Update: A Targeted Wealth Approach to Risk and Return in Climate Change Policy

Progress Report for the Global Risk Institute

Margaret Insley* Peter A. Forsyth †

April 3, 2017

This report, which is part of an ongoing research project, is an update to a previous article entitled "A Targeted Wealth Approach to Risk and Return in Climate Change Policy". This article can be viewed HERE.

1 Introduction

Climate change is widely recognized as a serious risk that requires worldwide governmental action. In a recent speech to Lloyd's of London, Bank of England Governor, Mark Carney, drew attention to the risks of failing to act on climate change including the risk to the world's financial systems. To quote from Carney's speech,

> "The combination of the weight of scientific evidence and the dynamics of the financial system suggest that, in the fullness of time, climate change will threaten financial resilience and longer-term prosperity. While there is still time to act, the window of opportunity is finite and shrinking." [1]

Unfortunately, there is no consensus amongst policymakers and researchers on the most appropriate policy responses, including how quickly carbon emissions should be reduced. Reaching consensus is made more difficult by the enormous uncertainty inherent in projections regarding the timing and extent of the future warming of the earth's climate, the level of damages, and humanity's ability to adapt to a warmer climate.

*Department of Economics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1. *minsley@uwaterloo.ca*

†David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1. *paforsyt@uwaterloo.ca*

1From "Breaking the Tragedy of the Horizon - Climate Change and Financial Stability", speech given by Mark Carney, Governor of the Bank of England, Lloyd's of London, 29 September 2015."

Adding to this is the fact that the most serious damages are expected to occur many decades into the future. Using typical discount rates, traditional economic analysis would recommend little action be taken at present. Further, world climate is affected by the cumulative stock of atmospheric carbon, so that any actions taken today to reduce carbon emissions will only have an impact after a considerable time lag. Policies to reduce carbon emissions, whether via a carbon tax, cap and trade, or other regulatory schemes, are costly as they entail a realignment of the economy away from carbon based fuels. Amongst governments and other decision makers, there is a natural tension between the desire to delay the implementation of these costly policies and the concern that a more precautionary approach is warranted in the face of these huge uncertainties and time lags.

1.1 Current Approaches

A significant portion of current research and policy discussions concerning climate change are based on analysis from Integrated Assessment Models (IAMs), which integrate climate science and economic models to assess the impact of a warming climate.[2] Of necessity, these models include a large number of parameters and functional relationships about which we have little information. As is pointed out by Pindyck (2013), different economic analyses of climate change can come to starkly different conclusions, depending on key assumptions about such factors as the levels of future greenhouse gas emissions under business-as-usual versus alternate scenarios, projections of average global temperatures and economic impacts, estimates of the costs of abating greenhouse gas emissions, and the rate of discount. This is exemplified by two very influential studies based on IAM's as described in Stern (2008) and Nordhaus (2008). The former argued that immediate action to reduce global carbon emissions would produce significant positive economic returns, while the latter recommended only limited reduction in carbon emissions. The Stern study was criticized for using a very low discount rate, which Stern justified on the ethical grounds that society should give equal weight to the well being of current and future generations.

A key component in any economic analysis of climate change is valuing the loss in societal welfare that results from reduced economic output with a warming climate. Societal welfare is typically assumed to depend on consumption levels via a Constant Relative Risk Aversion

[2] See Nordhaus (2008), Hope (2006) and Tol (2001*a*) and Tol (2001*b*) for a description of three influential IAMS.

(CRRA) social utility function which can be specified as follows:

$$U(C_t) = \frac{C_t^{1-\eta}}{(1-\eta)},$$

where U denotes utility which depends on consumption in period t, C_t, and η is the index of relative risk aversion. The impact of climate change on societal welfare is measured by the discounted sum of annual utilities over a long time period (such as 500 years) in different climate scenarios. The discount rate and the coefficient of relative risk aversion are critical to the results, and there is no general agreement on the correct values.

