

Strategic interactions and uncertainty in decisions to curb greenhouse gas emissions *

Margaret Insley[†] Tracy Snoddon[‡] Peter A. Forsyth[§]

This version: June 1, 2018

Working Paper

Abstract

This paper examines the strategic interactions of two large regions making choices about greenhouse gas emissions in the face of rising global temperatures. Optimal decisions are modelled in a fully dynamic, closed loop Stackelberg pollution game. Global average temperature is modelled as a mean reverting stochastic process. A numerical solution of a coupled system of HJB equations is implemented. We explore the impact of temperature volatility and regional asymmetries on emissions, contrasting the outcomes from the Stackelberg game with the choices made by a social planner. When players are identical, a classic tragedy of the commons is demonstrated in which players in the game choose higher carbon emissions and have lower utility as compared to the outcome with a social planner. Over certain values of state variables, the tragedy of the commons is shown to be exacerbated by increased temperature volatility and regional asymmetries in climate damages. Asymmetries in environmental preferences can, under certain conditions, result in a green paradox whereby green sentiments in one region cause the other region to increase emissions. Interestingly, we also find that a contrary “green bandwagon” effect is possible. At high levels of the carbon stock, green preferences in one region can cause the other region to reduce emissions.

Keywords: climate change, differential Stackelberg game, uncertainty, HJB equation

JEL codes: C73, Q52, Q54, Q58

*The authors gratefully acknowledge funding from the Global Risk Institute, globalriskinstitute.org

[†]Department of Economics, University of Waterloo, Waterloo, Ontario, Canada

[‡]Department of Economics, Wilfrid Laurier University, Waterloo, Ontario, Canada

[§]Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada

1 Introduction

Climate change caused by human activity represents a particularly intractable tragedy of the commons, which can only be addressed by cooperative actions of individual decision makers at both national and regional levels. The likely success of cooperative actions is hampered by the large incentives for free riding by decision makers who may delay making deep cuts in carbon emissions in hopes that others will do the “heavy lifting”. Further complicating the problem are the enormous uncertainties inherent in predicting climate responses to the buildup in atmospheric carbon stocks and resulting impacts on human welfare, including the prospects for adaptation and mitigation. These large uncertainties and the need for cooperative global action have been used by some as justification for delaying aggressive unilateral policy actions. And yet, many nations and sub-national jurisdictions have acted on their own to adopt policies to reduce carbon emissions even without national agreements or legislation in place. As a prominent example, since the Trump administration has reneged on the Paris Climate Accord, several states have vowed to go it alone and continue with aggressive climate policies. Other examples of jurisdictions taking unilateral carbon pricing initiatives are given in Kossey et al. (2015).

The observation that national or regional governments implement environmental regulations sooner or more aggressively than required by international agreements or national legislation has been studied by various researchers.¹ Local circumstances, including voter preferences, local damages from emissions, and strategic considerations regarding the actions of other jurisdictions, may play a role. A nation or region may be motivated to act ahead of others if it experiences relatively more severe local damages from emissions. Differences in environmental preferences may prompt some jurisdictions to take early action (Bednar-Friedl 2012). California and British Columbia (a province in Canada), both early adopters of carbon pricing, appear to have residents who are more environmentally aware, implying these governments acted in accordance with the preferences of a large segment of their voters. A survey of stakeholders involved in the introduction of the B.C. carbon tax

¹ Urpelainen (2009) and Williams (2012) examine the puzzle at a sub-regional level.

concluded that a number of factors were at work. These factors include: (i) a high priority given to environmental stewardship by B.C. residents and (ii) the fact that several other regional jurisdictions appeared to be poised in 2008 to take climate change more seriously (Clean Energy Canada 2015). Governments may choose environmental policies strategically to gain a competitive advantage or to shift emissions to other regions (Barcena-Ruiz 2006).

This paper examines the strategic interactions of decision makers responding to climate change, focusing on three central features of the problem: uncertainty, the incentive for free riding, and asymmetric characteristics of decision makers. We develop a model of a differential Stackelberg game involving two regions. Each region is a large emitter of green house gases which benefits from its own emissions, but also faces costs from the impact on global temperature of the cumulative emissions of both players. The modelling of the linkage between carbon emissions and global temperature is based on the assumptions of the well-known DICE model (Nordhaus & Sztorc 2013). To capture uncertainty, average global temperature is modelled as a stochastic process. We allow for differing damages of climate change for each region as well as differing preferences for reducing green house gas emissions. We explore the impact of these features on the optimal choice of emissions for each player and contrast with the choices made by a social planner.

It is well known that for differential games with closed loop strategies, only special classes of models result in well-posed mathematical problems for which it is possible to characterize Nash equilibria.² These include linear-quadratic games where the feedback controls depend linearly on the state variable, as well as certain forms of stochastic differential games where the state evolves according to an Ito process. This paper analyzes a stochastic differential game which we solve using numerical techniques. We make no prior assumption about the existence of a Nash equilibrium. Our approach is to discretize time to approximate the dynamic game as a series of one-shot games (which occur at discrete points in time) and solve for a Stackelberg equilibrium for each of these one-shot games. At each point in the state space, we can check if a Nash equilibrium is possible, or if the Stackelberg solution also represents a Nash equilibrium (See 4.2.2).

²Bressan (2011) provides an excellent overview of the mathematical theory of differential games.

There is a significant prior literature which examines the tragedy of the commons caused by polluting emissions in a differential game setting. The relevant differential game literature is reviewed in Section 2, but we note here two papers most closely related to our paper in their focus on asymmetry of players' utilities. Both employ economic models in a deterministic setting. Zagonari (1998) analyzes cooperative and non-cooperative games when the two players (countries) differ in the utility derived from a consumption good, the disutility caused by the pollution stock, and their concern for future generations as reflected in their discount rate. For the non-cooperative game, Zagonari examines the conditions under which the steady state stock of pollution might be less than under the cooperative game. Interestingly, Zagonari finds equilibria for which the steady state pollution stock is lower than in the cooperative game. In particular, this result holds if the country with stronger environmental preferences (the "eco-country") has sufficiently large disutility from pollution and either a relatively strong concern for future generations or relatively small utility from consumption goods.

Wirl (2011) also examines whether differences in environmental sentiments can mitigate the tragedy of the commons associated with a problem such as global warming. The author characterizes a multi-player game with green and brown players. Green players are distinguished from brown players by a penalty term in their objective function which depends on the extent to which their emissions exceed the social optimum. Wirl finds the presence of green players introduces discontinuous strategies with some interesting features. In the examples chosen, the effect of green players on total emissions is modest, as their actions increase the free riding of brown players. Wirl notes the possibility of a type of green paradox in which the increasing numbers of green players causes increased emissions, because brown players increase their emissions and more than offset the impact of green players' decisions.

Our paper contributes to this literature in several ways. We develop a more general model which includes uncertainty and closed loop strategies in a dynamic setting. The numerical results highlight the important influence of uncertainty in future temperature on optimal emissions choices and the carbon stock. We study the effect of asymmetry in damages and

environmental preferences on emissions choices and utility, contrasting the non-cooperative outcome with the outcome assuming a central planner empowered to make choices. Finally, we make a contribution in terms of the numerical methodology for solving a Stackelberg game under uncertainty with path dependent variables. We describe the method used to determine the closed loop optimal Stackelberg solution (which always exists) and then show how to determine if a Nash equilibrium exists. Our numerical solution procedure involves use of a finite difference discretization of the system of HJB equations. In contrast to much of the previous literature, the choice of damage function can be arbitrary function of state variables.

To preview our results, we highlight the crucial role of the damage function which specifies the harm from rising temperature, as has been noted by others (Weitzman 2012, Pindyck 2013). Very little reduction in carbon emissions occurs in the Stackelberg game or with the central planner using a conventional quadratic damage function. Exponentially increasing damages better reflect the catastrophic nature of damages anticipated if average global temperature should increase beyond 3°C above preindustrial levels. We also find that temperature uncertainty plays a key role. With a larger temperature volatility, optimal emissions are reduced for the players in the game as well as for the social planner. The social planner's response is relatively greater compared to the game for key values of the state variables (carbon stock and temperature) implying the benefit of cooperative action as demonstrated by a social planner increases at higher volatility. Asymmetric costs are also found to have an important effect on strategic interactions of players. Higher damage experienced by one player causes the other player to increase emissions over certain values of the state variables. This again implies that the advantage provided by a social planner is greater under asymmetric costs. Finally, we observe that an increase in green preferences by one player has an impact on the optimal actions of the other player, but the direction of that effect varies depending on the current temperature and stock of atmospheric carbon. We identify both a green paradox and a green bandwagon effect.

The remainder of the paper proceeds as follows. In Section 2, we provide a more detailed

literature review. The formulation of the climate change decision problem is described in Section 3. Section 4 provides a detailed description of the dynamic programming solution. Section 5 describes the detailed modelling assumptions and parameter values. Numerical results are described in Section 6, while Section 7 provides concluding comments.

2 Literature

This paper contributes to the literature on differential games dealing with trans-boundary pollution problems as well as to the developing literature on accounting for uncertainty in optimal policies to address climate change.

Economic models of climate change have long been criticized for arbitrary assumptions regarding functional forms and key parameter values as well as unsatisfactory treatment of key uncertainties including the possibility of catastrophic events.³ Of course, this is not surprising given the intractable nature of the climate change problem. Policies to address climate change have been extensively studied using the DICE (Dynamic Integrated Model of Climate and the Economy) model, a deterministic model developed in the 1990s, which has been revised and updated several times since then (Nordhaus 2013). Initially uncertainty was addressed through sensitivities or Monte Carlo analysis, but there has since been a significant research effort to address uncertainty using more robust methodologies. We mention only a sample of that literature. Kelly & Kolstad (1999) and Leach (2007) embed a model of learning into the DICE model to examine active learning by a social planner about key climate change parameters. More recent papers which incorporate stochastic components into one or more state variables in the DICE model include Crost & Traeger (2014), Ackerman et al. (2013) and Traeger (2014). Lemoine & Traeger (2014) extend the work of Traeger (2014) by incorporating the possibility of sudden shifts in system dynamics once parameters cross certain thresholds. Policy makers learn about the thresholds by observing the evolution of the climate system over time. Hambel et al. (2017) present a stochastic equilibrium model for optimal carbon emissions with key state variables, including carbon concentration,

³See Pindyck (2013) for a harsh critique.

temperature and GDP, modelled as stochastic differential equations. Chesney et al. (2017) examine optimal climate policies using a model in which global temperature is stochastic and assuming there is a known temperature threshold which will result in disastrous consequences if it is exceeded for a sustained period of time.

