible decisions, Γ, for the choice variables, a

feasible set of distribution Ω for ξ, and a time separable objective function U, then policy optimization IAMs in essence are standard social planning problems of the form:

$$W(S_t) = \max_{C_t \in \Gamma(S_t)} \left\{ \mathrm{E} \left[\sum_{t=0}^{\infty} \beta^t U(C_t, S_t) \right] \right\}, \tag{3.2}$$

subject to:

$$S_{t+1} = g(S_t, C_t, \xi_{t+1}), \tag{3.3}$$

$$\xi_{t+1} \sim \Omega, \tag{3.4}$$

$$\lim_{t \to \infty} \beta^t U_s(C_t, S_t) \cdot S_t = 0. \tag{3.5}$$

That is, the social planner's objective is to maximize expected present discounted utility subject to the transition equations and a transversality condition. Note that policy evaluation models and deterministic models are a sub case of the above framework for a well specified Γ and Ω, respectively.

Under standard assumptions (Santos, 1999), the solution to the this social planning problem can be written recursively as a Bellman equation as follows:

$$v(S) = \max_{C \in \Gamma(S)} \left\{ U(S, C) + \beta \mathrm{E} \left[v \left(g \left(S, C, \xi \right) \right) \right] \right\}, \tag{3.6}$$

where $v(S)$ is the value function. Bellman equation (3.6) is a functional equation from a computational point of view, and function v is a fixed point of this equation. The most common way to compute this fixed point is known as value function iteration. In the next section, we present a value function iteration method with multidimensional nonparametric kernel approximations.

3.3 Nonparametric Solution Algorithm

The difficulty in the recursive problem to find the value function v numerically as the fixed point. Once v is known, it is easy to calculate the policy function at any point in time by using standard optimization method to solve the right hand side of Equation (3.6) for the optimal choices.

Generally, the state space and the feasible set for the random variables and controls in IAMs contain a continuum of elements. However, since computers can not handle infinite dimensional objects directly, discretization using numerical methods is necessary.

3.3.1 Value Function Iteration

We rely on a numerical dynamic programming algorithm with value function iteration on a discretized space to solve the recursive problem (see for example, Santos, 1999). The algorithm uses the Bellman equation to iterate, except that we iterate using the approximation of the value function $\hat{v}$, rather than the unknown v. We follow the standard assumptions that v belongs to a measurable, bounded real-valued functional space from Γ to $\mathbb{R}$, and the metric is the sup-norm metric. Below is a summary of the algorithm:

Step 1: Initialization. Form a uniform grid $\tilde{S} = \{S_i\}_{i=1}^{n}$ in the feasible set of state variables S.[1] Choose an initial approximation parameters p_0 and therefore the initial approximated value function $\hat{v}_0(S, p_0)$. Pick a convergence criterion ϵ and then iterate through Step 2 and 3.

Step 2: Maximization. For each grid points S_i, we use $\hat{v}_0$ to find $\hat{v}_1$ using the Bellmen

[1] See Judd (1998) for a discussion on grid choices.

equation. Or more generally, for all $S_i \in \tilde{S}$, we use $\hat{v}_m$ to find v_{m+1} by solving:

$$v_{m+1}(S_i) = \max_{C \in \Gamma} \{U(S_i, C) + \beta \mathrm{E}\left[\hat{v}_n\left(f\left(S_i, C, \xi\right)\right), p_m\right]\}. \tag{3.7}$$

Step 3: Approximation. Using an appropriate approximation method, compute p_{m+1} to construct an approximated value function $\hat{v}_{m+1}(S, p_{m+1})$ based on $v_{m+1}(\tilde{S})$.

Step 4: Termination. If $\parallel v_{m+1}(\tilde{S}) - \hat{v}_m(\tilde{S}) \parallel \leqslant \epsilon$, stop; else, increment i by 1, let $\hat{v}_{m+1}(S, p_{m+1})$ be the new guess $\hat{v}_m(S, p_m)$, and return to step 2.

The solution algorithm based on numerical approximation has an advantage that once solved, the optimal actions for all the values in the feasible set S is known, which offers great flexibility in policy analyses and simulations. As explained in Section 3.2, because of the high dimensionality of IAMs, however, the major challenge in the solution algorithm comes from the approximation step. Next we provide a method to approximate the value function across the entire state space with monotonicity and concavity constraints.

3.3.2 Nonparametric Kernel Approximation

Here we follow the kernel smoothing method proposed in Du, Parmeter, and Racine (2013) to approximate the value function. Kernel smoothing methods provide a powerful tool to approximate an unknown function by using real data without assuming a parametric form of the unknown function. In addition, Du, Parmeter, and Racine (2013) provide the first method in the nonparametric econometrics literature that allows multiple shape constraints at a time on multivariate functions. The method enables us to exploit the theoretical properties of the value function such as monotonicity and concavity, which in principle will stabilize the iteration process and generate more accurate solutions as the optimization problem is now convex.

The nonparametric kernel approximation method can be roughly divided into two steps: first obtain an *unrestricted* kernel approximation given a bandwidth vector and then impose shape restrictions to get the *restricted* approximation.

Unrestricted Kernel Approximation

Suppose our objective is to approximate a d dimensional value function. That is, there are d variables in the state space, $S = [s_1, s_2, \ldots, s_d]$. Let $\{v(S_i), S_i\}_{i=1}^n$ denote the gridded data, where $v(S_i)$ is a scalar, S_i is of dimension d, and n is the size of grid points. The goal here is to approximate the unknown value function $v(S)$ subject to constraints on $v^{(q)}(S)$ where q is a d dimensional vector corresponding to the dimension of S. The elements of p represent the order of the partial derivative corresponding to each element of S. For example, $q = (0, 0, \ldots, 0)$ represents the function itself, while $q = (1, 0, \ldots, 0)$ represents $\partial v(S)/\partial s_1$.

We first explain how to compute the unrestricted kernel approximation. An approximation of the value based on standard kernel smoothing regression can be written as linear combinations of $v(S_i)$:

$$\hat{v}(S) = \sum_{i=1}^{n} A_i(S) v(S_i), \tag{3.8}$$

where $A_i(S)$ is a local weighting matrix.

Following Hall and Huang (2001), we consider a generalization of the standard kernel regression smoothers:

$$\hat{v}(S|p) = \sum_{i=1}^{n} p_i A_i(S) v(S_i), \tag{3.9}$$

and for what follows $\hat{v}^{(q)}(S|p) = \sum_{i=1}^{n} p_i A_i^{(q)}(S) v(S_i)$, where $A_i^{(q)}(S) = \frac{\partial^{q_1} A_i(S) \ldots \partial^{q_d} A_i(S)}{\partial^{q_1} S_1 \ldots \partial^{q_d} S_d}$

for S. Note that the well-known unrestricted Nadaraya-Waston (NW) estimator (see Nadaraya, 1965; Watson, 1964, for details) is a special case of Equation (3.9), by setting:

$$A_i(S) = \frac{nK_h(S_i, S)}{\sum_{i=1}^{n} K_h(S_i, S)} \text{ and } p_i = \frac{1}{n},\ i = 1, 2, \ldots, n, \tag{3.10}$$

where $K_h(\cdot)$ is a kernel product and h is a vector of bandwidths (see Racine and Li, 2004, for details). In our numerical examples in Section 3.4 we use the NW method to calculate the optimal bandwidth and obtain the unrestricted approximations.

Imposing Shape Restrictions

The second step is to impose the shape restrictions on the unrestricted approximation in Equation (3.9).

Given an approximation of the value function $\hat{v}(S)$, the shape constraints on $\hat{v}(S)$ could be written as:

$$l(S) \leqslant \hat{v}^{(q)}(S) \leqslant u(S), \tag{3.11}$$

where $l(\cdot)$ and $u(\cdot)$ denote the lower and upper bounds on $\hat{v}^{(q)}(S)$ respectively. In our applications, for example, coordinate-wise concavity can be expressed in the above form with $q = (2, 2, \ldots, 2)$, $l(S) = \hat{v}\left(\min(S)\right)$, and $u(S) = 0$.

The two sided constraint in Equation (3.11) is a special case of multiple simultaneous one-sided constraints of the form:

$$\sum_{q \in Q_r} \alpha_{q,r} \hat{v}^{(q)}(S) \geqslant 0,\ r = 0, 1, \ldots, R, \tag{3.12}$$

where R is the number of restrictions, and in each restriction the sum is taken over all vectors in Q_r that correspond to our constraints, and $\alpha_{q,r}$ is a set of constants

used to generate various constraints. Note that these constraints must be internally consistent.

The key here is the selection of weights p. Let $p_u = \frac{1}{n}$ be the n dimensional vector and let p be the vector of weights to be selected. In order to impose the shape restrictions, we choose p to minimize the distance from p to p_u, $D(p)$, in a L_2 metric space. That is, in the second step, our objective is to solve the quadratic programming problem below:

$$\min_p D(p) = (p_u - p)'(p_u - p), \tag{3.13}$$

subject to Equation (3.12). The problem can be solved using standard quadratic programming methods and packages.

3.4 Numerical Examples

In this section we apply our algorithm to two examples, a standard optimal growth example and a simple integrated assessment model.

3.4.1 Optimal Growth Model Example

We first illustrate our method with a discrete time optimal growth model. We assume a logarithmic utility function over consumption, a Cobb-Douglas production function, full capital depreciation every period, and no population and technology growth in the economy. A recursive version of the model can be written as:

$$v(k) = \max_{k'}[\log(k^\theta - k') + \beta v(k')], \tag{3.14}$$

subject to

$$k' = k^\theta - k', \tag{3.15}$$

where k is capital, c is consumption, θ is capital share, and $v(k)$ is the value function.

The above problem has a linear analytical solution of the form:

$$v(k) = M + N \log(k), \tag{3.16}$$

where

$$M = \frac{1}{1-\beta}\left[\log(1-\beta\theta) + \beta\left(\frac{\theta}{1-\beta\theta}\right)\log(\beta\theta)\right], \tag{3.17}$$

$$N = \frac{\theta}{1-\beta\theta}, \tag{3.18}$$

with policy function:

$$k' = \beta\theta k^{\theta}. \tag{3.19}$$

3.4.2 Integrated Assessment Model Example

We next illustrate our method with a simple integrated assessment model with two state variables and two choice variables, which is a simplified version of the model in Chapter 2. To decrease the dimension of the state space, here we exclude learning, and assume climate damages to output comes from the flow of unabated emissions directly. To derive a simple analytical solution, we further assume full capital depreciation every period, a logarithmic utility function, and no population and technology growth in the economy.

Assume a logarithmic utility function for the L identical households in the economy over consumption per capita C_t/L:

$$U\left(\frac{C_t}{L}\right) = \log\left(\frac{C_t}{L}\right). \tag{3.20}$$

A constant returns to scale technology exists that produces output Y_t from capital K_t and labor L_t:

$$Q_t = F(K_t, L) = K_t^{\psi} L^{1-\psi}. \tag{3.21}$$

Assume unabated pollution is an exogenous proportion $1/B$ of output. Let u_t denote the fraction of emissions abated, or the emission control policy, then $(1-u_t)/B$ is the emission intensity of output. Therefore unabated emissions E_t can be written as:

$$E_t = (1 - u_t)\frac{Q_t}{B}. \tag{3.22}$$

The abatement however, is not without a cost. We assume a convex cost function:

$$\Lambda(u_t) = 1 - (1 - u_t)^{\varepsilon}, \tag{3.23}$$

such that the output net of abatement costs, Y_t, is:

$$Y_t = (1 - \Lambda(u_t))Q_t. \tag{3.24}$$

Using equations (3.21)-(3.24) to substitute for Q_t and u_t implies output net of abatement cost is Cobb-Douglas:

$$Y_t = F(K_t, E_t, L) = K_t^{\theta}(BE_t)^{\varepsilon}(L)^{1-\theta-\varepsilon}, \tag{3.25}$$

where $\theta = \psi(1 - \varepsilon)$ is the capital share and ε can be interpreted as emissions share.