More recently, research has focused on the treatment of uncertainty in scientific estimates of the response of world climate to the increasing atmospheric concentration of greenhouse gases. Most IAMs are deterministic, with uncertainty analyzed by running Monte Carlo simulations on key model parameters (Pindyck 2013). A criticism of much current work is the failure to deal adequately with possible catastrophic outcomes whereby average global temperatures increase by 4 to 5 degrees Celsius by the turn of the next century. The probability of such an event is non-zero, and its impact is unknown. The damage functions used in IAM's are calibrated for small temperature changes and are uninformative about the impact of such an extreme event. Weitzman (2014) has pointed out that if the probability distribution of the negative consequences of climate change has a sufficiently fat tail, and if we adopt a sufficiently high degree of relative risk aversion, the optimal policy for society would be highly precautionary.

Some of the earlier papers to incorporate uncertainty into optimal decisions for climate change were Pizer (1999) and Kelly & Kolstad (1999, 2001). More recent work adapts the Nordhaus IAMs (DICE-2007) introducing uncertainty into models whereby the objective is to maximize the sum of discounted utility. Traeger (2014) and Lemoine & Traeger (2014) are examples which use a stochastic dynamic programming solution approach. The current project draws on this recent work.

1.2 An Alternative: A Targeted Wealth Approach

Society's choices about spending today's national income will determine the level of capital stock and income generating capacity available for future generations. In any given year, a portion of national income is spent on consumption goods and a portion is invested to augment the capital stock. However the world's ability to generate income is harmed by increasing global temperatures. National governments can choose to ameliorate the situation by devoting a portion of national income to reducing worldwide carbon emissions, but at the cost of reducing the shares available for current consumption and investment in the capital stock. Governments face the dilemma of balancing the desire for current consumption with the desire to maintain the economy's income generating capacity through investment in capital stock and reducing the anthropomorphic causes of climate change.

We adopt an alternative approach to this dilemma, utilizing concepts from the pension industry for dealing with risky decisions under extreme uncertainty. Rather than choosing policies to maximize a sum of discounted social utility, we argue that from the point of view of intergenerational fairness, we should choose climate change policies which enable future generations to enjoy a reasonable standard of living in the distant future by maintaining a minimum level of capital stock per capita. We call our approach the *targeted wealth approach* as it involves setting a target level for global wealth as measured by the world's capital stock, and choosing policies to minimize the risk of falling short of the target.

While this idea is not in conflict with the social utility function approach, we deal directly with the desired end state, rather than indirectly through the use of a utility function and a discount rate. Using this new approach, the trade off between the risk and reward of devoting economic resources towards slowing climate change is very clear and intuitive. In addition we avoid the need to make an arbitrary choice for the discount rate. We consider this approach to be a complement to the current approaches used in IAMs models.

In the following sections we develop a model of the interaction of climate and the economy based on standard assumptions in the DICE model as described in Nordhaus (2008) and in many papers which address the economics of climate change such as Leach (2007), Traeger (2014), Lemoine & Traeger (2014). Uncertainty in climate change is captured by modelling temperature as a mean reverting stochastic process. We formulate a stochastic optimal control problem whereby the social planner chooses a path for the abatement of carbon emissions to minimize the risk of failing to achieve a targeted capital stock per capita T years in the future. It is demonstrated in the literature that the optimal abatement strategy to achieve this target is also dynamically mean-variance optimal (Li & Ng 2000, Zhou & Li 2000, Dang & Forsyth 2016), meaning that no other strategy has a smaller variance. This relates to a familiar concept in finance: the mean variance efficient frontier. Higher future expected capital stock per capita comes at the cost of greater risk. More certainty about future outcomes is only possible at the expense of lower capital stock. The choice of where to be on this frontier is a policy decision. However, given any specified future mean capital stock per capita, we can determine the emission strategy which results in this mean value, with the smallest risk.

1.3 Modelling assumptions

The mathematical details of the model are provided in the next section, including further discussion of the mean-variance efficient frontier, but as an overview we summarize some basic modelling assumptions.