This previous work considers uncertainty in models with a single decision maker, abstracting from the strategic interactions of multiple decision makers which is a key feature of policy making for climate change. Differential game models have been used extensively to examine strategic interactions between players who benefit individually from polluting emissions but are also harmed by the cumulative emissions of all players. The literature addresses the strategic interactions of decision makers over time in deterministic and stochastic settings. Key assumptions, such as the information known to each player, determine whether the game can be described by a closed form mathematical solution.⁴ For example, open loop strategies, which depend solely on time, result when players know only the initial state of the system. Nash and Stackelberg equilibria for open loop strategies are well understood. In contrast, when players can directly observe the state of the system at every instant in time, feedback strategies (also called closed-loop or Markovian strategies) which depend on the state of the system may be employed. The resulting value functions satisfy a system of highly non-linear HJB partial differential equations. From the theory of partial differential equations it is known that if the system is non-stochastic, it should be hyperbolic in order for it to be well posed, in that it admits a unique solution depending continuously on the initial data (Bressan & Shen 2004). Our system of HJB equations is degenerate parabolic, which further complicates matters.

In games with feedback strategies only special classes of models are known to result in well-posed mathematical problems. These include zero-sum games, as well as linear-quadratic games. Linear-quadratic games have been used extensively in the economics literature to study pollution games, and some relevant papers, which admit closed form solutions, are detailed below. In this class of games, utility is a quadratic function of the state variable,

⁴See Bressan (2011) for a discussion of the challenges of finding appropriate mathematical models which result in closed form solutions.

while the state variable is linear in the control. Robust game models are also found with Nash feedback equilibria for stochastic differential games where the state evolves according to an Ito process such as

$$dx = f(t, x, u_1, u_2)dt + \sigma(t, x)d\mathcal{Z} \quad (1)$$

where x represents the state variable, t is time, u_1 and u_2 represent the controls of players 1 and 2, f and σ are known functions, and $d\mathcal{Z}$ is the increment of a Wiener process. As noted by Bressan (2011), for this case the value functions can be found by solving a Cauchy problem for a system of parabolic equations. The Cauchy problem is well posed if the diffusion tensor σ has full rank. In our case, the diffusion tensor is not of full rank (i.e. the system of PDEs is degenerate), hence we cannot expect that a Nash equilibrium will always exist. Additional discussion of the complexities of solving problems involving differential games can be found in Salo & Tahvonen (2001), Ludkovski & Sircar (2015), Harris et al. (2010), Cacace et al. (2013), Amarala (2015), and Ledvina & Sircar (2011).

Long (2010) and Dockner et al. (2000) provide surveys of the sizable literature addressing strategic interactions in the optimal control of pollution or natural resource exploitation using games, much of it in a deterministic setting. This literature focuses on the questions: (i) are players are better off with cooperative behaviour and (ii) how do the steady state levels of pollution compare under cooperative versus non-cooperative games.

Examples of dynamic differential pollution games in a non-stochastic setting include Dockner & Long (1993), Zagonari (1998), Wirl (2011), and List & Mason (2001). Under certain conditions, analytical closed-form solutions are found for linear and non-linear closed-loop strategies.

A few papers derive analytical solutions to differential pollution models in stochastic settings. Xepapadeas (1998) models global warming policy as a stochastic dynamic game in which damage is linear in the pollution stock and uncertainty in damage is described by geometric Brownian motion. An analytical solution is derived assuming exclusively linear strategies.

Wirl (2008) considers a dynamic pollution game with n symmetric players and differentiates between the cases of reversible and irreversible emissions. He characterizes the accumulation of pollution as an Ito process. The case of irreversible emissions requires a non-negativity constraint such that an explicit analytical solution is no longer possible. However, Wirl is able to analyze Nash equilibria under continuous Markov strategies using value matching and smooth pasting when optimal emissions are positive or zero. Wirl characterizes feasible stopping thresholds for the pollution stock at which emissions will cease. The author notes that multiple Nash equilibria cannot be ruled out. As expected, an increase in uncertainty reduces optimal emissions and this effect is more pronounced when emissions are irreversible. In the non-cooperative game, the stock of pollution grows in an unlimited fashion as the number of players increases. This contrasts with the cooperative game in which optimal emissions and pollution levels decline as the number of players increases.

Finally, Nkuiya (2015) analyzes a pollution control game with n symmetric players when there is a risk of sudden jump in damages or a catastrophe at an unknown future date. The switch from the low damage to high damage state is modelled as a Poisson process. This allows the stochastic differential game to be transformed into a deterministic game which admits an analytic solution. He distinguishes two cases: one in which damages may jump to a level such that all economic activity ceases (doomsday scenario) and the other in which damages jump to a higher level, but economic activity continues (non-doomsday scenario). For the non-doomsday scenario, linear strategies, and an exogenous hazard rate, an increase in the hazard rate causes players to reduce emissions. In the doomsday case under linear strategies, an increase in the hazard rate adds to the effective discount rate and causes players to increase emissions. These results do not necessarily hold if the probability of jumping to the doomsday state depends endogenously on players' emission decisions, nor if players commit to non-linear strategies, even when the jump probability is exogenous.

There is a developing literature on the numerical solution of dynamic games in the context of non-renewable resource markets. Some earlier papers developed models where two or more players extract from a common stock of resource. Examples include van der Ploeg

(1987) and Dockner et al. (1996). Salo & Tahvonen (2001) were among the first to explore oligopolistic natural resource markets in a differential Cournot game using closed loop strategies. Prior to that, the focus had been on open-loop strategies, because of their tractability. Of course, in general, open loop strategies are sub-optimal compared to closed loop strategies. Salo & Tahvonen (2001) developed an analytical solution of affine-quadratic specifications and demonstrated a numerical method for other functional forms using a Markov chain approximation.

Harris et al. (2010) study the extraction of an exhaustible resource as an N-player continuous time Cournot game when players have heterogeneous costs. They note that existence, uniqueness and regularity of value functions are not well understood and that numerical solutions represent a major challenge. Harris et al. (2010) present an asymptotic approximation for a low exhaustible case (i.e. small cost of the alternate technology) and a numerical solution given a restriction on the cost of the alternate technology. Ludkovski & Sircar (2012) extend the work of Harris et al. (2010) by adding exploration to the model. Ludkovski & Yang (2015) includes both exploration and stochastic demand.

3 Problem Formulation

This section provides an overview of the climate change decision model. Details of functional forms and parameter values are provided in Section 5. A summary of variable names is given in Table 1. We model the optimal timing and stringency of environmental regulations (in terms of the reduction of greenhouse gas emissions) as a stochastic optimal control problem. We consider two cases: a Stackelberg game and a social planner. The players in the Stackelberg game are two regions, each contributing to the atmospheric stock of greenhouse gases - which, for simplicity, we will refer to as the carbon stock. These regions may be thought of as single nations or groups of nations acting together, but each is a major contributor to the global carbon stock. Each region seeks to maximize discounted expected utility by making emission choices taking into account the optimal actions of the other region. The social planner chooses emission levels in each region so as to maximize the expected sum of

Table 1: List of Model Variables

Variable	Description
$E_p(t)$	Emissions in region p
\bar{E}_p	benchmark emissions for player p
e_1, e_2	Particular realizations of $E_p(t)$
ω, ω_2	any control choice by players 1 and 2
$S(t)$	Stock of pollution at time t , a state variable
s	A realization of $S(t)$
\bar{S}	preindustrial level of carbon
$\rho(X, S, t)$	Rate of natural removal of the pollution stock
σ	temperature volatility
$\eta(t)$	speed of mean reversion in temperature equation
$X(t)$	Average global temperature, a state variable
x	A realization of $X(t)$
\bar{X}	long run equilibrium level of temperature, °C above pre-industrial levels
$B_p(E_p, t)$	Damages from pollution
$C_p(X, t)$	Damages from pollution
$g_p(t)$	Emissions reduction in region p relative to a target
θ_p	Willingness to pay in region p for emissions reduction from a target
$A_p(g_p(t))$	Green reward benefits from emissions reductions
π_p	Flow of net benefits to region p
r	Discount rate

utilities from both regions.

Regions emit carbon in order to generate income. For simplicity we assume that there is a one to one relation between emissions and regional income. The two regions are indexed by $p = 1, 2$ and E_p refers to carbon emissions from region p . The stock of atmospheric carbon, S , is augmented by the emissions of each player and is reduced by a natural cycle whereby carbon is removed from the atmosphere and absorbed into other carbon sinks. The removal of carbon from the atmosphere can be described by decay function, $\rho(X, S, t)$, which in theory may depend on the average surface temperature, X , the stock of carbon, $S(t)$, and

time, t . $\rho(X, S, t)$ is referred to as the removal rate. For simplicity, as described in Section 5, we will later drop the dependence on X and S , assuming that ρ is a function only of time. However, our solution technique can easily accommodate more general functional forms for ρ . The evolution of the carbon stock over time is described by the deterministic differential equation:

$$\frac{dS(t)}{dt} = E_1 + E_2 + (\bar{S} - S(t))\rho(X, S, t); S(0) = s_0 \quad S \in [s_{min}, s_{max}]. \quad (2)$$

\bar{S} is the pre-industrial equilibrium level of atmospheric carbon.

The mean global increase in temperature above the pre-industrial level, denoted by X , is described by an Ornstein Uhlenbeck process:

$$dX(t) = \eta(t) \left[\bar{X}(S, t) - X(t) \right] dt + \sigma dZ. \quad (3)$$

where $\eta(t)$ represents the speed of mean reversion and is a deterministic function of time. \bar{X} represents the long run mean of global average temperature which depends on the stock of carbon and time. σ is the volatility parameter, assumed to be constant. The detailed specification of these functions and parameters is given in Section 5. dZ is the increment of a standard Weiner process, intended to capture the volatility in the earth's temperature due to random effects.