Damages to output are due to unabated emissions:

$$D(E_t) = 1 - \frac{1}{1 + E_t^2}. \tag{3.26}$$

The resource constraint then sets consumption plus net investment equal to production net of abatement costs after damages due to unabated emissions:

$$C_t = (1 - D(E_t))Y_t - K_{t+1}. \tag{3.27}$$

The social planner's objective is to maximize lifetime utility of all households:

$$V = \max_{C_t} \sum_{t=0}^{\infty} \beta^t L \log\left(\frac{C_t}{L}\right). \tag{3.28}$$

The above problem can be normalized with variables written in per productivity unit terms. Let $f(K, E) = F(K, E, 1)$, $k_t = K_t/L$ and similarly for c_t and y_t, then the recursive version of the social planning problem can be written as:

$$v(k) = \max_{k', E} \{U[(1 - D(E))f(k, E) - k'] + \beta v(k')\}, \tag{3.29}$$

subject to:

$$c = (1 - D(E))f(k, E) - k'. \tag{3.30}$$

In addition, maximum and minimum emissions exist corresponding to $u = 0$ and $u = 1$, so that:

$$0 \leqslant E \leqslant k^{\frac{\theta}{1-\varepsilon}}. \tag{3.31}$$

Denote $H = \frac{1}{2}\varepsilon^{\frac{\varepsilon}{2}}(2 - \varepsilon)^{1-\frac{\varepsilon}{2}}$. It is straightforward to show that the above problem has a linear analytical solution of the form:

$$v(k) = M + N \log(k), \tag{3.32}$$

where

$$M = \frac{1}{1 - \beta}\left[\log(H(1 - \beta\theta)) + \frac{\beta\theta}{1 - \beta\theta}\log(\beta\theta H)\right], \tag{3.33}$$

$$N = \frac{\theta}{1 - \beta\theta}, \tag{3.34}$$

with policy functions:

$$k' = \beta\theta H k^{\theta}, \tag{3.35}$$

$$E = \sqrt{\frac{\varepsilon}{2-\varepsilon}}. \tag{3.36}$$

3.5 Conclusion

In this paper we develop a new solution algorithm for discrete time stochastic models of climate and the economy, relying on a nonparametric approximation of the value function. It is known by the dynamic programming theory that the value function is globally increasing and concave, but such information is not exploited by conventional approximation methods. This presents a challenge for solving the integrated assessment models numerically because climate change models have a large state space. The curse of dimensionality limits the size of the grid used in typical solution methods. Without concavity, local maxima can form in areas of the state space where grid points are sparse, which slows convergence. Therefore we establish a general approach to impose shape preserving constraints based on nonparametric econometrics by solving a quadratic programming problem. Then we illustrate stability and accuracy of the algorithm using an optimal growth model and a simple integrated assessment model with analytical solutions.

CHAPTER 4
Neutralization in China

4.1 Background

On January 1, 1994 China identified its currency regime as a "managed floating exchange rate regime based on market supply and demand." However, in practice, the renminbi (RMB) exchange rate was rigidly pegged at 8.28 yuan per dollar from 1997 until July 21, 2005. Then China announced that the RMB would be more flexible, and managed "with reference to a basket of currencies" rather than pegged to the US dollar. However, the implementation of the new reform has been questioned. Some studies argue that, in practice, China still pegs its currency to the dollar (Goldstein and Lardy, 2006; Fidrmuc, 2010). Others suggest the euro has a higher weight in the basket (Frankel, 2009). Nevertheless, there is common agreement that China has pegged its exchange rate for nearly two decades, albeit appreciating at a gradual 3.4% per annum since 2005. The slow appreciation has been accompanied by a policy of buying the foreign exchange at set rates.[1]

Since 1998, China has registered surpluses on both of its current and capital accounts due to strong net exports and net capital inflows. As a result, the People's Bank of China has passively purchased foreign exchange from enterprises and financial entities at pegged rates under the exchange settlement and purchase system. Consequently, the PBC's holding of net foreign assets increased to US$2.2 trillion in

[1]Under the pegged exchange rate, China implements strict capital control. For example, among others, the People's Bank of China (PBC) forbids most domestic enterprises and entities other than import/export companies from opening and maintaining foreign exchange accounts. In addition, PBC also enforces the foreign exchange settlement and purchase system under which all domestic enterprises and entities without foreign exchange accounts should sell their foreign exchange revenues to the state via banks licensed to conduct foreign exchange transactions. When foreign exchange payments are to be made, they need to purchase foreign exchange funds from the banks, which in most cases, requires the advance approval of the State Administration of Foreign Exchange (SAFE).

2011.

While intervention in the foreign exchange market maintains a stable exchange rate, it also leads to increases in base money. This would normally increase the money supply in the economy, as illustrated in the balance sheet of the central bank (Table B.1).

Such increases in the money supply could lead to high inflation. Thus, the PBC attempts to neutralize the growth in the money supply from its purchases of foreign assets. He, Chu, Shu, and Wong (2005), Ouyang, Rajan, and Willet (2010) and Wang (2010) summarize neutralization tools used by the PBC.

The PBC has two ways to offset the impact on the money supply of a rise in its foreign assets. It can reduce the change in monetary base through open market operations, i.e. contracting domestic credit of the PBC. Alternatively, it can decrease the money multiplier by raising the reserve requirement ratio or imposing bank loan ceilings to slow growth in commercial bank credit. In fact, the PBC both neutralizes by selling central bank bills and raising the required reserve ratio (Wang, 2010). Nonetheless, as some point out (see for example Prasad and Wei, 2005), it is inherently difficult for the PBC to neutralize its purchases of foreign assets while experiencing growing trade surpluses and rising capital inflows.

In this paper, we ask the question: How much of China's foreign asset purchases are effectively neutralized? In other words, can China maintain a sovereign monetary policy under its pegged exchange rate regime? In theory, we know from Mundell (1968) that a central bank can either fix the exchange rate and let the money supply float, or fix the money supply and let the exchange rate float.

To answer this question, we perform empirical tests on monthly data from June 1999 to June 2011 based upon the balance sheet of the PBC. We first estimate the sterilization coefficient over the whole sample period. Second, we estimate neutraliza-

tion for two sub-periods, from 1999:6 to 2005:7, and 2005:8 to 2011:6, after the 2005 currency reform. The empirical results suggest that the central bank in China does not effectively neutralize the contemporaneous change in net foreign assets. Consequently, China's control on domestic credit under a pegged exchange rate regime is weak, confirming Mundell's argument. Indeed, results from the two subsamples imply that neutralization in China is becoming increasingly difficult.

In the literature, several studies are related to ours. Using monthly data from 2000:7 to 2008:9, Ouyang, Rajan, and Willet (2010) estimate a simultaneous equation model based on Brissimis, Gibson, and Tsakalotos (2002), finding China's sterilization is successful. Wang (2010) finds that China's neutralization is almost perfect in terms of monetary base (IFS line 14), but not in terms of M2 (IFS line 34 plus 35) using monthly data from 1999:6 to 2009:3 (from the International Financial Statistics (IFS) database). Using monthly data from 1998:1 to 2004:12, He, Chu, Shu, and Wong (2005) uses Vector Auto Regression (VAR) analysis and finds the PBC's complete and contemporaneous neutralization of its purchases of foreign exchange: a one yuan purchase of net foreign assets leads to a corresponding decline of one yuan in net domestic assets: i.e. a sterilization coefficient of one. Our results suggest that this is no longer the case.

This chapter is organized as follows. Section 4.2 presents the model. Section 4.3 reports our empirical results. Section 4.4 concludes.

4.2 Model Specification

We employ a simple model based on the balance sheet of the central bank as the theoretical basis for our empirical estimation. From the simplified balance sheet of the PBC, a change in monetary base, H, the liabilities of the central bank, equals the change in its assets, either caused by a change in net foreign assets, F, or a change

in net domestic assets, D.

$$\Delta H \equiv \Delta D + \Delta F, \tag{4.1}$$

where Δ is the first difference operator. This identity suggests that an increase in foreign assets, due to the PBC's purchase of foreign exchange, tends to increase the monetary base, and hence the overall money supply.

However, to limit the impact of an increase of net foreign assets on monetary base, the central bank can neutralize the monetary effects by reducing domestic credit. A direct way to measure the effectiveness of PBC's neutralization is to test how domestic assets respond to a change in net foreign assets.

Neutralization can be characterized by a sterilization coefficient, θ, representing the fraction of the change in net foreign assets that is neutralized:

$$\Delta D = -\theta \Delta F, \text{ where } 0 \leqslant \theta \leqslant 1. \tag{4.2}$$

The coefficient of $\theta = 1$ implies full monetary neutralization, or $\Delta D = -\Delta F$ so $\Delta H = 0$, while $\theta = 0$ implies no monetary neutralization. Ouyang, Rajan, and Willet (2010) and Wang (2010) define a "sterilization coefficient" as how much domestic money creation responds to a change in international reserves. Here we define the sterilization coefficient narrowly from the central bank balance sheet perspective. It is the fraction of a change in net foreign assets that is neutralized by a reduction in the PBC's domestic credit.

We also test the impact of an increase in net foreign assets on broader monetary supply such as M2. Recall that the money supply M equals monetary base times the money multiplier m:

$$M = mH, \tag{4.3}$$

which implies for a constant multiplier:

$$\Delta M = m\Delta H = m(\Delta F + \Delta D). \tag{4.4}$$

Substituting Equation (4.2) into Equation (4.4), we have:

$$\Delta M = m(1-\theta)\Delta F. \tag{4.5}$$

where ΔM denotes the change in commercial banks' domestic credit due to a change in net foreign assets, ΔF, and $m(1-\theta)$ is the effective money multiplier, which de-multiplies the effect of the purchase of net foreign assets on the money supply. Note that in the case of perfect neutralization $\theta = 1$, so the effective money multiplier $m(1-\theta) = 0$. In general, $m(1-\theta) > 0$.

We now estimate Equations (4.2) and (4.3) using monthly data from June 1999 to June 2011, then derive the effective money multiplier.

4.3 Empirical Results

Because the PBC passively absorbs the increase in its net foreign assets at pegged rates, we assume the changes in net foreign assets are exogenous, and we estimate Equation (4.2) by OLS regression from 1999:7 to 2011:6.[2] We then estimate the money multiplier based on Equation (4.3). We also run regressions on two subsample periods, from 1999:7 to 2005:7, and 2005:8 to 2011:6 after the 2005 currency reform of managed floating, gradual appreciation of the RMB.

We use M2 (IFS line 34 plus 35) as a measure of the overall money supply, M, in the economy. Data for monetary base, H (IFS line 14), and net foreign assets, F (IFS line 11n), are also from International Financial Statistics database. We calculate

[2]The estimation using quarterly data yields similarly results.

net domestic assets D by subtracting net foreign assets F from monetary base H using Equation (4.1). Following Ouyang, Rajan, and Willet (2010), Wang (2010) and others, we normalize the first difference of F and H by monthly GDP. Because monthly GDP data for China are not available, we construct the series following Wang (2010) by using Industrial value added and quarterly GDP data, which are from the China Statistics Bureau website and CEIC China Premium Database.

We apply the Augmented Dicky-Fuller (ADF) test with the lag length selection by the Akaike Information Criterion (AIC). The results indicate that all variables in our regressions are stationary, as shown in Table B.2, so we may perform a meaningful time series regression.

The inclusion of a lag variable is statistically insignificant. Thus, we regress on the contemporaneous variables for both equations. On the complete sample period, OLS regression gives an estimated sterilization coefficient of -0.66, significant at the 5% level, suggesting that the PBC only neutralizes 66% of the change in net foreign assets effectively. This is illustrated by the negative slope of the regression line in Figure E.1.