- Global climate is affected by the mean temperature, which is determined, in the long-run, by the cumulative stock of global atmospheric carbon.
- The global stock of carbon, if no anthropocentric carbon is emitted, decays slowly to a natural background mean level.
- Uncertainties in the response of the climate to the increasing atmospheric stock of carbon are captured by modelling global average temperature as a mean reverting stochastic process, which mean-reverts to a level determined by the global stock of atmospheric carbon.
- World income (GDP) depends positively on the stocks of labour and capital, but is reduced by damages due to rising temperatures and also by spending on carbon abatement.
- The labour supply grows deterministically to a given upper limit. We assume the labour supply is a constant proportion of world population.
- National income is either consumed, invested, or devoted to carbon abatement policies. The amount devoted to carbon abatement is determined as the optimal control. Consumption per capita is assumed to grow at an exogenously specified rate, intended to reflect expectations for acceptable growth in living standards. The remainder of national income is invested in the capital stock.

2 Detailed specification of the decision model

2.1 Modelling the interaction of climate and the economy

The climate-economy model presented here is based on Traeger (2014) who introduced a stochastic dynamic programming version of the DICE-2007 model and Lemoine & Traeger

(2014) who used a similar model to investigate tipping points in the strength of the carbon-temperature feedback and the earth's ability to remove carbon from the atmosphere. A key difference in our analysis is that temperature is modelled as an Ito process.

2.1.1 Temperature and carbon stock

Average global surface temperature, X, is described by a stochastic differential equation as follows:

$$dX = \phi_1 \left[\bar{X}(t) - X(t)\right] dt + \sigma X(t) dz, \tag{1}$$

where $X(t)$ refers to the increase in average surface temperature since pre-industrial times, ϕ_1 and σ are constant parameters. dz is the increment of a standard Weiner process and is intended to capture volatility in the earth's temperature. Global average temperature is assumed to be mean reverting to a long run value $\bar{X}$, which is specified as

$$\bar{X} = \phi_2 \left(\frac{log(S(t)/\bar{S})}{log(2)} \right) + \phi_3. \tag{2}$$

$S(t)$ is the stock of atmospheric carbon. $\bar{S}$ is the pre-industrial atmospheric carbon stock. ϕ_2, the climate sensitivity parameter, can be interpreted as the impact on global average temperature of a doubling of carbon stock. ϕ_3 gives the impact on temperature of green-house gases other than carbon. The mean reverting structure of Equation (1) is common in Integrated Assessment Models such as RICE (Nordhaus & Boyer 2000). Variations emerge in the specification of ϕ_1, ϕ_2, and ϕ_3. In Lemoine & Traeger (2014) these parameters are deterministic functions of temperature and time.

Carbon emissions from burning fossil fuels enter the global carbon cycle whereby carbon is exchanged between various carbon reservoirs including the atmosphere, oceans and terrestrial ecosystems. An important modelling assumption is how carbon emissions affect the atmospheric stock of carbon, which is the key driver in climate change. The evolution of the stock of atmospheric carbon, $S(t)$ over time is assumed to be given by:

$$dS(t) = [\epsilon(t)\delta_{\epsilon,t} - (S(t) - \bar{S})\delta_{S,t}]dt. \tag{3}$$

Here $\epsilon(t)$ refers to current emissions. $\delta_{\epsilon,t}$ refers to the portion of emissions that becomes part of atmospheric carbon stock. $\bar{S}$ is the pre-industrial accumulation of carbon. $\delta_{S,t}$ is the rate at which carbon is removed from the atmospheric stock by natural processes. The literature varies in its specification of $\delta_{S,t}$. It is sometimes specified as a deterministic function of time, or alternately as a function of time and the stock of carbon. In this analysis we adopt the specification of Golosov et al. (2014).

2.1.2 Economic output and carbon emissions

Gross economic output (Y^{gross}) is specified by a Cobb-Douglas production function:

$$Y^{\text{gross}}(t) = (A(t)L(t))^{(1-\xi)}K(t)^{\xi} \tag{4}$$

where $A(t)$ is a parameter capturing technological change, $L(t)$ denotes units of labour, $K(t)$ denotes units of capital, and ξ is a constant parameter. $A(t)$ and $L(t)$ are both deterministic functions of time, as is detailed in Section 3.3. In each period, gross output is reduced by damages from climate change and by the cost of resources devoted to reducing carbon emissions.

The damages caused by rising global temperature are captured by a damage function which specifies the fractional reduction in annual output associated with higher levels of X.

$$D(t) = \kappa_1 X(t)^{\kappa_2}. \tag{5}$$

where κ_1 and κ_2 are constant parameters. This form of damage function is common in the literature (Nordhaus (2008), Leach (2007), Crost & Traeger (2013), Golosov et al. (2014), Lemoine & Traeger (2014)).