The net benefits from carbon emissions are represented as a general function π_p :

$$\pi_p = \pi_p(E_1, E_2, X, S, t) \quad (4)$$

More specifically, π is composed of the benefits from emissions, $B_p(E_p, t)$, the damages from increasing temperature, $C_p(X, t)$, and a green reward that results from reducing emissions relative to a given target or baseline level, $A_p(g_p(t))$:

$$\pi_p = B_p(E_p, t) - C_p(X, t) + A_p(g_p(t)) \quad p = 1, 2; \quad (5)$$

where $g_p(t)$ refers to emissions reduction. The detailed specification of benefits, damages, and the green reward is left to Section 5

It is assumed that the control (choice of emissions) is adjusted at discrete decision times denoted by:

$$\mathcal{T} = \{t_0 = 0 < t_1 < \dots < t_m \dots < t_M = T\}. \quad (6)$$

Let t_m^- and t_m^+ denote instants just before and after t_m , with $t_m^- = t_m - \epsilon$ and $t_m^+ = t_m + \epsilon$, $\epsilon \rightarrow 0^+$, and where T is the time horizon of interest.

$e_1^+(E_1, E_2, X, S, t_m)$ and $e_2^+(E_1, E_2, X, S, t_m)$ denote the controls implemented by the players 1 and 2 respectively, which are contained within the set of admissible controls: $e_1^+ \in Z_1$ and $e_2^+ \in Z_2$. We can specify a control set which contains the optimal controls for all t_m .

$$K = \{(e_1^+, e_2^+)_{t_0=0}, (e_1^+, e_2^+)_{t_1=1}, \dots, (e_1^+, e_2^+)_{t_M=T}\}. \quad (7)$$

In this paper we will consider three possibilities for selection of the controls (e_1^+, e_2^+) at $t \in \mathcal{T}$: Stackelberg, Nash, and social planner. We delay the precise specification of how these controls are determined until Section 4.2.

Regardless of the control strategy, the value function for player p , $V_p(e_1, e_2, x, s, t)$ is defined as:

$$V_p(e_1, e_2, x, s, t) = \mathbb{E}_K \left[\int_{t'=t}^T e^{-rt'} \pi_p(E_1(t'), E_2(t'), X(t'), S(t')) dt' + e^{-r(T-t)} V(0, 0, X(T), S(T), T) \right. \\ \left. \Big| E_1(t) = e_1, E_2(t) = e_2, X(t) = x, S(t) = s \right], \quad (8)$$

where $\mathbb{E}_K[\cdot]$ is the expectation under control set K . Note that lower case letters e_1, e_2, x, s have been used to denote realizations of the state variables E_1, E_2, X, S . The value in the final time period, T , is assumed to be the present value of a perpetual stream of expected net benefits at given carbon stock, S , and temperature levels, X , with emissions set to zero. This is reflected in the term $V(0, 0, x, s, T)$ and is described in Section 4.1 as a boundary condition.

4 Dynamic Programming Solution

Using dynamic programming, we solve the problem represented by Equation (8) backwards in time, breaking the solution phases up into two components for $t \in (t_m^-, t_m^+)$ and (t_m^+, t_{m+1}^-) . In the interval (t_m^-, t_m^+) , we determine the optimal controls, while in the interval (t_m^+, t_{m+1}^-) , we solve a system of PDEs. As a visual aid, Equation (9) shows the noted time intervals going forward in time,

$$t_m^- \rightarrow t_m^+ \rightarrow t_{m+1}^- \rightarrow t_{m+1}^+ . \quad (9)$$

4.1 Advancing the solution from $t_{m+1}^- \rightarrow t_m^+$

The solution proceeds going backward in time from $t_{m+1}^- \rightarrow t_m^+$. Define the differential operator, \mathcal{L} for player p , in Equation (10). The arguments in the V_p function, as well as in η and ρ , have been suppressed when there is no ambiguity.

$$\mathcal{L}V_p \equiv \frac{(\sigma)^2}{2} \frac{\partial^2 V_p}{\partial x^2} + \eta(\bar{X} - x) \frac{\partial V_p}{\partial x} + [(e_1 + e_2) + \rho(\bar{S} - s)] \frac{\partial V_p}{\partial s} - rV_p; \quad p = 1, 2 . \quad (10)$$

where r is the discount rate. Then using standard techniques (Dixit & Pindyck 1994), the equation satisfied by the value function, V_p is expressed as:

$$\frac{\partial V_p}{\partial t} + \pi_p(e_1, e_2, x, s, t) + \mathcal{L}V_p = 0, \quad p = 1, 2 . \quad (11)$$

The domain of Equation (11) is $(e_1, e_2, x, s, t) \in \Omega^\infty$, where $\Omega^\infty \equiv Z_1 \times Z_2 \times [x^0, \infty] \times [\bar{S}, \infty] \times [0, \infty]$. x^0 would be zero degrees Kelvin in our units. For computational purposes, we truncate the domain Ω^∞ to Ω , where $\Omega \equiv Z_1 \times Z_2 \times [x_{min}, x_{max}] \times [\bar{S}, s_{max}] \times [0, T]$. T , \bar{S} , s_{max} , Z_1 , Z_2 , x_{min} , and x_{max} are specified based on reasonable values for the climate change problem, and are given in Section 5.

Remark 1 (Admissible sets Z_1, Z_2). *We will assume in the following that Z_1, Z_2 are compact discrete sets, which would be the only realistic situation.*

Boundary conditions for the PDEs are specified below.

- As $x \rightarrow x_{max}$, it is assumed that $\frac{\partial^2 V_p}{\partial x^2} \rightarrow 0$. This boundary condition is commonly used in the literature and implies that the impact of volatility at very high temperature levels is unimportant relative to the size of the damages. Assuming that $x_{max} > \bar{X}$, Equation (11) has outgoing characteristics with $V_{xx} = 0$ at $x = x_{max}$ and hence no other boundary conditions are required.
- As $x \rightarrow x_{min}$, where x_{min} is below the pre-industrial temperature, the effect of volatility is small compared to the drift term. Hence we set $\sigma = 0$ at $x = x_{min}$. Assuming $x_{min} < \bar{X}$ then Equation (11) has outgoing characteristics at $x = x_{min}$ and no other boundary conditions are required. Note that we will show that $\pi_p \geq 0$ at $x = x_{min}$.
- As $s \rightarrow s_{max}$, it is assumed that emissions do not increase s beyond the limit of s_{max} . s_{max} is set to be a large enough value so that there is no impact on utility or optimal emission choices for s levels of interest. We have verified this in our computational experiments. This amounts to dropping the term $\frac{\partial V_p}{\partial S}(e_1 + e_2)$ from Equation (10). This can be justified by noting that if $s_{max} \gg \bar{S}$ then $\rho(\bar{S} - S) \gg (e_1 + e_2)$ for reasonable values of e_1 and e_2 .
- As $s \rightarrow \bar{S}$, no extra boundary condition is needed as we assume $e_1, e_2 \geq 0$ hence the Equation has outgoing characteristics at $s = \bar{S}$.
- At $t = T$, it is assumed that V_p is equal to the present value of the infinite stream of benefits associated with a given temperature when emissions are set to zero. Essentially, it is assumed that we receive the costs associated with that temperature in perpetuity and T is large enough that we assume the world has decarbonized.

More details of the numerical solution of the system of PDEs are provided in Appendix A.

4.2 Advancing the solution from $t_m^+ \rightarrow t_m^-$

Going backward in time, the optimal control, is determined between $t_m^+ \rightarrow t_m^-$. We consider three possibilities for selection of the controls (e_1^+, e_2^+) at $t \in \mathcal{T}$: Stackelberg, Nash, and social planner. We include the Nash case for reference only. We remind the reader that our controls are assumed to be feedback, i.e. a function of state. However, to avoid notational clutter in the following, we will fix (e_1, e_2, s, x, t_m) , so that, if there is no ambiguity, we will write (e_1^+, e_2^+) which will be understood to mean $(e_1^+(e_1, e_2, s, x, t_m), e_2^+(e_1, e_2, s, x, t_m))$.

4.2.1 Stackelberg Game

In the case of a Stackelberg game, suppose that, in forward time, player 1 goes first, and then player 2. Conceptually, we can then think of the time intervals (in forward time) as $(t_m^-, t_m]$, (t_m, t_m^+) . Player 1 chooses control e_1^+ in $(t_m^-, t_m]$, then player 2 chooses control e_2^+ in (t_m, t_m^+) .

We suppose at t_m^+ , we have the value functions $V_1(e_1, e_2, s, x, t_m^+)$ and $V_2(e_1, e_2, s, x, t_m^+)$.

Definition 1 (Response set of player 2). *The best response set of player 2, $R_2(\omega_1, e_1, e_2, s, x, t_m)$ is defined to be the best response of player 2 to a control ω_1 of player 1.*

$$R_2(\omega_1, e_1, e_2, s, x, t_m) = \operatorname{argmax}_{\omega_2' \in Z_2} V_2(\omega_1, \omega_2', s, x, t_m^+) ; \omega_1 \in Z_1 . \quad (12)$$

Remark 2 (Tie breaking). *We break ties by (i) staying at the current emission level if possible, or (ii) choosing the lowest emission level. Rule (i) has priority over rule (ii). Note that rule (i) corresponds to an infinitesimal switching cost and rule (ii) to an infinitesimal green reward (see Section 5.3.3). Consequently there are no ties after applying either of these rules.*

Similarly, we define the best response set of player 1.

Definition 2 (Response set of player 1). *The best response set of player 1, $R_1(\omega_2, e_1, e_2, s, x, t_m)$ is defined to be the best response of player 1 to a control ω_2 of player 2.*

$$R_1(\omega_2, e_1, e_2, s, x, t_m) = \operatorname{argmax}_{\omega'_1 \in Z_1} V_1(\omega'_1, \omega_2, s, x, t_m^+) ; \omega_2 \in Z_2 . \quad (13)$$

Again, to avoid notational clutter, we will fix (e_1, e_2, s, x, t_m) so that we can write without ambiguity $R_1(\omega_2) = R_1(\omega_2, e_1, e_2, s, x, t_m)$ and $R_2(\omega_1) = R_2(\omega_1, e_1, e_2, s, x, t_m)$.