The estimated money multiplier is 4.04. This implies an effective money multiplier of 1.4, that is, a one yuan increase in net foreign assets would lead to a 1.4 yuan increase in the money supply due to sterilization. From July 1999 to June 2011, China's foreign reserves rose by 66.9 trillion yuan, at the nominal exchange rate. Given an effective money multiplier of 1.4, this led to an increase of 27.3 trillion yuan in the money supply, which is 41% of the total increase of 66.9 trillion yuan during the same period. In other words, 41% of money growth is accounted for by the growth of foreign reserves alone. This high percentage suggests that the PBC cannot effectively neutralize the change in its net foreign assets, contrary to previous results (He, Chu, Shu, and Wong, 2005; Ouyang, Rajan, and Willet, 2010; Wang, 2010).

Ouyang, Rajan, and Willet (2010) mentioned a selection problem in evaluating the effectiveness of neutralization. That is, if the central bank wants the money supply to grow to accommodate growth in money demand due to growth in real GDP, it would not completely neutralize the increase in reserves. But this does not necessarily mean the central bank has lost control of the domestic monetary process.

Our estimates, however, strongly suggest that the PBC's control over the money supply is weak. Next, we divide the whole sample into two subsamples using the 2005 currency reform as the break point. The first subsample contains 73 observations from July 1999 to July 2005. The second has 71 observations from August 2005 to June 2006. The estimates are significant at the 5% level except for the sterilization coefficient in the second subsample, whose p-value is less than 0.12. Estimated sterilization coefficients for these two subsamples are 0.82 and 0.57 respectively, and for the money multiplier are 5.53 and 3.59 respectively.

The size of the sterilization coefficient decreases considerably in the second period, and the degree of statistical significance is lower, indicating it has become more difficult for the PBC to pursue a policy of neutralization. The decrease in the money multiplier could be explained by the frequent increases of reserve requirements for commercial banks: from July 2006 to September 2008, the legal reserve ratio for commercial banks was raised 19 times, from 8% to 17.5% (Wang, 2010).

Given our estimates for the sterilization coefficient and money multiplier, we estimate the effective money multiplier to be 1.01 for the first period and 1.56 for the second. That is, an increase of 1 yuan in net foreign assets would lead to an increase of approximately 1 yuan in the money supply in period one, and a 1.56 yuan in period two. The substantial increase in the effective money multiplier suggests that neutralization is increasingly difficult for PBC. Growth of foreign reserves accounts for 29% of the increase in the money supply in the first period, but 64% in the second one.

4.4 Conclusion

The empirical results over the whole sample period from 1999:7 to 2011:6 indicate that the People's Bank of China does not effectively neutralize the contemporaneous change in net foreign assets, contrary to most studies. An estimated neutralization coefficient of 0.66 and money multiplier of 4.04 yields an effective money multiplier of 1.4 with sterilization. That is, a 1 yuan increase in net foreign assets leads to an increase of 1.4 yuan in the money supply. Thus, 41% of the increase in money supply during the whole sample period is accounted for by the rapid accumulation of foreign reserves. This strongly suggests that China cannot retain its control on the money supply under a pegged exchange rate regime, as Mundell (1968) hypothesized. Mundell noted that the monetary authorities can fix the exchange rate and let the money supply float, or fix the money supply and let the exchange rate float: but it cannot fix both the exchange rate and the money supply. Our results from two subsamples further suggest that a policy of neutralization is increasingly difficult in the face of the rapid buildup of foreign reserves of the PBC.

CHAPTER 5
China's Implicit Demand for Foreign Reserves

5.1 Background

In the aftermath of the Asian Financial Crisis, China has accumulated large stockpiles of foreign reserves. In 2007, foreign reserves in China topped $1,528 billion, compared to that of $52 billion in 1994. However, China's reserve hoardings cannot be compared through time unless they are scaled in some way, since the economy has also been developing rapidly during the same period. In Table C.1, we choose GDP, broad money (M2), imports and foreign debt as the scale variables. The ratios are widely used as rules of thumbs for the central banks to manage their reserves. As is shown in the table, the reserves-to-GDP ratio rises notably from 8.9% in 1994 to 43.6% in 2007, while the reserves-to-M2 ratio increases from 9.5% in 1994 to 26.7% in 2007. Reserves scaled by foreign debt boosts radically from 55.6% to 409% during the same period. In 1994, China's foreign reserve holdings could cover 23 weeks of imports, while in 2007 it is equivalent to 86 weeks of imports. In short, the reserve holdings in China have been surging dramatically even when considered in relative terms.

There were many significant changes in China's foreign exchange system in the 1990s. In 1994, the People's Bank of China (PBC) replaced dual exchange rates with a single exchange rate, i.e. the so-called managed floating exchange based on market supply and demand. In fact, it was actually a pegged exchange rate. The PBC also set up a centralized, automatic price-matching, inter-bank foreign exchange market and implemented the system of foreign exchange settlement and sales, under which companies are required to sell the foreign exchange earned to the central bank. Until

now, private Chinese citizens could only buy limited amounts of foreign exchange for travel or education abroad, with the permission of the authorities The State Administered Foreign Exchange (SAFE) system. Since then, China's balance on both the current account and the capital account has been positive, with the exception of a negative capital account in 1998. Under a pegged exchange rate, the PBC has to purchase foreign exchange when there is a surplus at the lower intervention point. At the same time, the purchase of foreign exchange injects base money, normally resulting in a multiplied increase in the money supply. To avoid inflation, the PBC uses many monetary policy tools such as repos, the required reserve rate, a benchmark interest rate, and open market operations, in order to neutralize the increase in money supply. However, the neutralization policy also boosts China's demand for foreign reserves by a multiplier a magnification effect determined by the neutralization coefficient of sterilization operations. While ultimately inconsistent, this leads to a continuous growth in foreign exchange reserves.

We use a monetary model for China's demand for foreign reserves that explicitly takes into account neutralization policy. Based on the above reasons, we use quarterly data from 1994 to 2007, the period of rising foreign exchange reserves. According to Chow (1985), when the sample observation is greater than 40, we could have serious time series results. The data are mainly from the International Monetary Fund's International Financial Statistics (IFS), the website of the PBC the National Bureau of Statistics of China, and China Stock Market Accounting Research (CSMAR) database.

This chapter is organized as follows. Section 5.2 briefly states the essential elements of the monetary model, and derives the equation to be tested for China's demand for international reserves with neutralization from 1994:1 to 2007:4. Section 5.3 reports empirical results for the model. Section 5.4 concludes.

5.2 Model Specification

Following Connolly and Silveira (1979), Obstfeld (1982), Siklos (2000) and Sarno and Taylor (2001),we establish a monetary model for China's demand for foreign reserves.

According to Keynes, there are three distinct motives of holding money balances the transactions motive, the precautionary motive and portfolio speculative motive. The demand for money is thus a function of real income and the opportunity cost of holding money:

$$L = PY^{\phi_1}e^{-\phi_2 i}, \tag{5.1}$$

where L is the nominal demand for money, P is price level, Y is real income and i is interest rate. The parameters ϕ_i and ϕ_2 are respectively the income elasticity of demand for money and the semi-elasticity of demand for money with respect to the interest rate.

Second, we specify the money supply process. The money supply equals high-powered money times the money multiplier:

$$M = mH = m(F + D), \tag{5.2}$$

where M is the money supply, m the money multiplier, H high-powered money, F the stock of international reserves in RMB and D domestic credit respectively.

Third, we assume that purchasing power parity holds for China:

$$P = sP^*, \tag{5.3}$$

where P is the Chinese price index, s the spot price of foreign exchange in terms of RMB, and P^* the world price level. Finally, we assume that money supply adjusts

rapidly to the quantity demanded, so that monetary equilibrium holds:

$$M = L. \tag{5.4}$$

Under this simple monetary model, foreign reserves become endogenous:

$$F = g(P, Y, i, m, D). \tag{5.5}$$

From the above equations we can see that the domestic credit and foreign exchange reserves are negatively related. A domestic credit contraction by the central bank leads to an inflow of foreign reserves with a stable exchange rate. On the other hand, an expansionary credit policy would be offset by the decrease of foreign reserves.

In general, the change in base money is indicated by:

$$\Delta H_t = \Delta F_t + \Delta D_t. \tag{5.6}$$

When the PBC perfectly neutralizes a purchase in foreign reserves, $\Delta H = 0$, we have:

$$\Delta D = -\Delta F. \tag{5.7}$$

Partial neutralization can be characterized by a sterilization coefficient, θ, representing the fraction of foreign exchange purchases that are neutralized. That is,

$$\Delta D = -\theta \Delta F, \text{ where } 0 \leqslant \theta \leqslant 1. \tag{5.8}$$

Combining other factors that would influence ΔF, we estimate Equation (5.9) in the

following section:

$$\Delta F = \alpha_0 + \alpha_1 \Delta D_1 + \alpha_2 \Delta Y_t + \alpha_3 \Delta P_t + \alpha_4 \Delta Def_t + \alpha_5 \Delta i_t + \mu_t, \tag{5.9}$$

where ΔF is the change of foreign reserves in billions of RMB. We use the data from IFS. They are converted to RMB at the annual average exchange rate. Following Xu (2003), ΔD is proxied by the changes of currency in circulation. ΔY_t is the change in real GDP. Because the data published by the National Bureau of Statistics of China are not seasonally adjusted, we adjust them using the multiplicative moving average method. We use the Consumer Price Index (CPI) data to measure the inflation rate ΔP_t. ΔDef_t Chinese government deficit, where Def_t is the difference between fiscal revenue and fiscal expenditure in year t, which comes from the CSMAR database. $\Delta i_t = i_t - i_t^*$ is the difference between the Chinese and American interest rates, where the it is the one year lending rate in China, and i_t is the US Treasury bill rate percentage per annum from the IFS.

Theoretically, we can denote the neutralization coefficient as:

$$\theta = -\frac{1}{m \times a_1}, \tag{5.10}$$

where $0 \leqslant \theta \leqslant 1$. $\theta = 1$ stands for perfect neutralization; $0 < \theta < 1$ stands for partial neutralization; when there is no neutralization operation, $\theta = 0$.

From the above model we have:

$$\Delta F = \frac{\Delta L}{m} - \Delta D. \tag{5.11}$$

Considering the neutralization policy $\Delta D = -\theta \Delta D$, we have:

$$\Delta F = \frac{\Delta L}{m(1-\theta)}. \tag{5.12}$$

Because of its foreign exchange management system, the policy of neutralization in China boosts foreign reserve accumulation by a magnification or multiplier effect equal to $1/(1-\theta)$.

5.3 Empirical Results

In literature, there are mainly three ways to estimate the neutralization coefficient: the reduced model, the two-stage least squares regression (2SLS) model, and the vector autoregression (VAR) model. Since there is an endogeneity problem in the OLS method, and a misspecification problem in the 2SLS model Obstfeld (1982), here we choose the VAR model to estimate Equation (9). VAR modeling is a dynamic system using lagged variables in reduced form. In our estimation, the minimum information criteria will be used to identify the correct lag length of the VAR model. Misspecification tests will be applied including lag length. The individual F-statistic value and log likelihood (AIC and SIC) of the system will be considered when choosing the appropriate model. Since the individual coefficients in the VAR model are difficult to interpret, the impulse respond function (IRF) is used to test the theory's predictions. The IRF traces out the response of the dependent variable in the VAR system to the shock in the error terms. The IRF is constructed using the estimated coefficients, so it reports the confidence level.

As is shown in Table C.2, results from the augmented Dickey-Fuller (ADF) unit root tests show that all the variables are stationary at the 5% level. The lag length tests show the minimum information value to be three quarterly lags. The estimation results estimate the R2 value for the reserve demand equation to be approximately

0.63. The IRF using VAR(3) is reported in Figure F.1.

According to the theoretic model, the IRF shows domestic credit has a negative relationship with foreign reserves with a coefficient of 0.73, as the theory of monetary approach expected. The significant negative effect lasts until the 10th quarterly lag. Since the average money multiplier for the sample period is approximately 2.43, according to our theoretical model, we have a neutralization coefficient of approximately 0.57. According to Ouyang, Rajan, and Willet (2010), since December 2002 net domestic assets of PBC's balance sheet have remained low, if not negative. Open market operations and raising reserve requirements are often used by the PBC to control domestic credit and money supply. As the theory suggests, the neutralization policy contracts domestic credit and leads to a further increase in foreign reserve accumulation. The magnification effect, $1/(1-\theta)$, equals 2.23.