Remark 1 (Criticism of equation (5))**.** *Some researchers have argued that equation (5) is only valid for small changes in global temperature, and does not model the disastorous effects of large temperature effects. Since we consider the temperature change to have deterministic and random components, it will be necessary to modify equation (5) for large (i.e.* $> 5°$*) temperature changes. We expect that a suitable modification will permit modelling the insurance effect of early emission reductions.*

Abatement cost as a fraction of gross GDP is given by:

$$\Lambda(\mu(t)) = \psi(t)\mu(t)^{a_2}. \tag{6}$$

$\mu(t)$ is the fractional reduction in emissions compared to the business-as-usual case, and is a state variable, as discussed in Section 3.5. a_2 is a constant parameter, and $\psi(t)$ is assumed to fall deterministically over time as is detailed in Section 3.3.

Net GDP can then be defined as

$$
\begin{aligned}
Y(t) &= \frac{1-\Lambda(\mu(t))}{1+D(X)} Y^{\text{gross}}(t) \\
&= \frac{1-\Lambda(\mu(t))}{1+D(X)} (A(t)L(t))^{(1-\xi)} K(t)^{\xi}
\end{aligned} \tag{7}
$$

Production results in carbon emissions. Total emissions at time t is the sum of emissions from production, $Y(t)$, as well as emissions from other sources such as land use changes, $B(t)$.

$$
\epsilon(t) = \eta(t)(1-\mu(t))Y(t) + B(t) \tag{8}
$$

with $\eta(t)$ and B_t deterministic functions of time as given in Section 3.3.

2.2 Consumption and the evolution of the capital stock

Consumption per capita is assumed to grow at rate δ_c, which is intended to reflect a minimum desired rate of increase in the standard of living. Consumption is also constrained to be no larger than output in any year, so that the capital stock cannot be consumed. Denote aggregate consumption as $C(t)$ and consumption per capita as $c(t) \equiv C(t)/L(t)$, then

$$
C(t) = \min\left[L(t)c_0 e^{\delta_c}, Y(t)\right] \tag{9}
$$

$c_0 \equiv C_0/L_0$ is consumption per capita at $t = 0$.

The capital stock evolves according to:

$$
dK = \Big(Y(t) - C(t) - \delta_K K\Big)dt \tag{10}
$$

where δ_K reflects the depreciation of the capital stock.

Since our objective function will be specified in terms of capital per person, we define the variable $\Omega(t) \equiv K(t)/L(t)$.

2.3 Evolution of deterministic variables

Following Lemoine & Traeger (2014) there are several variables which are assumed to evolve as deterministic functions of time.

- Labour supply, $L(t)$, in Equation (4):

$$dL = \delta_L(L_\infty - L(t))dt \tag{11}$$

where L_∞ is a long term equilibrium value for the labour supply.

- Technology, $A(t)$ in Equation (4):

$$dA = g_{A,t}A(t)dt \tag{12}$$

where the growth rate in technology follows:

$$dg_{At} = \delta_A g_{A,t}dt \tag{13}$$

- Emissions intensity of production, in Equation (8):

$$d\eta(t) = \eta(t)g_{n,t}dt \tag{14}$$

where

$$dg_{n,t} = -\delta_\eta g_{\eta,t}dt \tag{15}$$

- Nonindustrial carbon emissions, in Equation (8):

$$B(t) = B_0 e^{tg_B} \tag{16}$$

- Abatement cost in Equation (7):

$$\psi(t) = \frac{a_0\eta(t)}{a_2}(1 - \frac{(1 - e^{tg_\phi})}{a_1}) \tag{17}$$

2.4 Objective function

Our objective is specified in terms of capital stock per person, $\Omega(t) \equiv K(t)/L(t)$, and we denote the desired level of expected future capital stock per capita as Ω_{des}. We aim to achieve Ω_{des} by specifying the fractional reduction in annual carbon emissions compared to a "business as usual" (BAU) case. This fractional reduction, denoted $\hat{\mu}(t)$ is the control variable in the decision model. Mathematically our objective is to determine $\hat{\mu}(t)$ to minimize the risk that we fall below a particular target, denoted Ω^*, at time T in the future. In addition our objective function must be specified so that there is no penalty for exceeding the target. This is done by defining the objective function in terms of $\bar{\Omega}$, referred to as "effective capital" where $\bar{\Omega} = \min(\Omega_T, \Omega^*)$. The objective function is then given as:

$$
\begin{aligned}
&\min_{\hat{\mu}(t)} E\left[\left(\bar{\Omega} - \Omega^* \right)^2 \right]; \\
&\text{subject to} \quad \bar{\Omega} = \min(\Omega_T, \Omega^*) \quad \text{and} \quad E[\bar{\Omega}] = \Omega_{des} \\
&\qquad E[\cdot] = \text{ Expectation} \\
&\qquad \hat{\mu}(t) = \text{ Choice of fractional reduction in emissions compared to BAU case} \\
&\qquad \Omega^* = \text{ Target capital stock per capita at time } T \\
&\qquad \Omega_T = \text{ Observed capital stock per capita at time } T\ .
\end{aligned} \tag{18}
$$

Note that according to Equation (18), a large shortfall is penalized more than a small shortfall. It turns out that in order to achieve $E[\bar{\Omega}_T] = \Omega_{des}$, we must aim for a target, Ω^*, which is slightly higher than Ω_{des}.[3] In Equation (18) the expected value of $\bar{\Omega}$ at time T will be less than Ω^*, ie, $E[\bar{\Omega}_T] < \Omega^*$. We therefore choose Ω^* so that $E[\bar{\Omega}_T] = \Omega_{des}$. This idea is intuitive: as noted above in order to achieve Ω_{des} on average, our target capital stock must exceed our desired capital stock.

2.5 State and control variables

Carbon abatement is achieved through policies such as a carbon tax or regulation, which in reality can only be adjusted at discrete time points. In our model, the optimal control is chosen at discrete decision times denoted by:

$$\mathcal{T} = \{t_0 = 0 < t_1 < ..., t_i, ... < t_I = T\} \tag{19}$$

The state of the system is determined by four variables: the stock of carbon, $S(t)$, the temperature, $X(t)$, the capital stock per person, $\Omega(t)$, and abatement $\mu(t)$. All other variables are deterministic functions of the state variables. We define $P_t = (S(t), X(t), \Omega(t), \mu(t))$ as the process specifying the state of the system and $p = (s, x, \omega, \mu)$ as the realization of the state of the system at time t for $t \in \mathcal{T}$.

[3] "We aim above the mark to hit the mark." Ralph Waldo Emerson

Let t^- and t^+ respectively denote the instants in time just before and just after time, t. We use $\hat{\mu}(\cdot)$ to denote the control as a function of the current state at t^- for $t \in \mathcal{T}$, i.e.

$$\hat{\mu}(\cdot) : (P^{t-}, t^-) \rightarrow \hat{\mu} = \hat{\mu}(P^{t-}, t^-), \ t \in \mathcal{T} \tag{20}$$

The control $\hat{\mu}(\cdot)$ represents the abatement choice in each period. Note that abatement is both a state variable ($\mu(t)$), representing the current abatement level, and a control variable representing the optimal choice in the next instant, $\hat{\mu}(\cdot)$, given the current value of state variables. To allow for constraints on the control, we specify a set of admissible controls, $\mathcal{Z}$ and $\hat{\mu} \in \mathcal{Z}$.

The state prevailing at time t^-, $t \in \mathcal{T}$, is given as:

$$p = (s, x, \omega, \mu) = (S(t^-), X(t^-), \Omega(t^-), \mu(t^-)). \tag{21}$$

At time t^+ immediately after the imposition of the control, the state will be given by

$$\begin{aligned} &(S^+, X^+, \Omega^+, \mu^+) \equiv \\ &\quad (S^+(s, x, \omega, \mu; \hat{\mu}(\cdot)), X^+(s, x, \omega, \mu; \hat{\mu}(\cdot)), \Omega^+(s, x, \omega, \mu; \hat{\mu}(\cdot)), \mu^+(s, x, \omega, \mu; \hat{\mu}(\cdot))). \end{aligned} \tag{22}$$