Remark 3 (Dependence on states e_1, e_2). *In Equations (12) and (13) the tie breaking rule induces dependence on the initial state, e_1, e_2 .*

Definition 3 (Stackelberg Game: Player 1 first). *The optimal controls (e_1^+, e_2^+) assuming player 1 goes first are given by*

$$\begin{aligned} e_1^+ &= \operatorname{argmax}_{\omega'_1 \in Z_1} V_1(\omega'_1, R_2(\omega'_1), s, x, t_m^+) , \\ e_2^+ &= R_2(e_1^+) . \end{aligned} \quad (14)$$

Since we use dynamic programming, we determine the optimal controls using the following Algorithm.

Stackelberg Control: Player 1 first

Input: $V_1(e_1, e_2, s, x, t_m^+), V_2(e_1, e_2, s, x, t_m^+)$.

Step 1: Compute the best response set for player 2 assuming player 1 chooses control ω_1 first, $\forall \omega_1 \in Z_1$, using Equation (12), giving $R_2(\omega_1)$.

Step 2: Determine an optimal pair (e_1^+, e_2^+) using Equation (14).

Determine solution at t_m^-

$$\begin{aligned} V_1(e_1, e_2, s, x, t_m^-) &= V_1(e_1^+(\cdot), e_2^+(\cdot), s, x, t_m^+) ; \\ V_2(e_1, e_2, s, x, t_m^-) &= V_2(e_1^+(\cdot), e_2^+(\cdot), s, x, t_m^+) ; . \end{aligned} \quad (15)$$

Output: $V_1(e_1, e_2, s, x, t_m^-), V_2(e_1, e_2, s, x, t_m^-)$

4.2.2 Nash Equilibrium

We again fix (e_1, e_2, s, x, t_m) , so that we understand that $e_p^+ = e_p^+(e_1, e_2, s, x, t_m), R_p(\omega) = R_p(\omega, e_1, e_2, s, x, t_m)$.

Definition 4 (Nash Equilibrium). *Given the best response sets $R_2(\omega_1), R_1(\omega_2)$ defined in Equations (12)-(13), then the pair (e_1^+, e_2^+) is a Nash equilibrium point if and only if*

$$e_1^+ = R_1(e_2^+) \quad ; \quad e_2^+ = R_2(e_1^+) . \quad (16)$$

From Definition 3 of a Stackelberg game, if player 1 goes first, we have the optimal pair $(\hat{e}_1^+, \hat{e}_2^+)$

$$\begin{aligned} \hat{e}_1^+ &= \operatorname{argmax}_{\omega'_1 \in Z_1} V_1(\omega'_1, R_2(\omega'_1), s, x, t_m^+) , \\ \hat{e}_2^+ &= R_2(\hat{e}_1^+) . \end{aligned} \quad (17)$$

Similarly, we have the pair $(\bar{e}_1^+, \bar{e}_2^+)$ if player 2 goes first

$$\begin{aligned} \bar{e}_2^+ &= \operatorname{argmax}_{\omega'_2 \in Z_2} V_2(R_1(\omega'_2), \omega'_2, s, x, t_m^+) , \\ \bar{e}_1^+ &= R_1(\bar{e}_2^+) . \end{aligned} \quad (18)$$

Suppose $(\hat{e}_1^+, \hat{e}_2^+) = (\bar{e}_1^+, \bar{e}_2^+)$. Consequently, we have $(e_1^+, e_2^+) = (\hat{e}_1^+, \hat{e}_2^+) = (\bar{e}_1^+, \bar{e}_2^+)$ and we replace the \hat{e}_p^+ by e_p^+ and \bar{e}_p^+ by e_p^+ in Equations (17) - (18) giving

$$\begin{aligned} e_1^+ &= \operatorname{argmax}_{\omega'_1 \in Z_1} V_1(\omega'_1, R_2(\omega'_1), s, x, t_m^+) , \\ e_2^+ &= \operatorname{argmax}_{\omega'_2 \in Z_2} V_2(R_1(\omega'_2), \omega'_2, s, x, t_m^+) , \\ e_1^+ &= R_1(e_2^+) \quad ; \quad e_2^+ = R_2(e_1^+) , \end{aligned} \quad (19)$$

which is a Nash equilibrium from Definition 4. We can summarize this result in the following

Proposition 1 (Sufficient condition for a Nash Equilibrium). *A Nash equilibrium exists at a point (e_1, e_2, s, x, t_m) if $(\hat{e}_1^+, \hat{e}_2^+) = (\bar{e}_1^+, \bar{e}_2^+)$.*

Remark 4 (Checking for a Nash equilibrium). *A necessary and sufficient condition for a Nash Equilibrium is given by condition (16). However a sufficient condition for a Nash equilibrium in the Stackelberg game is that the optimal control of either player is independent of who goes first.*

In our numerical experiments we find Nash equilibria exist only at some points (not all) over the state space. This is, of course, not surprising since the system of PDEs is degenerate. Insley & Forsyth (2018) examine this issue, along with other possible games, such as leader-leader, follower-follower games, and interleaved games.

4.2.3 Social Planner

For the social planner case, we have that an optimal pair (e_1^+, e_2^+) is given by

$$(e_1^+, e_2^+) = \operatorname{argmax}_{\substack{\omega_1 \in Z_1 \\ \omega_2 \in Z_2}} \left\{ V_1(\omega_1, \omega_2, s, x, t_m^+) + V_2(\omega_1, \omega_2, s, x, t_m^+) \right\}. \quad (20)$$

and as a result

$$V_1(e_1, e_2, s, x, t_m^-) = V_1(e_1^+, e_2^+, s, x, t_m^+) \quad ; \quad V_2(e_1, e_2, s, x, t_m^-) = V_2(e_1^+, e_2^+, s, x, t_m^+) . \quad (21)$$

Ties are broken by minimizing $|V_1(e_1^+, e_2^+, s, x, t_m^+) - V_2(e_1^+, e_2^+, s, x, t_m^+)|$. In other words, the social planner picks the emissions choices which give the most equal distribution of welfare across the two players.

5 Detailed model specification and parameter values

This section describes the functional forms and parameter values used in the numerical application. Assumed parameter values are summarized in Table 2.

5.1 Carbon stock details

The evolution of the carbon stock is described in Equation (2). In Integrated Assessment Models, there is typically a detailed specification of the exchange of carbon emissions between the various carbon reservoirs: the atmosphere, the terrestrial biosphere and different ocean layers (Nordhaus 2013, Lemoine & Traeger 2014, Traeger 2014, Golosov et al. 2014). In Equation (2) the removal function is given as $\rho(X, S, T)$. In our numerical application, we use a simplified specification, based on Traeger (2014), to avoid the creation of additional path dependent variables which increases computational complexity. We denote the rate at which carbon is removed from the atmosphere by $\rho(t)$ and assume it is a deterministic function of time which approximates the removal rates in the DICE 2016 model.

$$\rho(t) = \bar{\rho} + (\rho_0 - \bar{\rho})e^{-\rho^*t}$$

ρ_0 is the initial removal rate per year of atmospheric carbon, $\bar{\rho}$ is a long run equilibrium rate of removal, and ρ^* is the rate of change in the removal rate. Specific parameter assumptions for this Equation are given in Table 2. The resulting removal rate starts at 0.01 per year and falls to 0.0003 per year within 100 years.

The pre-industrial equilibrium level of carbon, \bar{S} in Equation (2), is assumed to be 588 gigatonnes (GT) based on estimates used in the DICE (2016)⁵ model for the year 1750. The allowable range of carbon stock is given by $s_{min} = 588$ and $s_{max} = 10000$. s_{max} is set well

⁵The 2013 version of the DICE model is described in Nordhaus & Sztorc (2013). GAMS and Excel versions for the updated 2016 version are available from William Nordhaus's website: <http://www.econ.yale.edu/nordhaus/homepage/>.

Table 2: Base Case Parameter Values

Parameter	Description	Equation Reference	Assigned Value
\bar{S}	Pre-industrial atmospheric carbon stock	(2)	588 GT carbon
s_{min}	Minimum carbon stock	(2)	588 GT carbon
s_{max}	Maximum carbon stock	(2)	10000 GT carbon
$\bar{\rho}, \rho_0, \rho^*$	Parameters for carbon removal equation	(22)	0.0003, 0.01, 0.01
ϕ_1, ϕ_2, ϕ_3	Parameters of temperature equation	(22)	0.02, 1.1817, 0.088
ϕ_4	Forcings at CO2 doubling	(24)	3.681
$F_{EX}(0)$ $F_{EX}(100)$	Parameters from forcing equation	(24)	0.5 1
α_1, α_2	Ratio of the deep ocean to surface temp, $\alpha(t) = \alpha_1 + \alpha_2 \times t$, t is time in years with 2015 set as year 0	(22)	0.008, 0.0021
σ	Temperature volatility	(22)	0.1
x_{min}, x_{max}	Upper and lower limits on average temperature, °C	(22)	-3, 20
a_1, a_2	Parameter in benefit function, player p	(26)	10
Z_1, Z_2	Admissible controls	(7)	0, 3, 7, 10
\bar{E}	Baseline emissions	(29)	10
b_1, b_2	Cost scaling parameter, players 1 & 2 respectively	(28)	15, 15
κ_1	Linear parameter in cost function for both players	(28)	0.05
κ_2	Exponent in cost function for both players	(28)	2 or 3
κ_3	Term in exponential cost function for both players	(27)	1
θ_P	WTP for emissions reduction by player p	(5)	0 or 3
T	terminal time		150 years
r	risk free rate	(10)	0.01

above the 6000 GT carbon in Nordhaus (2013) and will not be a binding constraint in the numerical examples.⁶ A 2014 estimate of the atmospheric carbon level is 840 GT.⁷