As shown in Figure F.1, the empirical evidence on price level and interest rate is also consistent with the monetary model. The price level is positively related to foreign reserves, while foreign reserves respond negatively to the interest rate, at least initially. However, the results indicate that real income has a negative relationship to foreign reserves, which supports the prediction of the Keynesian school. It suggests that when income increases, spending will increase also, thus leading to reserve outflow. The response of the fiscal deficit is economically insignificant.

Another useful tool to analyze the results generated by the VAR model is variance decomposition. It gives the relative importance of each shock to the variables in the VAR. The variance decomposition for ΔF shows its variance can be explained by the innovations of itself and other variables in the model, especially domestic credit and interest rate shocks. The variance decomposition for ΔD indicates that in the forecast period, the variance of domestic credit changes are mainly influenced by the innovations of itself and foreign reserve changes. These results are consistent with

the monetary approach, which states that foreign reserves and domestic credit are negatively related in an open economy with a stable exchange rate.

5.4 Conclusion

Because of China's unique foreign exchange system and its neutralization policy, contraction in domestic credit boosts foreign reserve accumulation by a "magnification or multiplier" effect equal to approximately 2.3 times the purchase of foreign exchange by the PBC due to a neutralization coefficient of 0.57. While ultimately inconsistent, the neutralization policy leads to growth in foreign exchange reserves that seems limitless.

REFERENCES

Andronova, N., and M. Schlesinger, 2001, "Objective Estimation of the Probability Density Function for Climate Sensitivity," *Journal of Geophysical Research*, 106(22), 605–22.

Baker, M., and G. Roe, 2009, "The Shape of Things to Come: Why Is Climate Change So Predictable?," *Journal of Climate*, 22, 4574–89.

Bartz, S., and D. Kelly, 2008, "Economic Growth and the Environment: Theory and Facts," *Resource and Energy Economics*, 30, 115–49.

Brissimis, S., H. Gibson, and E. Tsakalotos, 2002, "A Unifying Framework for Analyzing Offsetting Capital Flows and Sterilization: Germany and the ERM," *International Journal of Finance and Economics*, 7, 63–78.

Cai, Y., and K. L. Judd, 2012, "Shape-preserving Dynamic Programming," *Mathematical Methods of Operations Research*, forthcoming.

Chow, G., 1985, "A Model of Chinese National Income Determination," *Journal of Political Economy*, 93, 356–376.

Connolly, M., and J. Silveira, 1979, "Exchange Market Pressure in Postwar Brazil: An Application of the Girton-Roper Monetary Model," *Journal of Political Economy*, 69, 448–454.

Cooley, T., and E. Prescott, 1995, "Economic Growth and Business Cycles," in *Frontiers of Business Cycle Research*, ed. by T. Cooley. Princeton University Press, Princeton, NJ, chap. 1, pp. 1–38.

Costello, C., M. Neubert, S. Polasky, and A. Solow, 2010, "Bounded Uncertainty and Climate Change Economics," *Proceedings of the National Academy of Sciences*, 107(18), 8108–10.

Du, P., C. F. Parmeter, and J. S. Racine, 2013, "Nonparametric Kernel Regression with Multiple Predictors and Multiple Shape Constraints," *Statistica Sinica*, forthcoming.

Fidrmuc, J., 2010, "Time-varying Exchange Rate Basket in China from 2005 to 2009," *Comparative Economic Studies*, 52(4), 515–529.

Forest, C., P. Stone, and A. Sokolov, 2005, "Estimated PDFs of Climate System Properties Including Natural and Anthropogenic Forcings," working paper Report 125, MIT Joint Program on the Science and Policy of Global Change.

Frankel, J. A., 2009, "New Estimation of Chinas Exchange Rate Regime," *Pacific Economic Review*, 14(3), 346–360.

Geweke, J., 2001, "A Note on Some Limitations of CRRA Utility," *Economic Letters*, 71, 341–5.

Goldstein, M., and N. R. Lardy, 2006, "Chinas Exchange Rate Policy Dilemma," *American Economic Review*, 96, 422–426.

Gregory, J., et al., 2002, "An Observationally Based Estimate of the Climate Sensitivity," *Journal of Climate*, 15(22), 3117–21.

Gregory, J., R. Stouffer, S. Raper, P. Stott, and N. Rayner, 2002, "An Observationally Based Estimate of the Climate Sensitivity," *Journal of Climate*, 15(22), 3117–21.

Hall, P., and L.-S. Huang, 2001, "Nonparametric Kernel Regression Subject to Monotonicity Constraints," *Annals of Statistics*, 28, 624–647.

He, D., C. Chu, C. Shu, and A. Wong, 2005, "Monetary Management in Mainland China in the Face of Large Capital Inflows," Working Papers 0507, Hong Kong Monetary Authority.

Intergovernmental Panel on Climate Change, 2007, "Climate Change 2007: The Physical Science Basis: Summary for Policy Makers, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change," working paper, IPCC Secretariat, World Meteorological Organization, Geneva, Switzerland, `http://www.ipcc.ch/ipccreports/ar4-wg1.htm`.

Judd, K. L., 1998, *Numerical Methods in Economics*. Cambridge University Press, Cambridge, MA.

Keller, K., B. Bolker, and D. Bradford, 2004, "Uncertain Climate Thresholds and Optimal Economic Growth," *Journal of Environmental Economics and Management*, 48, 723–41.

Kelly, D. L., and C. D. Kolstad, 1999a, "Bayesian Learning, Pollution, and Growth," *Journal of Economic Dynamics and Control*, 23(4), 491–518.

——— , 1999b, "Integrated Assessment Models for Climate Change Control," in *International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues*, ed. by T. Tietenberg, and H. Folmer. Edward Elgar, Cheltenham, UK, chap. 5, pp. 171–97.

——— , 2001, "Solving Infinite Horizon Growth Models with an Environmental Sector," *Journal of Computational Economics*, 18, 217–31.

Kelly, D. L., C. D. Kolstad, M. Schlesinger, and N. G. Andronova, 1998, "Learning About Climate Sensitivity From the Instrumental Temperature Record," University of Mimai.

Leach, A. J., 2007, "The Climate Change Learning Curve," *Journal of Economic Dynamics and Control*, 31, 1728–52.

Lemoine, D. M., 2010, "Climate Sensitivity Distributions Depend on the Possibility that Models Share Biases," *Journal of Climate*, 23(16), 4395–4415.

Lemoine, D. M., and C. P. Traeger, 2013, "Watch Your Step: Optimal Policy in a Tipping Climate," *American Economic Journal: Economic Policy*, forthcoming.

Manne, A., R. Mendelsohn, and R. Richels, 1993, "MERGE: A Model for Evaluation Regional and Global Effects of GHG Reduction Policies," *Energy Policy*, 23(1), 17–34.

Millner, A., 2011, "On Welfare Frameworks and Catastrophic Climate Risks," University of California at Berkeley Working Paper.

Mundell, R. A., 1968, *Barter Theory and the Monetary Mechanism of Adjustment*. Macmillan, New York.

Nadaraya, E., 1965, "On Non-parametric Estimates of Density Functions and Regression Curves," *Theory of Probability & Its Applications*, 10(1), 186–190.

Newbold, S., and A. Daigneault, 2009, "Climate Response Uncertainty and the Benefits of Greenhouse Gas Emissions Reductions," *Environmental and Resource Economics*, 44, 351–77.

Nicholls, R., R. Tol, and A. Vafeidis, 2008, "Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet: an Application of FUND," *Climatic Change*, 91, 171–91.

Nordhaus, W., 2007a, "DICE Model Version as of May 22, 2007," `http://www.econ.yale.edu/nordhaus/DICEGRAMS/DICE2007.htm`.

——— , 2007b, "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, 45, 686–702.

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

——— , 2010, "RICE Model Version as of March 20, 2012," `http://nordhaus.econ.yale.edu/RICEmodels.htm`.

——— , 2011, "The Economics of Tail Events with an Application to Climate Change," *Review of Environmental Economics and Policy*, 5(2), 240–57.

Obstfeld, M., 1982, "Can we Sterilize? Theory and Evidence," *American Economic Review*, 72(2), 45–50.

Ouyang, Y. A., S. R. Rajan, and T. D. Willet, 2010, "China as a Reserve Sink: the Evidence from Offset and Sterilization Coefficients," *Journal of International Money and Finance*, 29(10), 951–972.

Pindyck, R., 2011, "Fat Tails, Thin Tails, and Climate Change Policy," *Review of Environmental Economics and Policy*, 5(2), 258–74.

Prasad, E., and S.-J. Wei, 2005, "The Chinese Approach to Capital Inflows: Patterns and Possible Explanations," IMF Working Papers 05/79, International Monetary Fund.

Racine, J., and Q. Li, 2004, "Nonparametric Estimation of Regression Functions with Both Categorical and Continuous Data," *Journal of Econometrics*, 119(1), 99–130.

Roe, G., and M. Baker, 2007, "Why is Climate Sensitivity so Unpredictable?," *Science*, 318(5850), 629–32.

Santos, M. S., 1999, "Numerical solution of dynamic economic models," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford. Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 5, pp. 311–386.

Sarno, L., and M. P. Taylor, 2001, "Official Intervention in the Foreign Exchange Market: Is It Effective and, If So, How Does It Work?," *China and World Economy*, 39(3), 839–868.

Schwartz, S., 2007, "Heat Capacity, Time Constraint, and Sensitivity of the Earth's Climate System," *Journal of Geophysical Research*, 112, D24S05.

Siklos, P. L., 2000, "Capital Flows in a Transitional Economy and the Sterilization Dilemma: the Hungarian Experience, 1992-1997," *Policy Reform*, 3, 373–392.

Tol, R., 2009, "The Economic Effects of Climate Change," *Journal of Economic Perspectives*, 23(2), 29–51.

Traeger, C. P., 2012, "A 4-Stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change," working paper 1130, CUDARE Working Paper.

Wang, Y., 2010, "Effectiveness of Capital Controls and Sterilizations in China," *China and World Economy*, 18(3), 106–124.

Watson, G. S., 1964, "Smooth Regression Analysis," *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372.

Weitzman, M., 2009, "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *Review of Economics and Statistics*, 91, 1–19.

——— , 2011, "Fat Tailed Uncertainty and the Economics of Catastrophic Climate Change," *Review of Environmental Economics and Policy*, 5(2), 275–92.

Xu, C., 2003, *Issues of China's Foreign Exchange Reserves*. China Statistics Press, Beijing.

APPENDIX A
Learning and Climate Feedbacks Tables

Variable Definitions

Variable Definitions

	State Variables
k	Productivity adjusted capital stock per capital
m	GHG concentration over preindustrial level
T	Atmospheric temperature
μ	Mean estimate of feedback parameter
S	Variance of the prior distribution of feedback parameter
	Control Variables
E	GHG emissions
k'	Gross investment
	Random Variables
ν'	Weather shock
β_1	Feedback parameter

Table A.1 lists the definitions of the variables in the integrated assessment model.

Calibrated Parameter Values

Parameter	Description	Calibrated Value
β	Discount factor	0.952
σ	Coefficient of risk aversion	1.5
ψ	Capital share	0.402
ϵ	Emissions share	0.0057
δ_k	Capital depreciation rate	0.046
η	Population growth rate	0.011
ϕ	Productivity growth rate	0.018
δ_m	GHG stock decay	0.0083
b_1	Damage function parameter	0.003
b_2	Damage function parameter	2
γ	Initial emissions intensity (GtC/trillion 2005 $)	4.66
MB	Preindustrial GHG concentrations (GtC)	596.4
α	Heat capacity of the ocean	0.22^{-1}
Γ	Preindustrial temperature	-0.4607
Ω	Forcing parameter (W/M^2)	4.39
σ_ν	Standard deviation of weather shocks	0.11
$\max \Delta T_{2\times}$	Maximum temperature change	26.83
ζ	Heat transfer from atmosphere to ocean	0.05
τ	Heat transfer from ocean to atmosphere	0.30

Table A.2 lists the calibrated parameter values.