The instant after the control is applied, a new abatement level, $\hat{\mu}$, may be chosen; while S, X, and Ω are unchanged:

$$S^+ = s, \quad X^+ = x, \quad \Omega^+ = \omega, \quad \mu^+ = \hat{\mu} \tag{23}$$

2.6 Mean-Variance Efficient Frontiers

This section details how the mean-variance frontier is determined, and is based on Forsyth & Vetzal (forthcoming). The achievable M-V objective set is defined solely in terms of $\bar{\Omega} = \min(\Omega_T, \Omega^*)$, as specified in Equation (18) . We denote by $E_{\hat{\mu}}^{p,t}(\cdot)[\bar{\Omega}]$ and $Var_{\hat{\mu}(\cdot)}^{p,t}[\bar{\Omega}]$, respectively, the expectation and the variance of effective capital, conditional on the state (p, t) and the control $\hat{\mu}(\cdot)$, $t \in \mathcal{T}$. Let $(p_0, 0) \equiv (P(t = 0), t = 0)$ denote the initial state. The achievable M-V objective is:

$$\mathcal{Y} = \{(Var_{\hat{\mu}(\cdot)}^{p_0,0}[\bar{\Omega}], E_{\hat{\mu}(\cdot)}^{p_0,0}[(\bar{\Omega}]) \ : \ \hat{\mu} \in \mathcal{Z}\}. \tag{24}$$

For each point $(\mathcal{V},\mathcal{E}) \in \mathcal{Y}$ and for an arbitrary scalar $\rho > 0$, the set of points $\mathcal{Y}_{P(\rho)}$ is defined to be:

$$\mathcal{Y}_{P(\rho)} = \{(\mathcal{V}_*,\mathcal{E}_*) \in \mathcal{Y} : \rho\mathcal{V}_* - \mathcal{E}_* = \min_{(\mathcal{V},\mathcal{E})\in\mathcal{Y}} (\rho\mathcal{V} - \mathcal{E})\}. \tag{25}$$

ρ is a parameter which governs the decision maker's preferences regarding the trade off between expected Y and its variance. For a given ρ, $\mathcal{Y}_P(\rho)$ is an efficient point in that it offers the highest expected $\bar{\Omega}$ for a given variance. The efficient frontier, $\mathcal{Y}_P$, maps out all the efficient points for all values of ρ.

The variance term in Equation (25) implies that we cannot determine the associated value function using dynamic programming.[4] Using the embedding result of Li & Ng (2000) and Zhou & Li (2000), we define the value function $V(p,t)$:

$$V(p,t) = \min_{\hat{\mu}\in Z}\{E^{p,t}_{\hat{\mu}}\left[\left(\bar{\Omega} - \Omega^*\right)^2\right]\} \tag{26}$$

where the parameter $\Omega^* \in (-\infty,+\infty)$. As discussed in Forsyth & Vetzal (forthcoming), the embedding result implies that there exists a $\Omega^* \equiv \Omega^*(p,t,\rho)$ such that for a given positive ρ, a control $\hat{\mu}$ which minimizes Equation (25) also minimizes Equation (26). Unlike Equation (25), dynamic programming can be implemented to determine the optimal control. The value function $V(p,t)$ can be found by solving the associated Hamilton-Jacobi-Bellman (HJB) equation backward in time with the terminal condition:

$$V(p,T) = [\min(\omega,\Omega^*) - \Omega^*]^2 \tag{27}$$

Equation (26) is solved for a given value of Ω^*. Once the optimal control is known it is straightforward to determine the expected value of wealth (capital stock per person) under the optimal control $E^{p,t}_{\hat{\mu}}[\bar{\Omega}_T]$, and hence a point on the efficient frontier, $\left(Var^{p_0,t=0}_{\hat{\mu}}[\bar{\Omega}_T], E^{p_0,t=0}_{\hat{\mu}}[\bar{\Omega}_T]\right)$. This is repeated for numerous values of Ω^* to trace on a curve in the $(\mathcal{V},\mathcal{E})$ plane.