5.2 Stochastic process temperature: details

Equation (3) specifies the stochastic differential equation which describes temperature X and includes the parameters $\eta(t)$ and $\bar{X}(t)$. To relate Equation (3) to common forms used in the climate change literature, we rewrite it in the following format:

$$dX = \phi_1 \left[F(S, t) - \phi_2 X(t) - \phi_3 [1 - \alpha(t)] X(t) \right] dt + \sigma dZ \quad (22)$$

where ϕ_1 , ϕ_2 , ϕ_3 and σ are constant parameters.⁸ The drift term in Equation (22) is a simplified version of temperature models typical in Integrated Assessment Models, based on Lemoine & Traeger (2014). $\alpha(t)$ represents the ratio of the deep ocean temperature to the mean surface temperature and, for simplicity, is specified as a deterministic function of time.⁹ Equation (22) is equivalent to Equation (3) with:

$$\begin{aligned} \eta(t) &\equiv \phi_1 \left(\phi_2 + \phi_3 (1 - \alpha(t)) \right) \\ \bar{X}(t) &\equiv \frac{F(S, t)}{(\phi_2 + \phi_3 (1 - \alpha(t)))}. \end{aligned} \quad (23)$$

$F(S, t)$ refers to radiative forcing, and it measures additional energy trapped at the earth's surface due to the accumulation of carbon in the atmosphere compared to preindustrial levels and also includes other greenhouse gases,

$$F(S, t) = \phi_4 \left(\frac{\ln(S(t)/\bar{S})}{\ln(2)} \right) + F_{EX}(t). \quad (24)$$

⁶Golosov et al. (2014) chose a maximum atmospheric carbon stock of 3000 GT which is intended to reflect the carbon stock that results if most of the predicted stocks of fossil fuels are burned in “a fairly short period of time” (page 67).

⁷According to the Global Carbon Project, 2014 global atmospheric CO₂ concentration was 397.15 ± 0.10 ppm on average over 2014. At 2.21 GT carbon per 1 ppm CO₂, this amounts to 840 GT carbon. (www.globalcarbonproject.org)

⁸ ϕ_1 , ϕ_2 , ϕ_3 are denoted as ξ_1 , ξ_2 , and ξ_3 in Nordhaus (2013).

⁹We are able to get a good match to the DICE2016 results using a simple linear function of time.

ϕ_4 indicates the forcing from doubling atmospheric carbon.¹⁰ $F_{EX}(t)$ is forcing from causes other than carbon and is modelled as an exogenous function of time as specified in Lemoine & Traeger (2014) as follows:

$$F_{EX}(t) = F_{EX}(0) + 0.01(F_{EX}(100) - F_{EX}(0)) \min\{t, 100\} \quad (25)$$

The values for the parameters in Equation (22) are taken from the DICE (2016) model. Note that $\phi_1 = 0.02$ which is the value reported in DICE (2016) divided by five to convert to an annual basis from the five year time steps used in the DICE (2016) model. $F_{EX}(0)$ and $F_{EX}(100)$ (Equation (24)) are also from the DICE (2016) model. The ratio of the deep ocean temperature to surface temperature, $\alpha(t)$, is modelled as a linear function of time. This function approximates the average values from the DICE (2016) base and optimal tax cases.

Useful intuition about the temperature model can be gleaned by substituting parameter values from Table 2 to determine implied values for the speed of mean reversion $\eta(t)$ and the long run temperature mean $\bar{X}(t)$ in Equation (3) for 2015. Using the definitions in Equation (23) it can be determined that $\eta(t) = 0.02$ and $\bar{X} = 1.9^\circ\text{C}$. This value for η implies that, ignoring volatility, temperature would revert to its long run mean in about 50 years. The long run temperature of 1.9°C is above today's value of 1°C above preindustrial levels. This temperature model and assumed parameter values imply considerable momentum in the temperature trajectory.

Figure 1 shows the changes in global surface temperature relative to 1951 to 1980 averages.¹¹ Based on this data the volatility parameter was estimated using maximum likelihood to be approximately $\sigma = 0.1/\sqrt{\text{year}}$. For the numerical solution we choose $x_{min} = -3$ and $x_{max} = 20$.

As time tends to infinity, the probability density of an Ornstein-Uhlenbeck process is Gaussian with mean \bar{X} and variance $\sigma^2/2\eta$. Our assumed parameter values therefore give

¹⁰ ϕ_4 translates to Nordhaus's η (Nordhaus & Sztorc 2013).

¹¹The data is from NASA's Goddard Institute for Space Studies and is available on NASA's web site Global Climate Change: <https://climate.nasa.gov/vital-signs/global-temperature/>.

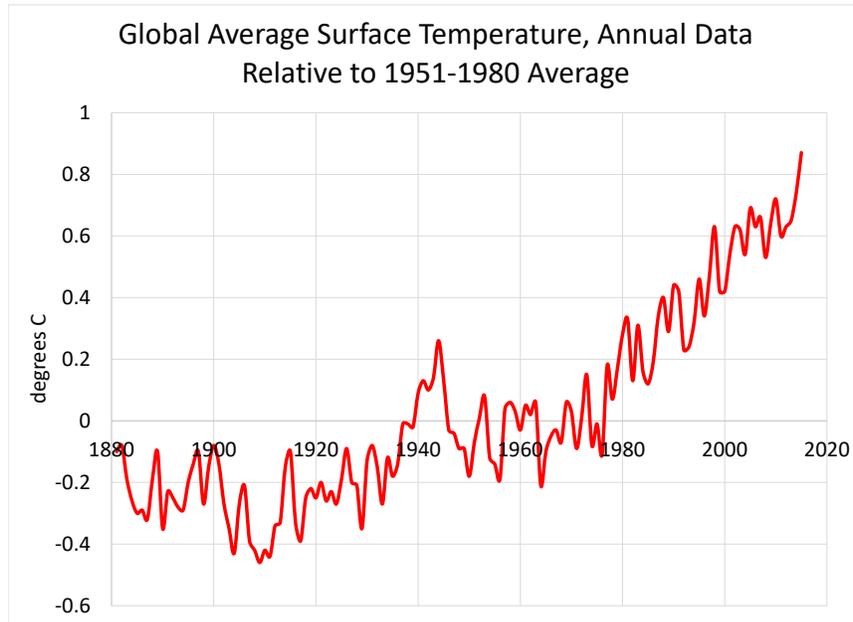


Figure 1: Global land-ocean temperature index, degrees C, annual averages since 1880 relative to the 1951-1980 average

a long run standard deviation of 0.44°C and mean of 1.9°C . This implies there is a 2.3% probability that temperature could rise by 2 standard deviations (0.88°C) due solely to randomness, independent of carbon emissions. We conclude that volatility should be an important consideration in any analysis of climate change policies. For further intuition, Figure 2 presents simulations of the temperature path given our assumed model. We compare 15 realizations of the temperature process if the volatility term $\sigma = 0.1/\sqrt{\text{year}}$ and if it is doubled to $0.2/\sqrt{\text{year}}$.

5.3 Benefits, Damages and the Green Reward

The term π_p in Equation (5) comprises benefits and damages from emissions as well as the green reward. This section describes these components.

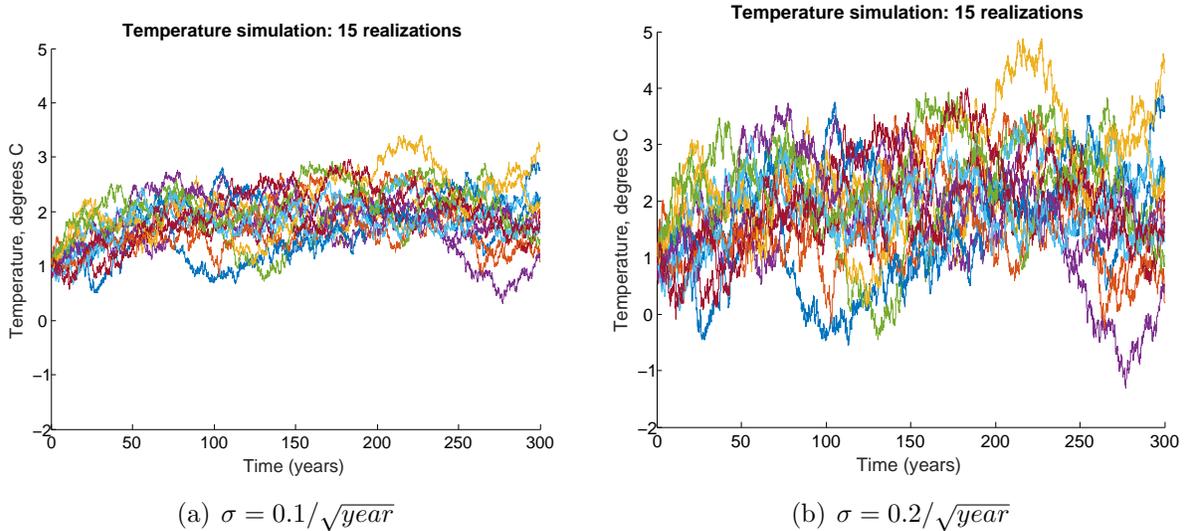


Figure 2: Simulations of Ornstein-Uhlenbeck process for temperature, Equation (3)

5.3.1 Benefits

As is common in the pollution game literature, the benefits of emissions are quadratic according the following utility function:

$$B_p(E_p) = a_p E_p(t) - E_p^2(t)/2, \quad p = 1, 2 \quad (26)$$

a_p is a constant parameter which may be different for different players. As in List & Mason (2001), $E_p \in [0, a_p]$ so that the marginal benefit from emissions is always positive. In the numerical example, there are four possible emissions levels for each player $E_p \in \{0, 3, 7, 10\}$ in gigatonnes (Gt) per year of carbon and we set $a_1 = a_2 = 10$.

5.3.2 Damages

Assumptions regarding damages from increasing temperatures are speculative, and this is a highly criticized element of climate change models. The DICE model (and others) specify damages as a multiple of GDP and a quadratic function of temperature, implying that damages never exceed 100% of GDP. This formulation ignores possible catastrophic effects.

Damage function calibrations are generally based on estimates for the zero to 3°C range above preindustrial temperatures.