Initial Conditions

State Variable	Initial Value	Units
K	137	Trillions of 2005 dollars
m	1.3563	Fraction relative to preindustrial
T	0.731	°C above preindustrial
μ	0.65	Watts per square meter (W/M^2)
$\sqrt{S}$	0.13	Watts per meter squared (W/M^2)

Table A.3 shows the initial conditions corresponding to 2005. Dollar units are adjusted for purchasing power parity. Initial values for K, m, and T are from Nordhaus (2008). Initial values for μ and S are from Roe and Baker (2007).

Expected Learning Time

Test	99%	99.9%	Full
First reject, not subsequently not rejected	8.99	14.58	85.35
First reject, not subsequently not rejected, with ocean	15.84	24.39	110.52
First reject	6.86	10.43	71.06
First Reject, with ocean	11.18	17.43	94.94

Table A.4 shows the expected learning time in years conditional on current information. Expected number of years until the hypothesis $\Delta T_{2\times} \geq \Delta T^*_{2\times} + 1.5$ is rejected at the given confidence level, where $\Delta T^*_{2\times}$ is the true value. Column 4 is the expected number of years until full learning is complete.

Difference in Optimal Emissions Policy

Year	2005	2010	2015	2020	2050
Emissions: Learning (% vs certainty)	19.30	7.68	3.69	1.42	-0.49
Emissions: No learning (% vs certainty)	38.31	46.3	48.17	48.02	45.72
Control rate: Learning (% vs no learning)	49.55	81.36	81.36	75.54	54.76
Carbon tax: Learning (% vs no learning)	23.44	40.71	45.50	46.72	45.96
Probability $T \geq 4.26$	0.165	0.013	0.0002	$< 10^{-6}$	$< 10^{-6}$
95% confidence interval lower bound, $\Delta T_{2\times}$	0.97	2.07	2.39	2.51	2.71
95% confidence interval upper bound, $\Delta T_{2\times}$	7.09	3.57	3.30	3.19	3.00

Table A.5 shows the difference in optimal emissions policy, learning, no learning, and perfect information. The true value is the prior. The first two rows give the percent difference between emissions under certainty and emissions under learning and no learning. That is the first cell of the table indicates emissions under learning are 19.3% lower than emissions under certainty. Rows 3-4 give the percent difference in policies under no learning and the policies under learning. The last two rows give the progress of partial learning versus full learning. All results are the mean of 150 simulations.

Difference in Optimal Emissions Policy (High β)

Year	2005	2010	2015	2020	2050
Emissions: Learning (% vs certainty)	35.62	1.54	0.26	-0.11	-0.24
Emissions: No learning (% vs certainty)	47.31	46.28	45.31	44.63	42.96
Control rate: Learning (% vs no learning)	11.99	38.08	34.22	31.42	25.07
Carbon tax: Learning (% vs no learning)	18.07	42.91	44.11	44.01	42.95
Probability $T \geq 4.26$	0.165	0.014	0.0003	$< 10^{-5}$	$< 10^{-6}$
95% confidence interval lower bound, $\Delta T_{2\times}$	0.97	2.05	2.37	2.5	2.69
95% confidence interval upper bound, $\Delta T_{2\times}$	7.09	3.58	3.31	3.21	3.02

Table A.6 shows the difference in optimal emissions policy, learning, no learning, and perfect information, when $\beta = 0.97$. The true value is the prior. The first two rows give the percent difference between emissions under certainty and emissions under learning and no learning. Rows 3-4 give the percent difference in policies under no learning and the policies under learning. The last two rows give the progress of partial learning versus full learning. All results are the mean of 150 simulations.

APPENDIX B
Neutralization In China Tables

Balance Sheet of Monetary Authority

Assets	Liabilities and Equity
Foreign Assets	Currency in Circulation and Deposits (Monetary Base)
Domestic Assets	Foreign Liabilities
Other Assets	Domestic Liabilities
	Other Liabilities
	Equity

Table B.1 illustrates the balance sheet of the central bank. Source: IMF's International Financial Statistics (IFS) database.

Augmented Dickey-Fuller Unit Root Tests

Variable	Deterministic Terms	Lags	Test Value
ΔF	Constant	1	-9.57***
ΔD	Constant	1	-5.20***
H	Constant, Trend	1	-3.55***
M	Constant, Trend	1	-7.03***

Table B.2 provides results from augmented Dickey-Fuller unit root tests. Statistical significance at the .10, .05, .01 level is denoted by *, **, ***, respectively.

APPENDIX C

China's Implicit Demand for Foreign Reserves Tables

Relative Ratios of China's Foreign Reserves

	Reserves	Reserves/GDP	Reserves/M2	Reserves/Debt	Weeks
1994	51,620	8.9%	9.5%	55.6%	33
1995	73,579	9,7%	10.1%	69.0%	30
1996	105,029	11.8%	11.5%	90.3%	41
1997	139,890	14.2%	12.7%	106.8%	53
1998	144,959	13.9%	11.5%	99.3%	56
1999	154,675	14.1%	10.7%	101.9%	50
2000	165,574	13.9%	10.2%	113.6%	40
2001	212,165	16.1%	11.1%	114.8%	47
2002	286,407	19.7%	12.8%	153.7%	52
2003	403,251	24.5%	15.1%	193.2%	53
2004	609,932	31.5%	19.9%	246.4%	59
2005	818,872	36.0%	22.5%	291.4%	67
2006	1,066,340	40.6%	24.6%	330.1%	73
2007	l,528,250	43.6%	26.7%	409.0%	86

Table C.1 shows China's foreign reserves in Millions of USD, the ratio of foreign reserves to GDP, the ratio of foreign reserves to M2, the ratio of foreign reserves to foreign debt, and the number of weeks of imports of goods and services reserves could pay for.

Augmented Dickey-Fuller Unit Root Tests

Variable	Deterministic Terms	Lags	Test Value
ΔF	Constant	0	-3.33***
ΔD	Constant	4	-3.04**
ΔY	Constant	0	-6.69***
ΔP	Constant, Trend	7	-4.66***
ΔDef	Constant, Trend	2	-66.73***
Δi	Constant	1	-3.37***

Table C.2 provides results from augmented Dickey-Fuller unit root tests. Statistical significance at the .10, .05, .01 level is denoted by *, **, ***, respectively.

APPENDIX D
Learning and Climate Feedbacks Figures

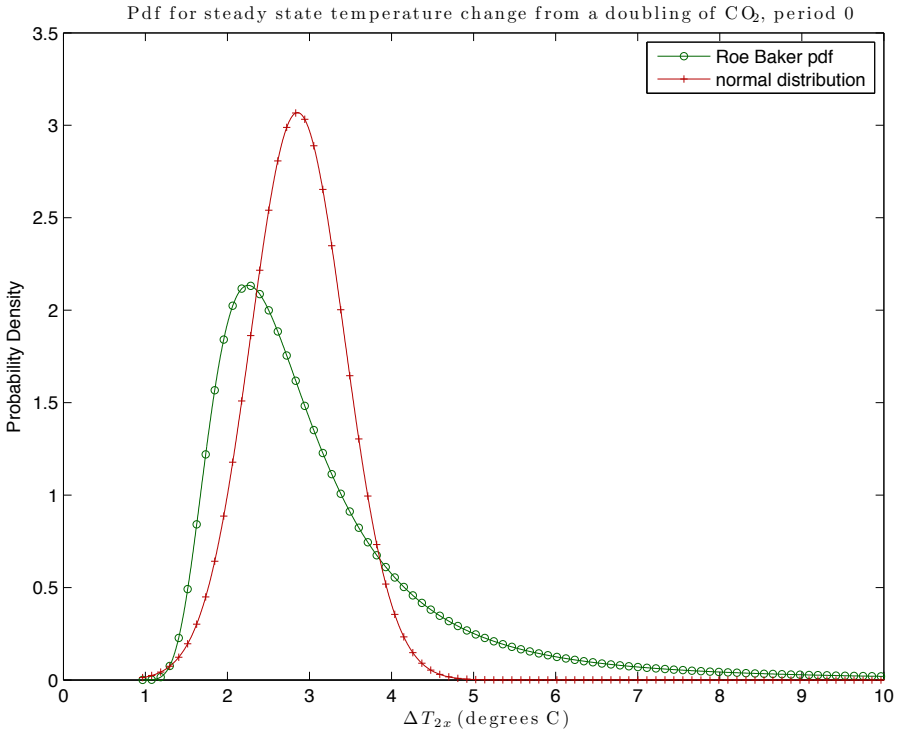

Figure D.1 shows the prior PDF for the climate sensitivity with $\beta_1 \sim N(0.65, 0.013)$.

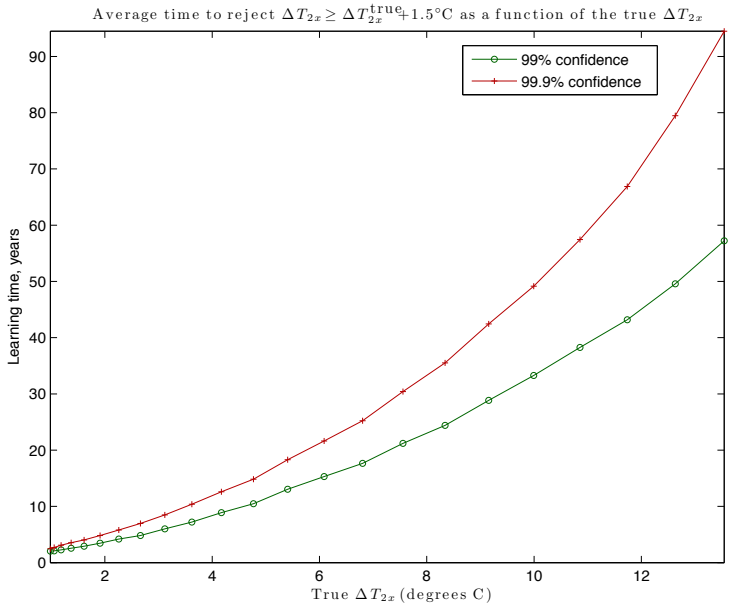

Figure D.2 shows the mean learning time to reject $\Delta T_{2\times} \geq \Delta T_{2\times}^{*} + 1.5°\text{C}$ at 1% and 0.1% critical values as a function of the true $\Delta T_{2\times}$ (an average of 150 simulations for each value of $\Delta T_{2\times}^{*}$).

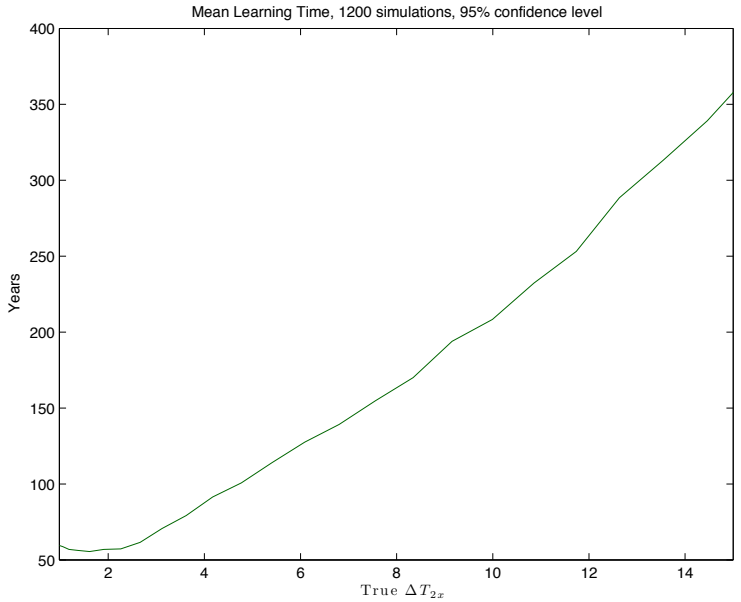

Figure D.3 shows the mean learning time required to reject the hypothesis that $\Delta T_{2\times} \leq 0.95 \cdot \Delta T_{2\times}^{*}$ and $\Delta T_{2\times} \geq \Delta 1.05 \cdot T_{2\times}^{*}$ with 95% confidence (an average of 1200 simulations).