The expected value of the effective capital stock per person, $E[\bar{\Omega}_T]$ can be found by solving an additional HJB equation for $U(p,t) = E[\bar{\Omega}_T]$ with the terminal condition $U(p,t) = \min(\omega,\Omega^*)$.

[4] This is discussed in Basak & Chabakauri (2010), Li & Ng (2000) and Zhou & Li (2000).

2.7 The HJB Equation and Optimal Control

Using dynamic programming, the problem is solved backwards in time. For times $t \notin \mathcal{T}$, V is found from the solution of an HJB equation, which follows by applying Ito's lemma to Equation (26) and taking expectations. Note that there is no term associated with the discount rate, since the objective function is posed in terms of the real terminal capital stock per capita. Define the diffusion operator as:

$$\begin{aligned} \mathcal{L}V \quad &\equiv \quad \frac{(\sigma x)^2}{2}\frac{\partial^2 V}{\partial x^2} + \phi_1(\bar{X} - x)\frac{\partial V}{\partial x} + \left((\bar{S} - s)\delta_{S,t} + \epsilon\delta_{\epsilon,t}\right)\frac{\partial V}{\partial s} + \alpha_\omega \frac{\partial V}{\partial \omega}. \\ & \alpha_\omega = \left[Y - C - \delta_L(L_\infty - L)\right] \,. \end{aligned} \tag{28}$$

We remind the reader that $Y(\cdot)$ is defined in equation (7) and $C(t)$ is defined in equation (9). To reduce notational complexity we have suppressed the dependence of state variables on time as well as the dependence of V on the state variable in vector p. Then the final HJB equation is

$$\frac{\partial V}{\partial t} + \mathcal{L}V = 0 \ , t \notin \mathcal{T} \ . \tag{29}$$

At discrete decision dates, $t = t_i \in \mathcal{T}$, define the intervention operator as:

$$\mathcal{M}(\mu)V(p,t) = V(S^+, X^+, \Omega^+, \mu^+, t). \tag{30}$$

The intervention operator applies the optimal control by choosing the level of abatement, $\mu^+ = \hat{\mu}$ as specified in Equation (23).

At each time $t = t_i \in \mathcal{T}$, the optimal control is chosen to minimize the objective function by applying the intervention operator,

$$V(p, t_i^-) = \min_{\hat{\mu} \in \mathcal{Z}} \mathcal{M}(\mu)V(p, t_i^+). \tag{31}$$

At the terminal time $t = T$ we have

$$V(p,T) = [\min(\omega, \Omega^*) - \Omega^*]^2. \tag{32}$$

In order to determine the mean and variance of $\bar{\Omega}_T$, we need to solve an additional equation, for $E[\bar{\Omega}_T]$ Given a set of optimal controls, $\hat{\mu}^* = \hat{\mu}(p,t)$, $t \in \mathcal{T}$, we use these controls to solve for:

$$U(p,t) = E^{p,t}_{\bar{\Omega}^*}[\bar{\Omega}_T], \tag{33}$$

by solving the PDE

$$U_t + \mathcal{L}U = 0, \quad t \notin \mathcal{T}, \tag{34}$$

and the optimal conditions at $t \in \mathcal{T}$

$$\mathcal{M}(\hat{\mu}^*)U(p,t) = U(S^+, S^+, \Omega^+, \mu^+), \tag{35}$$

with the terminal condition

$$U(p,T) = \min(\omega, \Omega^*) \quad . \tag{36}$$

Defining

$$p_0 = (S_0, X_0, \Omega_0, \mu_0) \tag{37}$$

then, the mean and variance as seen at $t = 0$ are given by:

$$\begin{aligned} V(p_0, 0) &= E^{p_0,t}_{\hat{\mu}(\cdot)}[(\bar{\Omega} - \Omega^*)^2] \\ U(p_0, 0) &= E^{p_0,t}_{\hat{\mu}(\cdot)}[\bar{\Omega}] \\ E^{p_0,t}_{\hat{\mu}(\cdot)}[\bar{\Omega}^2] &= V(p_0, 0) + 2\Omega^*\, U(p_0, 0) - (\Omega^*)^2. \end{aligned} \tag{38}$$

Therefore,

$$Var^{p_0,t}_{\hat{\mu}(\cdot)}[\bar{\Omega}] = E^{p_0,t}_{\hat{\mu}(\cdot)}[(\bar{\Omega})^2] - (U(p_0, 0))^2 \ . \tag{39}$$

2.8 Numerical solution of the HJB equation

The HJB equation (29) is a highly nonlinear PDE, and closed form solutions are elusive. In order to avoid making unrealistic simplifications, we have developed a numerical algorithm for solving equation (29). This equation is a four dimensional problem, and can be expected to be CPU intensive. The method we will use is based on an implicit finite difference technique, combined with semi-Lagrangian timestepping (Chen & Forsyth 2007).