A multiplicative formulation is not appropriate for the model used in this paper in which benefits are zero if emissions are zero (Equation (26)). This is because the multiplicative damage function implies that choosing zero emissions would reduce damages immediately to zero. For this analysis an additive damage function is adopted in which damages rise exponentially with temperature:

$$C_p(t) = \kappa_1 e^{\kappa_3 X(t)} \quad p = 1, 2, . \quad (27)$$

where κ_3 is a constant and $p = 1, 2$ refers to the two players. We also explore results with quadratic or cubic forms of the cost function

$$C_p(X, t) = \kappa_1 X(t)^{\kappa_2} \quad p = 1, 2, . \quad (28)$$

where κ_1 and κ_2 are constants.

We choose the parameters in the damage functions (Equation (28) and (27)) so that damages represent a reasonable portion of benefits at current temperatures levels (i.e. at 0.86 degrees C over preindustrial levels). Base case values for κ_1 , κ_2 and κ_3 imply damages of about 1% of benefits at current temperature levels. Figure 3 compares the three cost functions as a percentage of benefits. The comparison is for the exponential function, Equation (27), compared with the power function, Equation (28), with the exponent set to 2 or 3. We observe that the three cost functions are virtually indistinguishable up to 3 °C above preindustrial levels. After 3 °C the cost functions diverge dramatically. We choose the exponential cost function as our base case as it implies that for temperature increases above 3 °C, damages from climate change would be disastrous, which seems a reasonable supposition.

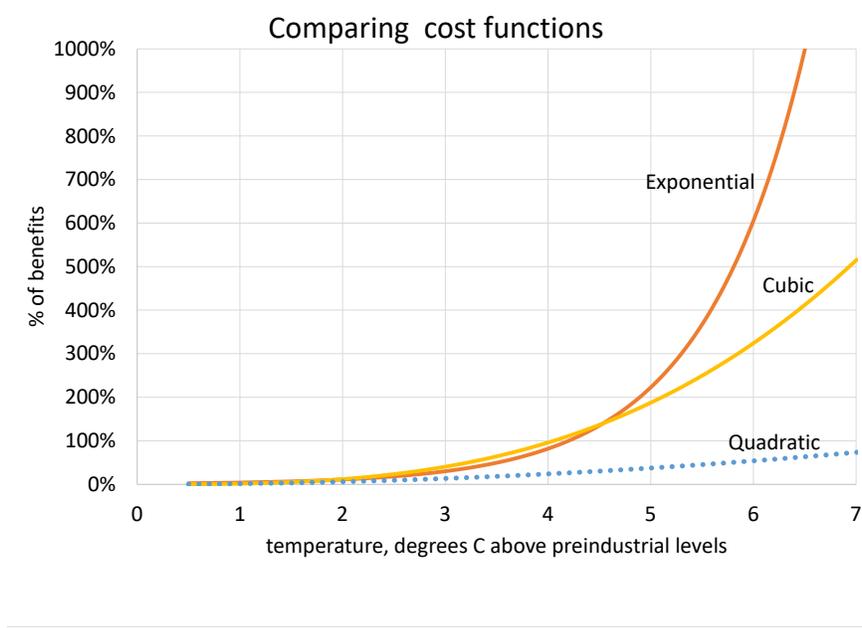


Figure 3: Comparing costs of increased temperatures as a % of benefits for different cost functions

5.3.3 Green Reward

We define emissions reduction, $g_p(t)$, relative to a baseline level of emissions level, \bar{E} , for each region.

$$g_p(t) = \max(\bar{E}_p - E_p(t), 0), \quad p = 1, 2 \quad (29)$$

Citizens of each region are assumed to value emissions reduction as contributing to the public good. We denote the degree of environmental awareness in a region by θ_p which represents a willingness to pay for emissions reduction because of a desire to be good environmental citizens, distinct from the expressions for the benefits and costs of emissions as defined in Equations (26) and (28) or (27).

The benefit from emissions reduction, called the green reward, A_p , depends on environmental awareness as well as emissions reduction in both regions:

$$A_p(t) = \theta_p g_p(t), \quad p = 1, 2. \quad (30)$$

In our base case, $\theta_p = 0$ for both players initially. We then explore differential green preferences by setting $\theta_p = 3$ for one of the players. In future work, we will explore the possibility that environmental preferences may evolve randomly over time and may depend on environmental actions taken in the other region.

6 Numerical results

In this section we analyze the results for five different cases. In the first three cases players are identical. We contrast a base case in which parameters are described as in Table 2 with two other cases - one with a higher volatility and the other with an alternative damage function. We then consider two cases in which players are asymmetric - one in which players differ in terms of damage functions and in the other, players differ in terms of preferences for emissions reduction (i.e. green preferences).

6.1 Base case: identical players

We begin with consideration of the case in which players are identical, the willingness to pay for emissions reduction due to the green reward is zero, and the damage function is exponential (Equation (27)). Figure 4(a) contrasts total discounted expected utility versus temperature at $S = 800$ for the game compared to the social planner. The figures show $V_p(e_1, e_2, x, s, t)$ for $p = 1, 2$ in the game, the sum of players 1 and 2 utilities in the game, and the sum of utilities for the players as chosen by the social planner. Figure 4(b) zooms in on the leader and follower utilities in the game. We observe, as expected, that utility declines with temperature and total utility is greater in the social planner case. Under the symmetric game, the leader is better off than the follower. Recall that this is a repeated game which is played (i.e. optimal emissions chosen) every two years over the 150 year time line of the analysis. Since the leader is able to choose an emission level first, with knowledge of how the other player will react, this imparts some advantage to the leader depending on the values of the state variables. When the stock of carbon or temperature are very high,

the leader's advantage is diminished.

Figure 5 gives a different perspective on the game with plots of total discounted expected utility for the follower and leader versus temperature, with several different curves shown reflecting different levels of carbon stock, S . Utility is uniformly lower at higher carbon stocks.

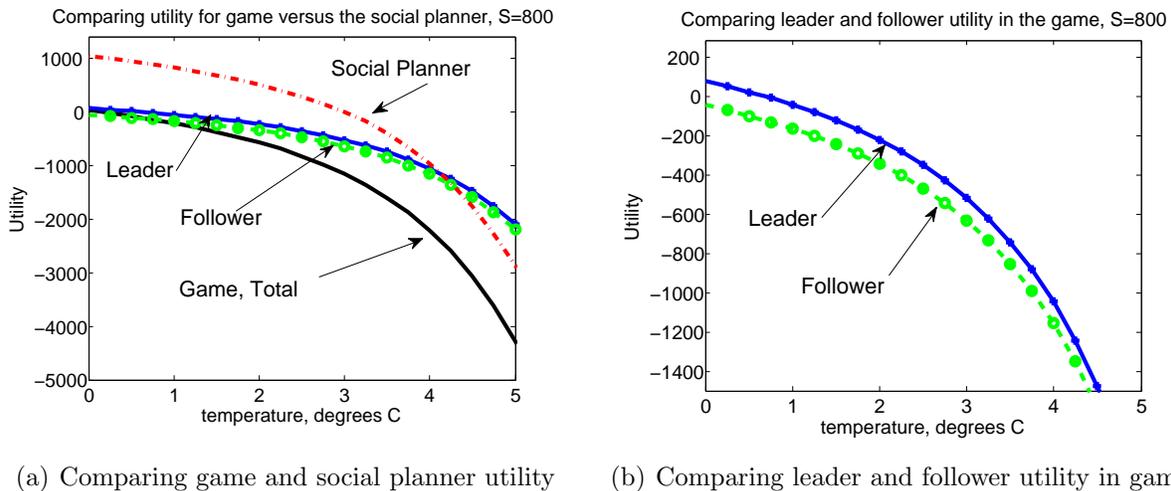


Figure 4: Base case utility for both players, comparing game and social planner. Exponential damage function. Stock of carbon at 800 GT.

Figure 6 depicts the optimal controls for the game and the social planner versus the stock of carbon, conditional on a temperature of 1.0 °C (close to the current value). For reference, recall that the stock of carbon in 2017 was about 870 GT. We see that total emissions fall as the stock of carbon is increased. In the game, both players choose emissions of 7 until the carbon stock reaches 1700 GT at which point the leader increases emissions to 10 while the follower cuts back to 3. By the time the stock reaches 1800 GT both leader and follower choose a lower emission level of 3. This reflects an advantage to the leader who can choose a higher emissions level, knowing that the follower will respond appropriately with emissions reductions which benefit both players. The social planner cuts back emissions more aggressively and at a lower carbon stock than the players in the game. At a stock of 1000 GT of carbon, the Planner chooses total emissions of 7 compared to 14 for the total emissions

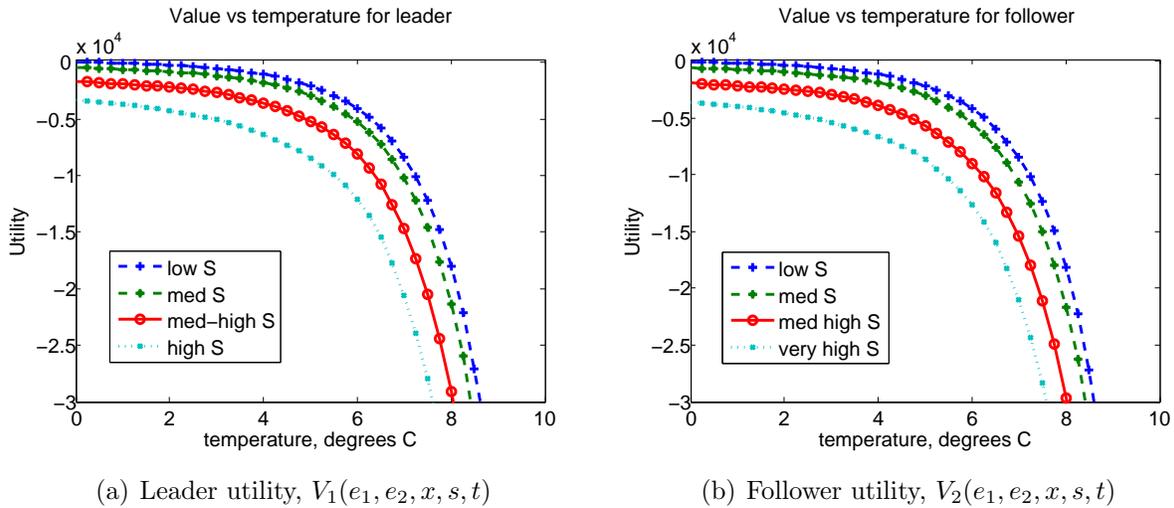
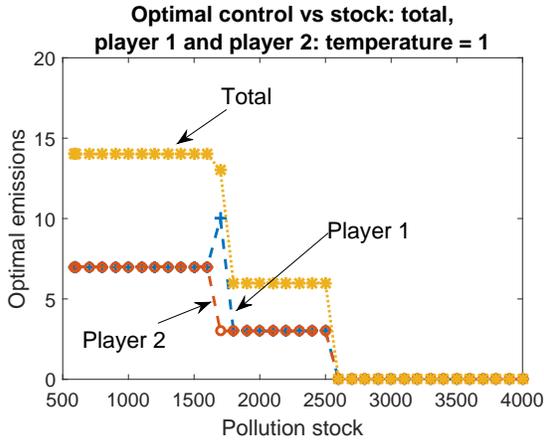