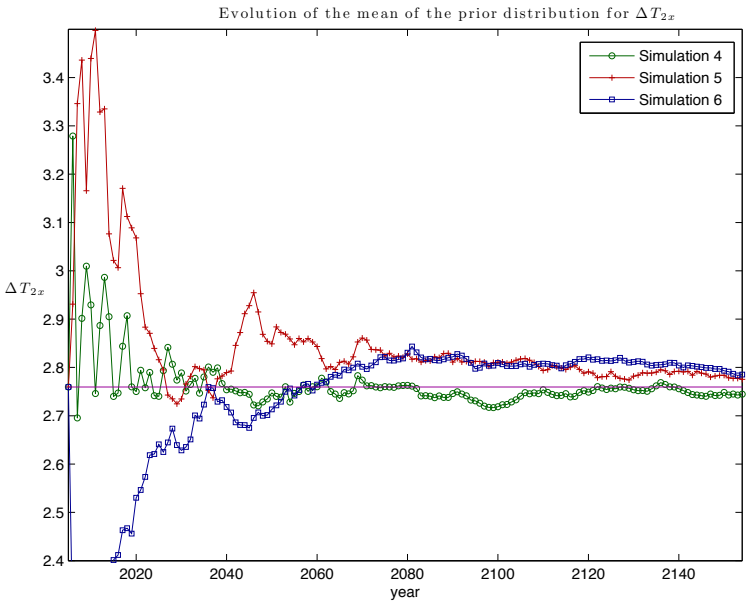

Figure D.4 shows the learning dynamics for the mean of the feedback parameter. The solid line is the true value and is also the ideal case of learning when all weather shocks are zero.

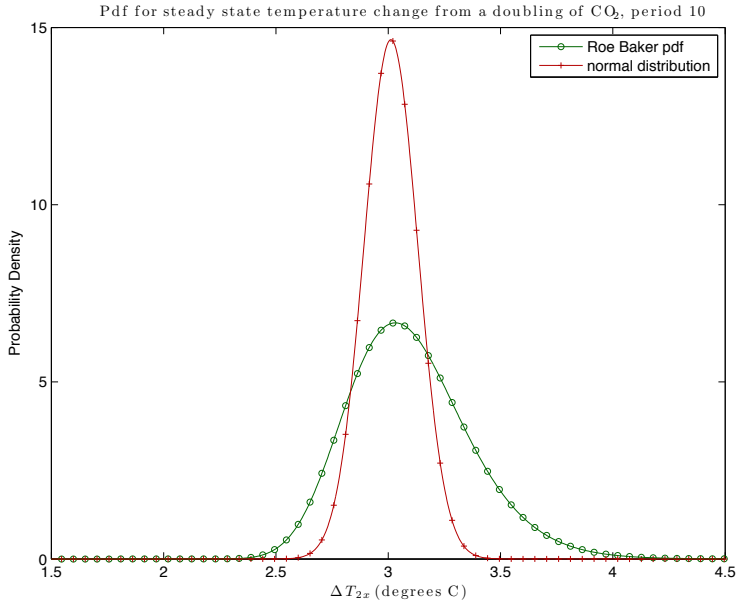

Figure D.5 shows the posterior PDF for the climate sensitivity after 10 observations in one simulation. The true value equals the prior.

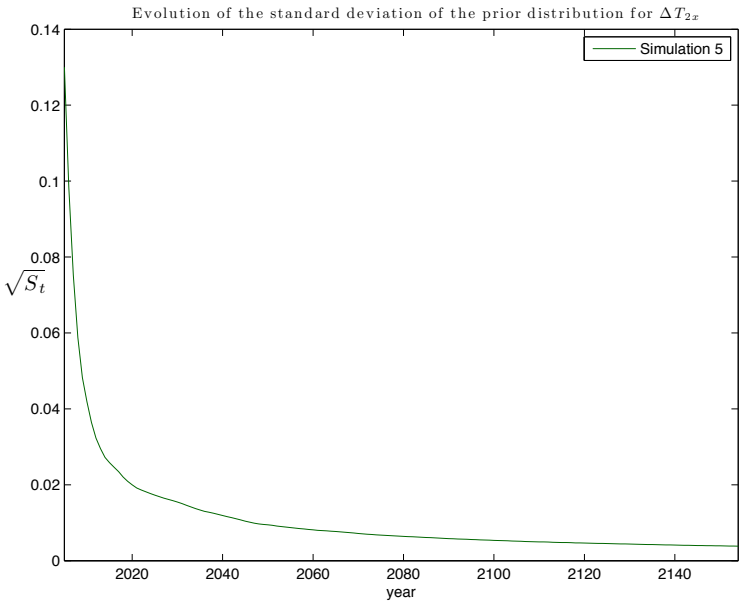

Figure D.6 shows the learning dynamics for the variance of the feedback parameter. The true value equals the prior.

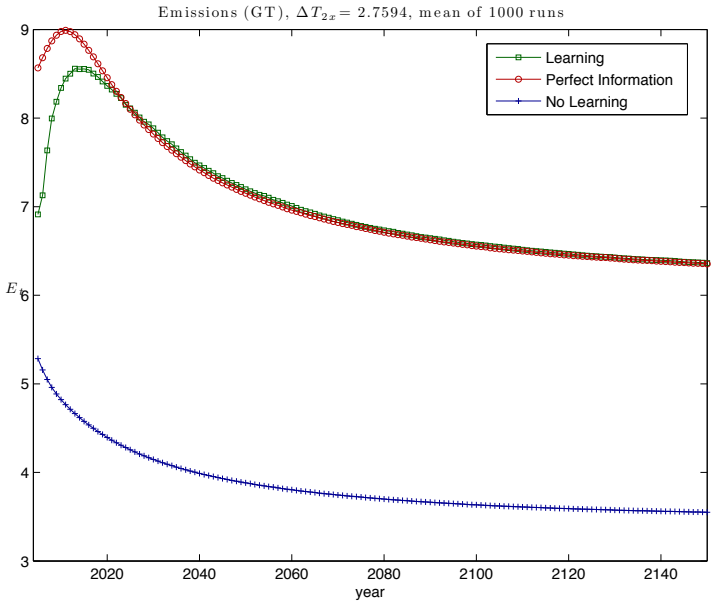

Figure D.7 shows the optimal emissions policy in gigatons (GT), with true $\Delta T_{2\times} = 2.76$ (mean of 1000 runs).

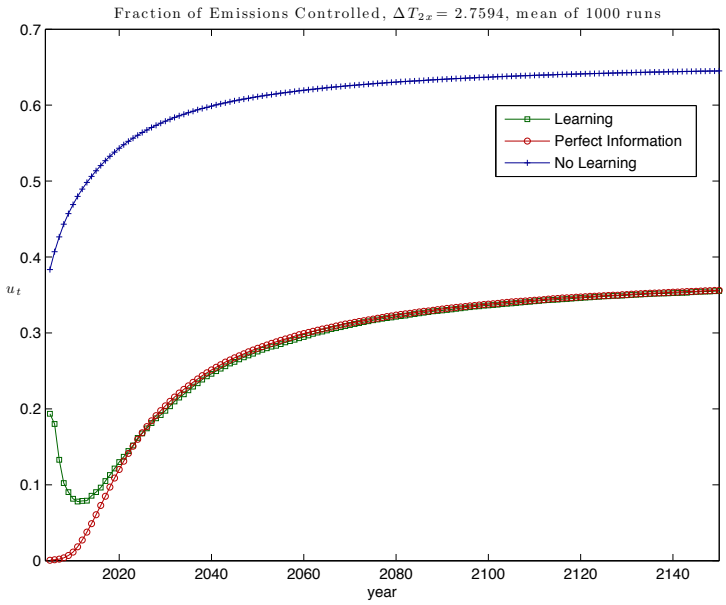

Figure D.8 shows the optimal emissions control rate, with true $\Delta T_{2\times} = 2.76$ (mean of 1000 runs).

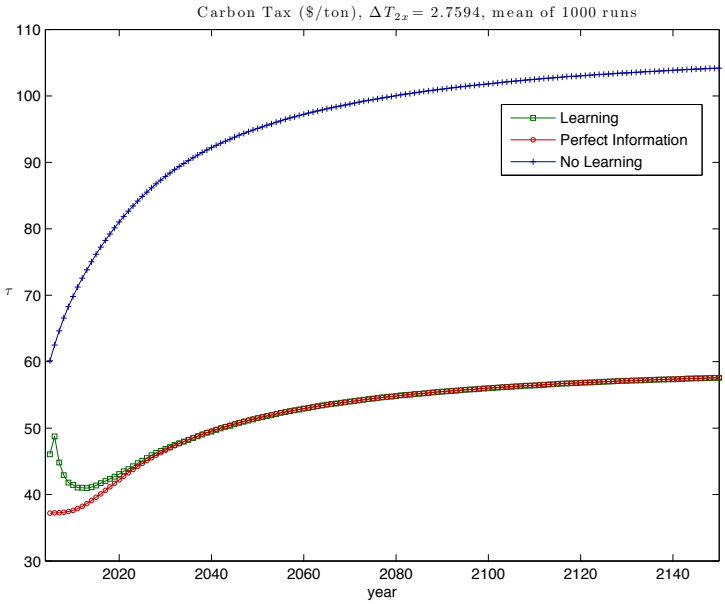

Figure D.9 shows the optimal carbon tax ($/ton), with true $\Delta T_{2\times} = 2.76$ (mean of 1000 runs).

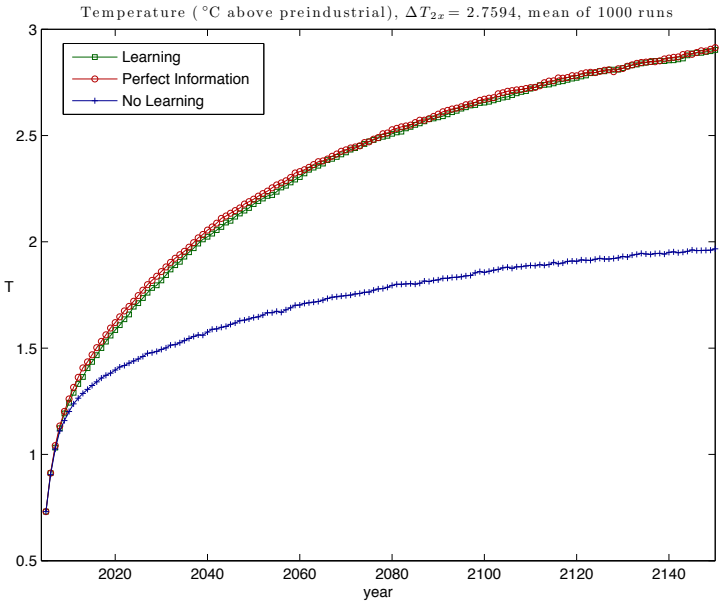

Figure D.10 shows the temperature, degrees C above preindustrial, with true $\Delta T_{2\times} = 2.76$ (mean of 1000 runs).