3 Data Gathering

The data for the parameters in this model will be largely from the DICE-2007 model described in Nordhaus (2008). Nordhaus based his analysis on the most recently available reports from the Intergovernmental Panel on Climate Change. We will update Nordhaus's assumptions using more recent reports.

References

Basak, S. & Chabakauri, G. (2010), 'Dynamic mean variance asset allocation', *Review of Financial Studies* **23**, 2970–3016.

Chen, Z. & Forsyth, P. (2007), 'A semi-lagrangian approach for natural gas storage valuation and optimal operation', *SIAM Journal on Scientific Computing* **30**, 339–368.

Crost, B. & Traeger, C. P. (2013), 'Optimal climate policy: Uncertainty versus monte carlo', *Economics Letters* **120**(3), 552–558.

Dang, D. M. & Forsyth, P. A. (2016), 'Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton Jacobi Bellman approach', *European Journal Operational Research* **250**, 827–841.

Forsyth, P. & Vetzal, K. (forthcoming), 'Long term asset allocation for the patient investor', *International Journal of Theoretical and Applied Finance* .

Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2014), 'Optimal taxes on fossil fuel in general equilibrium', *Econometrica* **82**(1), 41–88.

Hope, C. (2006), 'The marginal impact of CO2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern', *Integrated Assessment* **10**(1), 19–56.

Kelly, D. & Kolstad, C. (1999), 'Bayesian learning, growth, and pollution', *Journal of Economic Dynamics & Control* **23**(4), 491–518.

Kelly, D. & Kolstad, C. (2001), 'Solving growth models with with an environmental sector', *Journal of Computational Economics* **18**.

Leach, A. (2007), 'The climate change learning curve', *Journal of Economic Dynamics and Control* **31**, 1728–1752.

Lemoine, D. & Traeger, C. (2014), 'Watch your step: Optimal policy in a tipping climate', *American Economic Journal: Economic Policy* **6**(2), 137–166.

Li, D. & Ng, W.-L. (2000), 'Optimal dynamic portfolio selection: Multiperiod mean-variance formulation', *Mathematical Finance* **10**, 387–406.

Nordhaus, W. (2008), *A question of balance: Weighing the options on global warming policies*, Yale University Press.

Nordhaus, W. D. & Boyer, J. (2000), *Warming the World: Economic Modeling of Global Warming*, MIT Press.

Pindyck, R. S. (2013), 'Climate change policy: What do the models tell us?', *Journal of Economic Literature* **51**, 860–872.

Pizer, W. (1999), 'The optimal choice of climate change policy in the presence of uncertainty', *Resource and Energy Economics* **21**, 255–287.

Stern, N. (2008), 'The economics of climate change', *American Economic Review* **98**(2), 1–37.

Tol, R. (2001*a*), 'Estimates of the damage costs of climate change, part i: Benchmark estimates', *Environmental and Resource Economics* **21**(1), 47–73.

Tol, R. (2001*b*), 'Estimates of the damage costs of climate change, part ii: Dynamic estimates', *Environmental and Resource Economics* **21**(2), 135–60.

Traeger, C. (2014), 'A 4-stated dice: Quantitatively addressing uncertainty effects in climate change', *Environmental and Resource Economics* **59**(2), 1–37.

Weitzman, M. (2014), 'Fat tails and the social cost of carbon', *American Economic Review: Papers & Proceedings* **104**(5), 544–546.

Zhou, X. Y. & Li, D. (2000), 'Continuous-time mean-variance portfolio selection: A stochastic LQ framework', *Applied Mathematics and Optimization* **42**, 19–33.