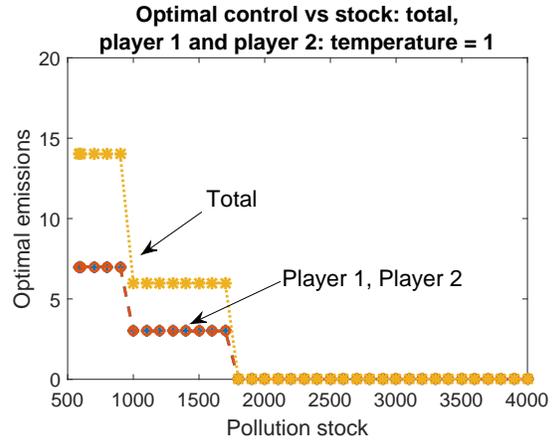
Figure 5: Values versus temperature for various carbon stock levels for the game, base case. Current emission levels at 3 GtC.

under the game. (Recall that the social planner chooses emissions to maximize total utility for the two players, which implies equalizing emissions since players are symmetric in the benefits received from emissions.) The model is describing a clear tragedy of the commons where strategic interactions of the two decision makers leave both worse off (in terms of utility) than when decisions are made by a central planner. In addition, for most values of state variables, emissions are higher under the game assumptions.

Figure 7 provides another perspective by contrasting the optimal controls for the game and the social planner versus temperature when the carbon stock is fixed at 800 GT. Here the difference between the game and social planner is even more stark. Under the game the players both choose to emit 7 GT per year (compared to a maximum possible of 10) even when the temperature reaches a very high level. The social planner is, in contrast, much more responsive to temperature and reduces emissions when temperature exceeds 2 °C.

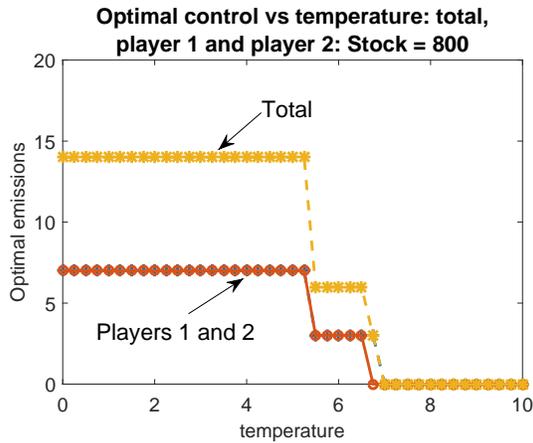


(a) Game

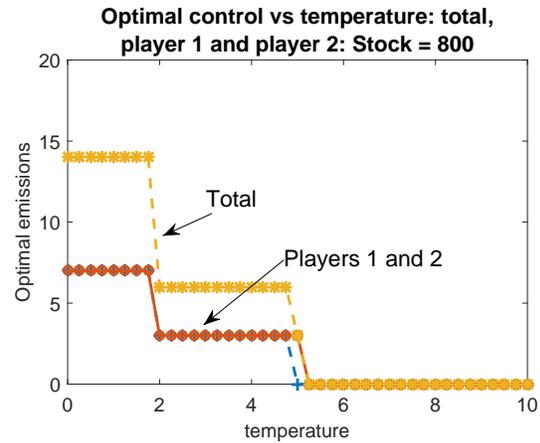


(b) Social planner

Figure 6: Optimal control versus pollution stock, contrasting game and social planner, base case, exponential damages, current temperature = 1 degrees C above preindustrial levels



(a) Game



(b) Social planner

Figure 7: Optimal control versus temperature, game (symmetric players) and social planner, exponential damage, carbon stock at 800 GT

6.2 Importance of volatility

As noted in Section 5.2, average global temperature exhibits significant volatility. We are interested in the effect on optimal emissions of an increase in volatility. Figure 8 compares

the optimal controls of the game versus the social planner when volatility is tripled to 0.3. For contrast, total emissions for the base case volatility ($\sigma = 0.1$) are also shown. Higher volatility results in earlier emissions reduction for both the game and the planner, but the impact on the planner is more marked. In the game, total emissions are the same in the high and low volatility cases until the carbon stock reaches 1200 GT. In contrast, the planner chooses lower emissions in the high volatility case over most levels of the carbon stock. The social planner chooses zero emission for carbon stock levels for 1300 Gt and beyond which is much reduced compared to the social planner choice under lower volatility, or in the game under either volatility scenario. In general at lower carbon stock levels between 800 to 1000 Gt, the increase in volatility causes the social planner to react to a greater extent than in the game.

We might consider the difference in utilities in the social planner case versus the game as another indication of the extent of the tragedy of the commons. For clarity, the difference $V^{planner}(\cdot) - V^{game}(\cdot)$ will be referred to as the social planner advantage (denoted SPA). At $S = 800$ Gt, the SPA in the high volatility case is higher than in the low volatility case, with the ratio $SPA^{\text{high vol}}/SPA^{\text{low vol}}$ ranging from 1.35 to 1.45 over temperature levels of interest. A plot of the planner and game utilities in the high and low volatility cases is shown for $S = 800$ in Figure 9. This ratio of SPAs is reduced as the carbon stock increases and disappears completely at very high carbon stock when no amount of cooperative action affect the bad outcomes.

These results suggest that high temperature volatility makes cooperative action as provided by a social planner even more important in order to avoid the potentially higher damages that can result from the high volatility. As carbon stocks increase this effect is reduced as the social planner has less leeway to avoid those damages.

6.3 Asymmetric damages

An important feature of global warming is the distribution of damages across nations, with some of the world's poorer regions suffering disproportionately. In this section we explore

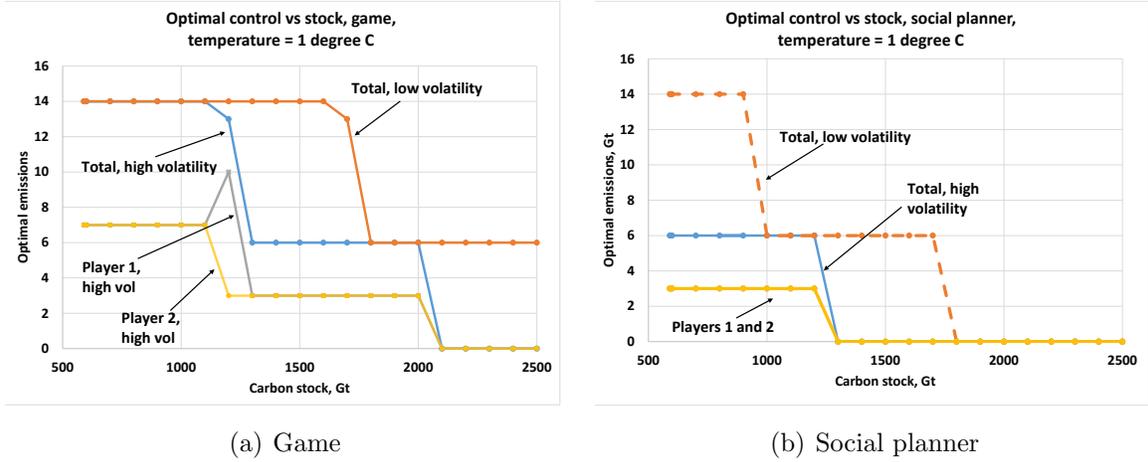


Figure 8: Optimal control versus carbon stock for high volatility ($\sigma = 0.3$) and low volatility ($\sigma = 0.1$) cases, game (symmetric players) and social planner, exponential damage.

the effect of asymmetric damages on strategic interactions by considering a case in which the follower has much higher sensitivity to increasing temperatures than the leader. Specifically, we compare a case with symmetric damage functions for the players ($\kappa_3 = 0.8$ in Equation (27)) with a case in which the follower's damages rise more quickly with temperature (κ_3 for player 2 is set at 1.1). We refer to $\kappa_3 = 1.1$ as high damage and $\kappa_3 = 0.8$ as low damage. We also contrast with the social planner case when damages are asymmetric. The optimal controls in these cases for various carbon stocks and at a temperature of 1°C are shown in Figure 10. In Figure 10(a) we observe that when the follower has high damages (dashed line) it starkly curtails emissions, choosing emissions of zero for all carbon stock levels above 600 Gt. In contrast in the symmetric case (dotted line), when the follower has low damages, it never reduces emission below 7 GtC. Figure 10(b) depicts the leader's optimal controls - the leader always has low damages, but responds differently depending on whether the follower has low or high damages. In particular, if the follower has low damages (dotted line), the leader reduces emissions from 10 Gt to 7 Gt when the carbon stock reaches 1000 Gt. However when the follower has high damages, the leader maintains emissions at 10 GtC until the carbon stock reaches 1600 Gt. The leader benefits from the fact that the follower has high damages, as the leader is able to maintain high emissions for longer. In Figure

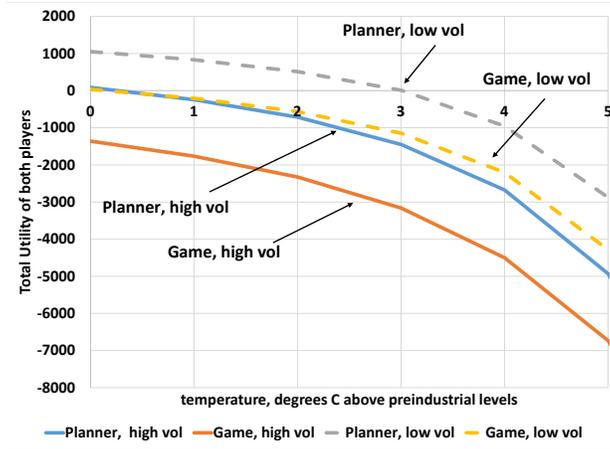


Figure 9: Comparing utilities for the game and the social planner for the high and low volatility cases. Stock of carbon at 800 Gt.