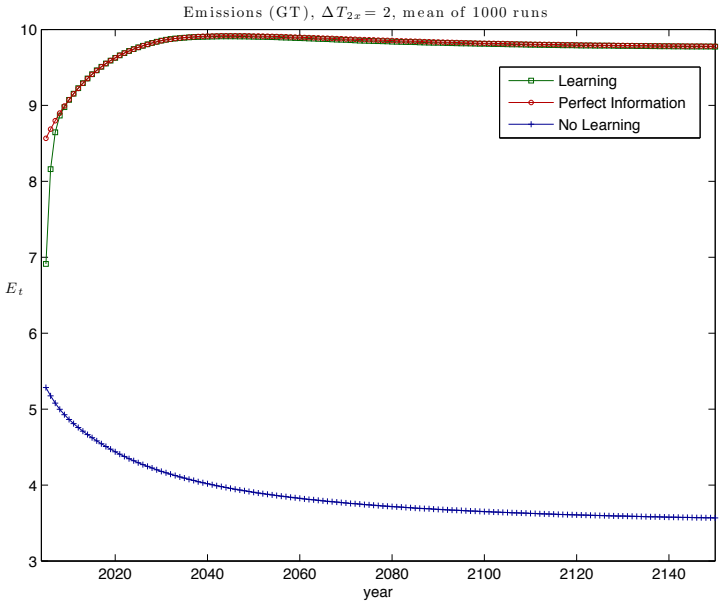

Figure D.11 shows the optimal emissions policy (GT), with true $\Delta T_{2\times} = 2$ (mean of 1000 runs).

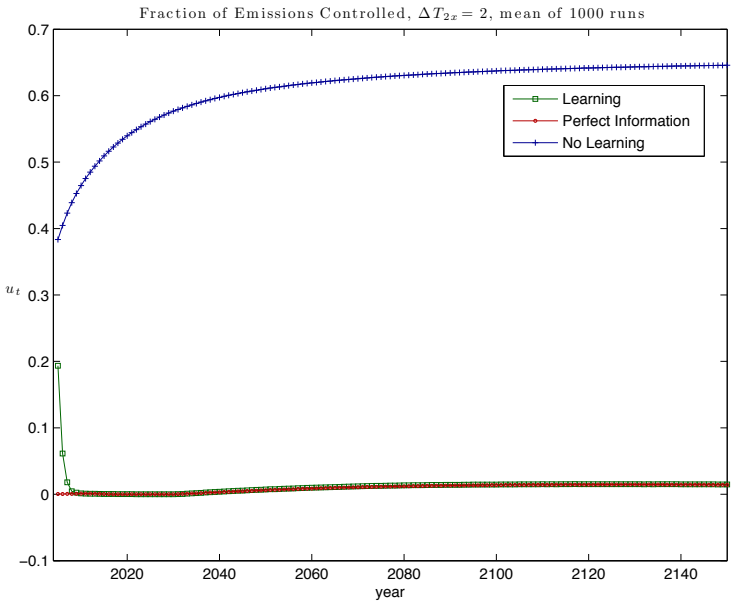

Figure D.12 shows the optimal emissions control rate, with true $\Delta T_{2\times} = 2$ (mean of 1000 runs).

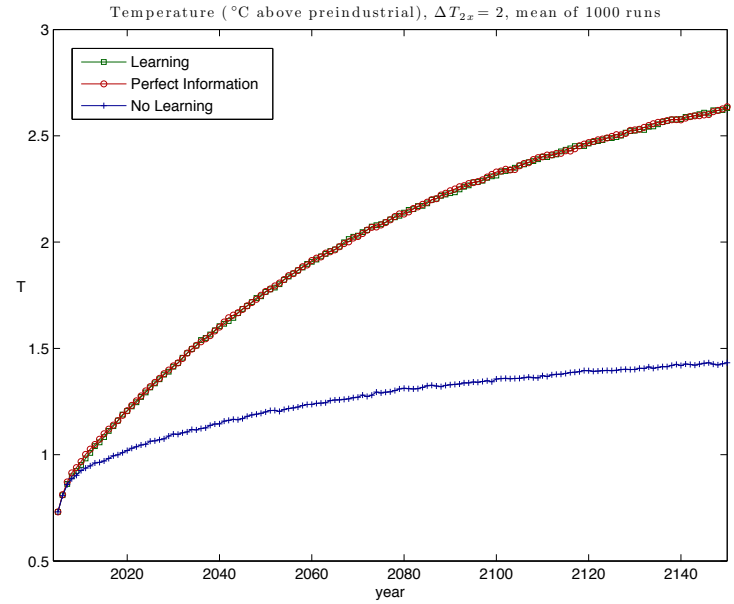

Figure D.13 shows the temperature, degrees C above preindustrial, with true $\Delta T_{2\times} = 2$ (mean of 1000 runs).

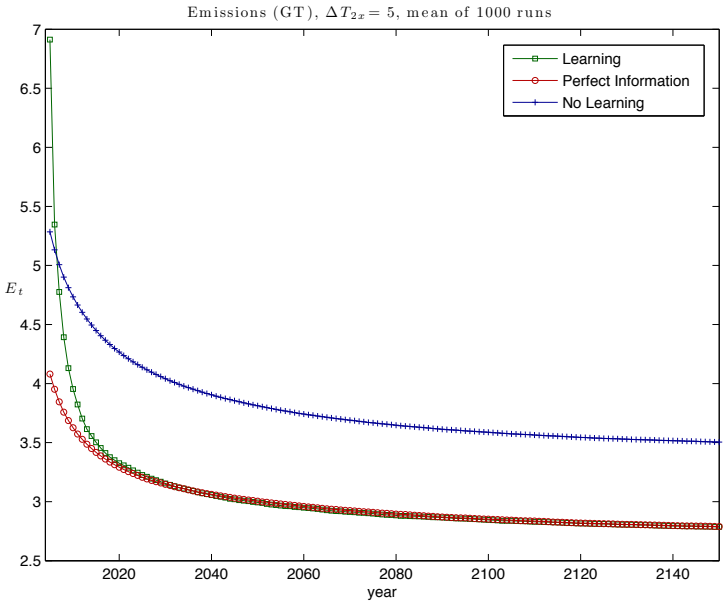

Figure D.14 shows the optimal emissions policy (GT), with true $\Delta T_{2\times} = 5$ (mean of 1000 runs).

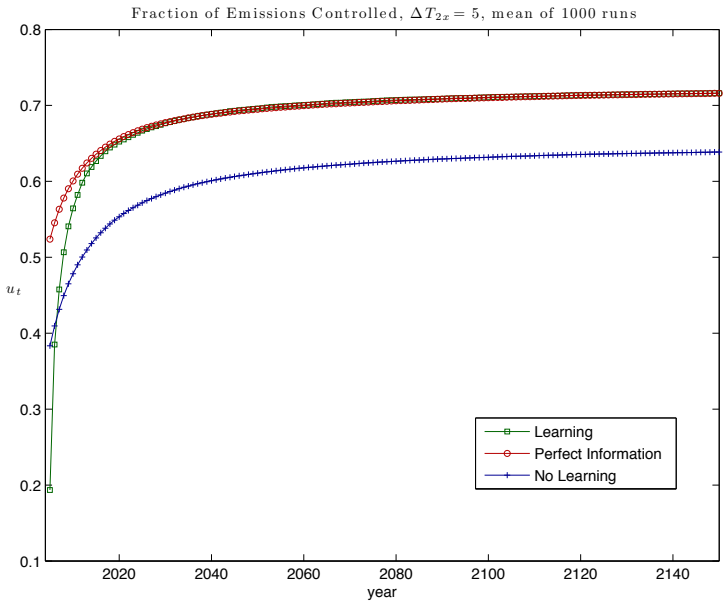

Figure D.15 shows the optimal emissions control rate, with true $\Delta T_{2\times} = 5$ (mean of 1000 runs).

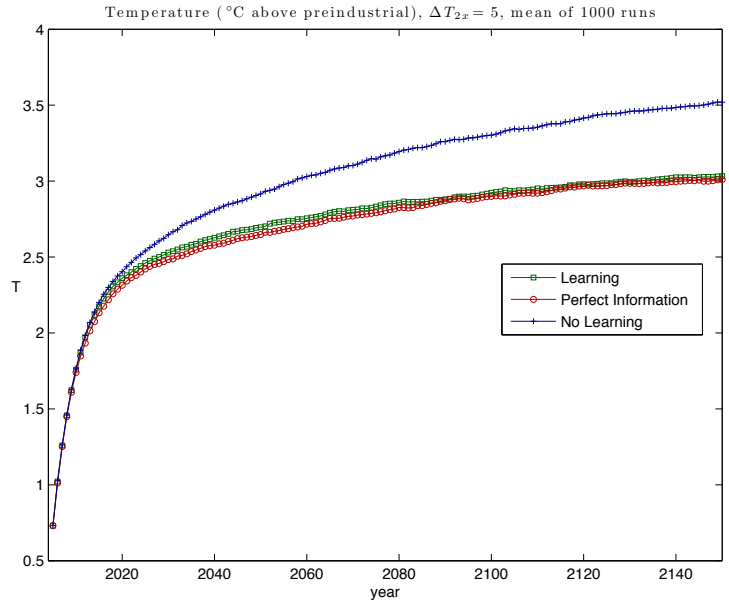

Figure D.16 shows the temperature, degrees C above preindustrial, with true $\Delta T_{2\times} = 5$ (mean of 1000 runs).

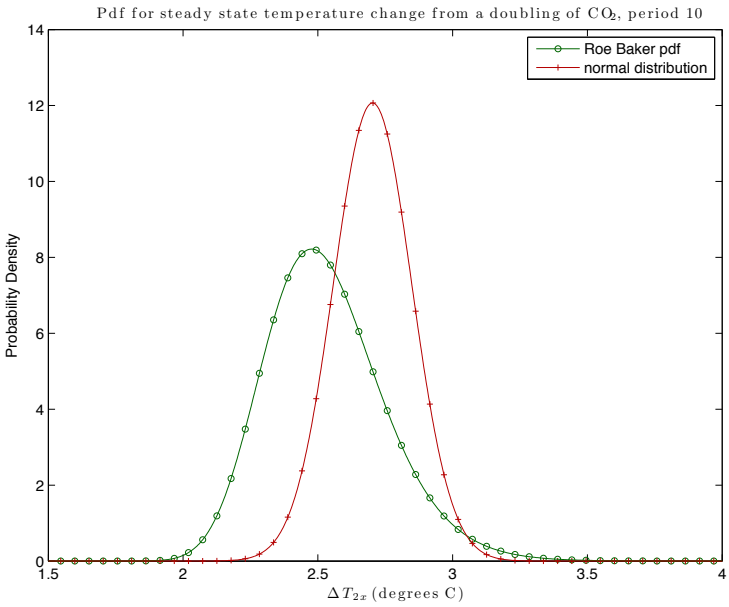

Figure D.17 shows the posterior PDF of $\Delta T_{2\times}$, with true value 2.76 after 10 periods.

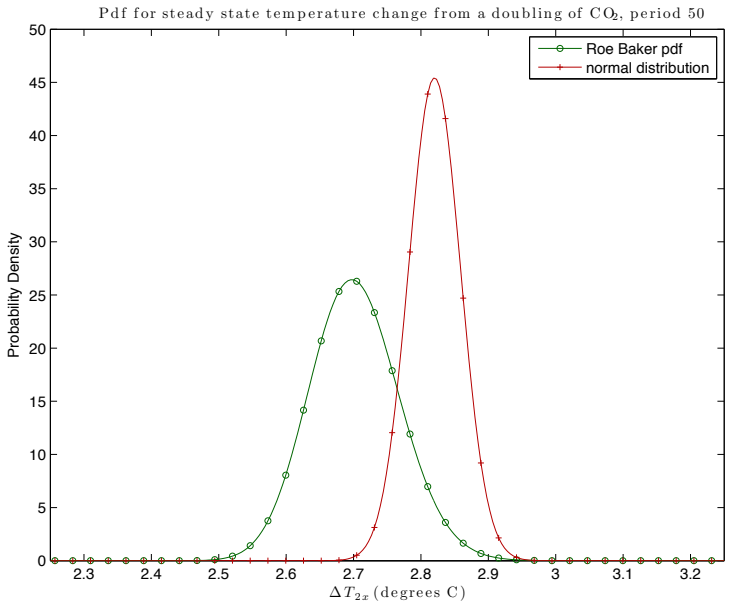

Figure D.18 shows the posterior PDF of $\Delta T_{2\times}$, with true value 2.76 after 50 periods.

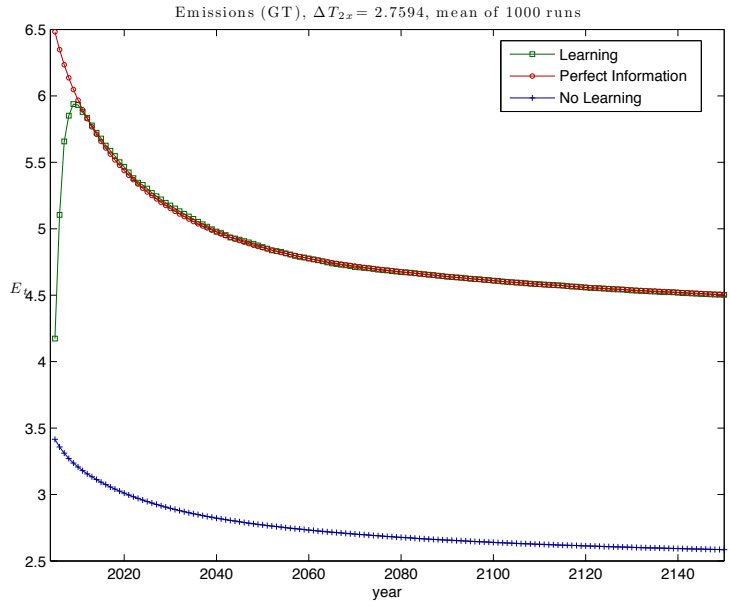

Figure D.19 shows the optimal emissions policy in gigatons (GT), with true $\Delta T_{2\times} = 2.76$, $\beta = 0.971$ (mean of 1000 runs).

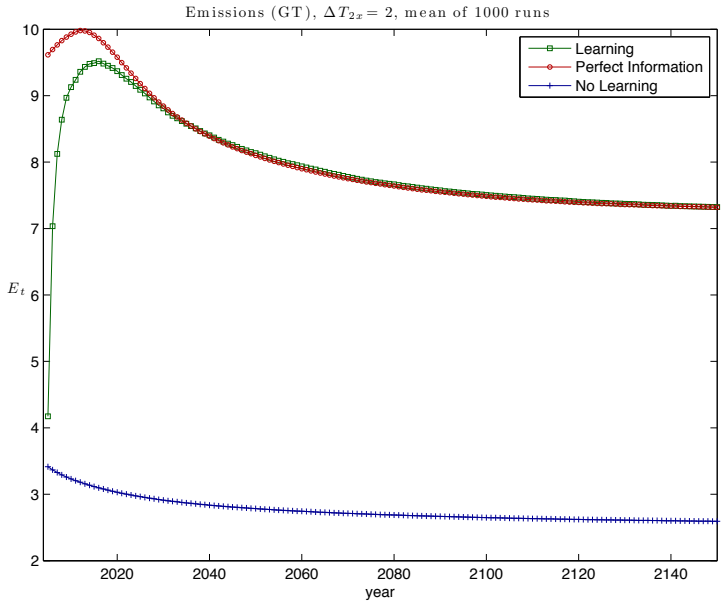

Figure D.20 shows the optimal emissions policy in gigatons (GT), with true $\Delta T_{2\times} = 2$, $\beta = 0.971$ (mean of 1000 runs).

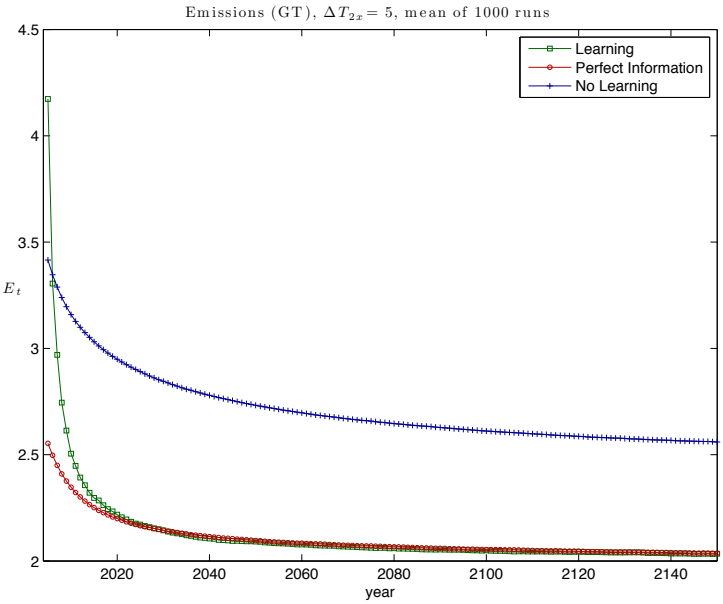

Figure D.21 shows the optimal emissions policy in gigatons (GT), with true $\Delta T_{2\times} = 5$, $\beta = 0.971$ (mean of 1000 runs).

APPENDIX E
Neutralization In China Figures

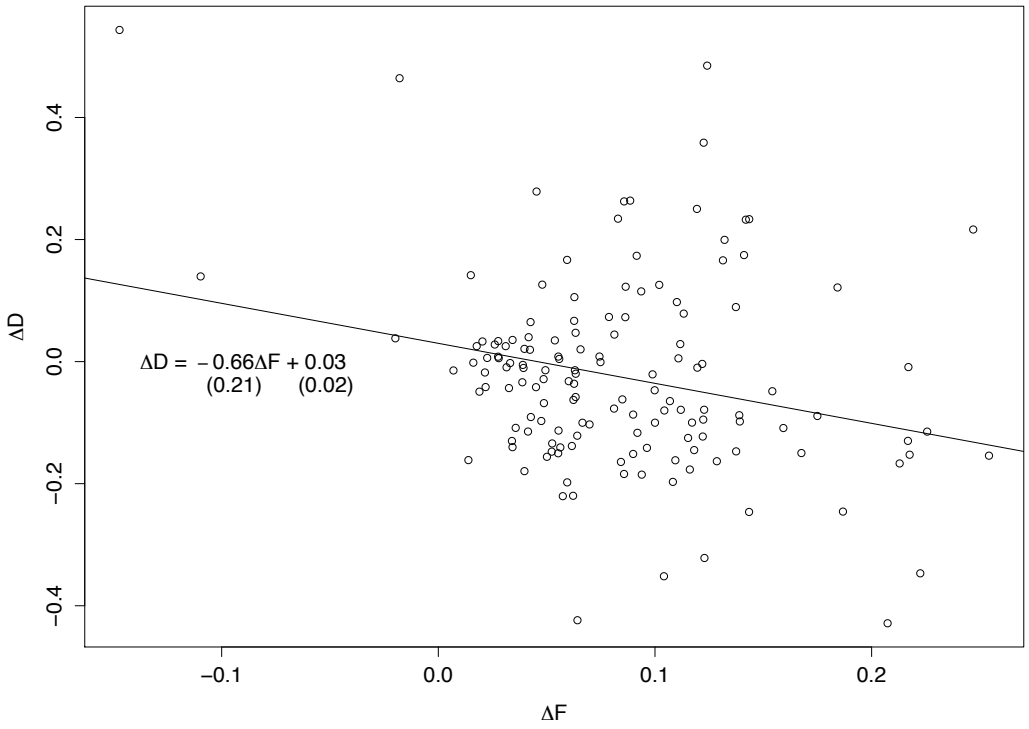

Figure E.1 shows the regression line of China's sterilization coefficient estimation.

APPENDIX F

China's Implicit Demand for Foreign Reserves Figures

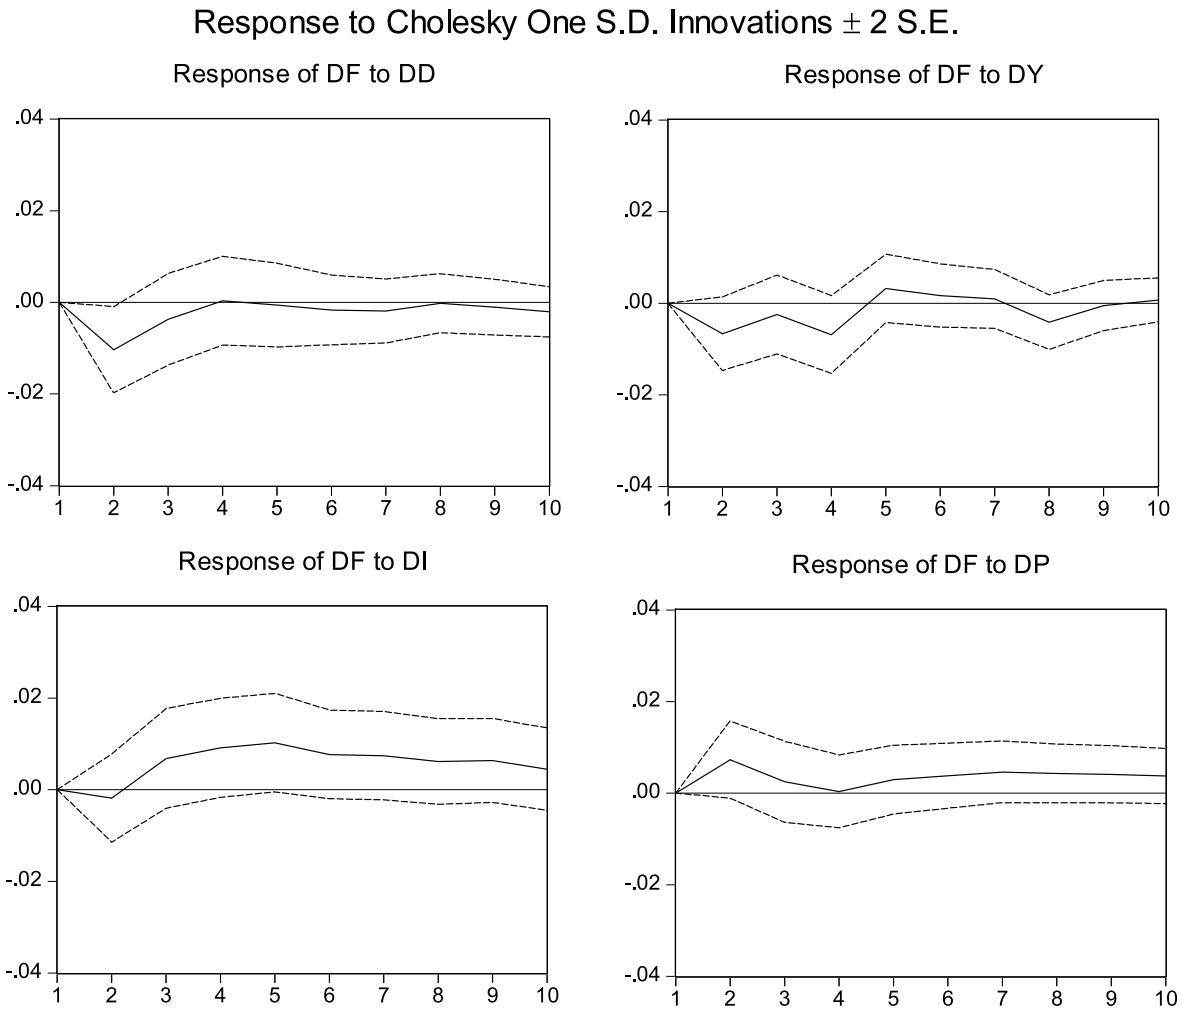

Figure F.1 shows the impulse responses of China's foreign reserve changes to changes of domestic credit, real income, interest rate and price level.