10(c), which shows total emissions, we observe that although the leader takes advantage of the higher costs of the follower, total emissions are still lower under asymmetric damages (dashed line) than under low symmetric damages (dotted line). Note that these results also hold when the leader has the higher damages. In this case (not shown) the follower takes advantage and increases its own emissions.

Figure 10(c) also shows the social planner choice for total emissions for the symmetric damages (both low damages) and asymmetric damages cases (solid lines). We see that when both players have low damages the social planner chooses fairly high emissions (14 GtC) over all carbon stock levels shown. In the asymmetric damages case, the social planner reduces emissions for carbon stocks levels above 800 Gt choosing lower total emissions than in the game. However, for lower carbon stock levels ($S < 800$) with asymmetric damages, the social planner chooses higher total emissions than in the game. This results from the fact that in the game, the leader is choosing its maximum possible emissions of 10 Gt while the follower chooses zero emissions over these carbon stock levels. The social planner chooses emissions of 14 Gt which are distributed equally between the players.

A comparison of utilities is shown in Figure 10(d). Not surprisingly, the relative difference of total utility for the social planner versus the game (the social planner advantage) is larger under asymmetric damages than under symmetric low damages

These results indicate the even greater need for cooperation in the case of asymmetric damages compared to symmetric damages. The social planner makes significantly different emissions choices than the outcome of the game in which the player with low damages takes advantage of the fact that the other player has relatively high damages. The social planner advantage is relatively larger in the case of asymmetric damages.

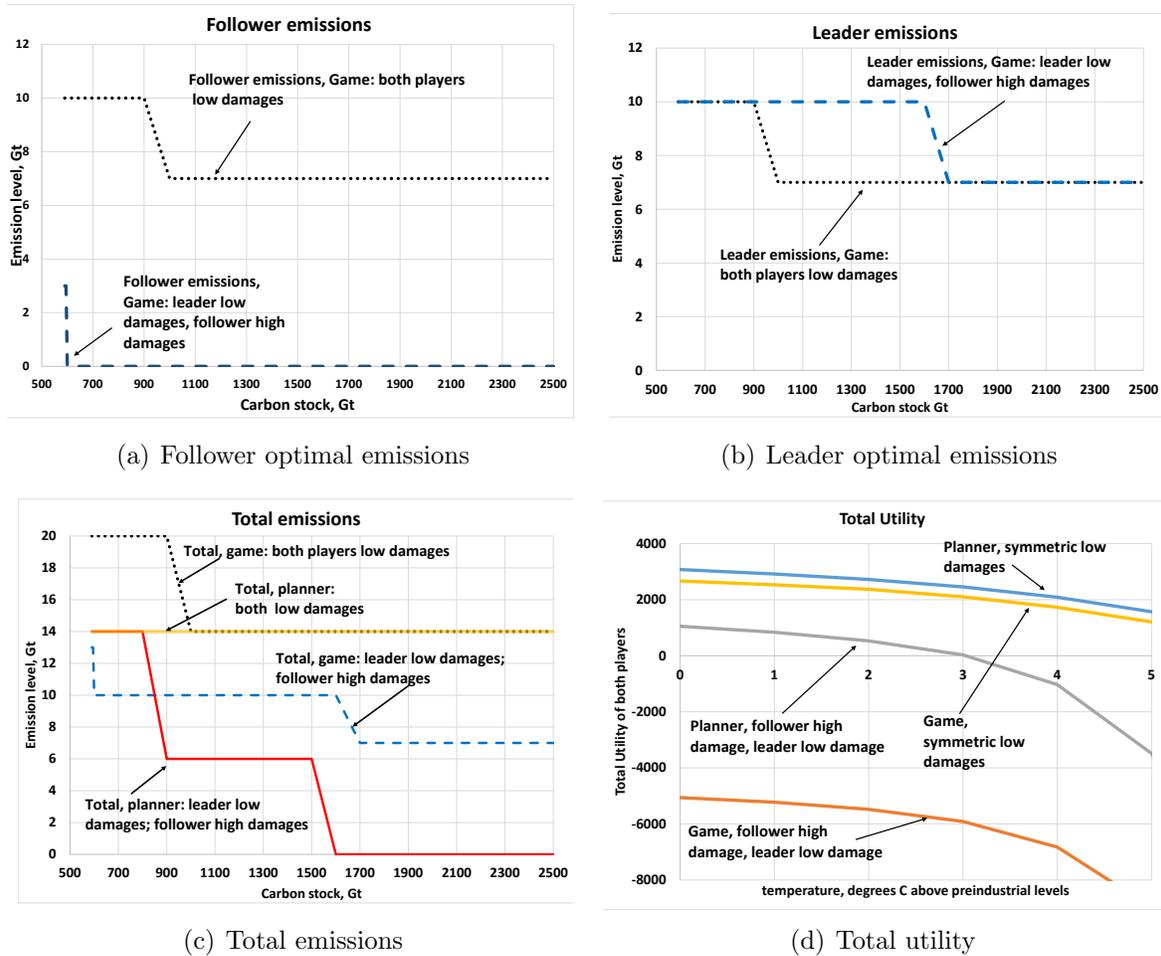


Figure 10: Contrasting symmetric and asymmetric damage functions, current temperature = 1 °C, Symmetric case: $\kappa_3 = 0.8$ for both players; Asymmetric case: $\kappa_3 = 0.8$ for player 1 and $\kappa_3 = 1.1$ for player 2, Optimal control versus pollution stock (Figures a,b,and c) and utility versus pollution stock (Figure d)

6.4 Alternate damage functions

Sensitivities were conducted using the alternate damage function given by Equation (28). Using a quadratic function, ($\kappa_2 = 2$), the optimal choice of emissions is the maximum possible (10 for each player) in both the game and the social planner. In contrast a cubic damage function, ($\kappa_2 = 3$) results in some curtailment of emissions, but emissions are still at higher levels than with exponential damages. We consider the exponential damage function to be the most reasonable as the damages quickly become very large at temperature above 3°C.

6.5 Asymmetric preferences

This section examines the impact of asymmetric preferences by considering a case in which one of the players gains a psychic benefit for reducing emissions relative to a given benchmark. We assume that the environmental friendly player is willing to pay 3 utility units, ($\theta_p = 3$) for reductions in emissions below the benchmark \bar{E} . The results are shown in Figure 11 which depicts optimal emissions choices for different carbon stock levels conditional on a temperature of 1 °C. We observe from Figure 11(b) that, as expected, when the leader has greener sentiments, it chooses a lower level of emissions than in the base case over nearly all carbon stock levels. In Figure 11(a) it is shown that when the leader has green sentiments, the follower maintains the same emissions (7 GtC) compared to the base case up to an atmospheric stock of 1600 GtC. Then for carbon stock levels ranging from 1700 to 1900, the follower increases its emissions relative to the base case, taking advantage of the reduced emissions of the leader. This may be characterized as a form of green paradox where increased green sentiments of one player causes the other player to free ride by increasing emissions.

In Figure 11(a), there is a contrasting effect observed at high carbon stock levels. The follower chooses zero emissions at 2400 GtC and beyond, which is at a lower stock level than in the base case. In the base case the follower chooses zero emissions for $S \geq 2600$ GtC. We call this the green bandwagon effect - the opposite of the green paradox effect noted above. An explanation is that at high carbon levels, the environmentally friendly policies of the leader make it worthwhile for the follower to also choose environmentally friendly policies,

because the follower knows this choice will help avert highly damaging consequences.

Figure 11(c) depicts total emissions for the base case and the green reward case. Over most values of carbon stock, total emissions are less under the green reward than in the base case. However, between 1800 and 1900 Gt of carbon, it may be observed that total emissions are higher under the green reward. The green sentiments of the leader have resulted in higher world emissions overall due to increased free riding of the follower.

Figure 11 shows results for the state variable temperature at 1°C above preindustrial levels. The green paradox effect and green band wagon effects are also observed for higher temperature levels of interest (up to 4°C). The green paradox is observed at medium high levels of carbon stock - 1200 to 1900 Gt depending on the temperature. The green bandwagon effect is observed at the upper end of carbon stock levels of interest, 2000 Gt and above. These effects are also observed when the follower has a green reward, and the leader responds via the green bandwagon or green paradox.

The impact of green sentiments in one player requires further investigation. In future work we will explicitly model the preferences of players and consider the possibility of players switching from brown to green preferences.

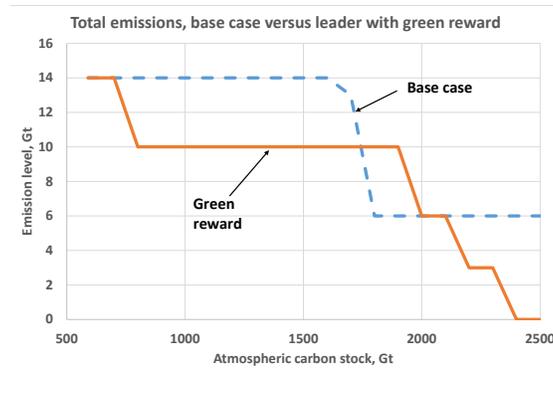
6.6 Checking for Nash equilibria

Recall that we are solving for a repeated series of Stackelberg games which happen every 2 years over the 150 year time span of the analysis. It is of interest to note whether Nash equilibria exist for these repeated games. In each of the cases described above, we check for the existence of Nash equilibria across all state variables and at each of the 75 decision times. We find that at each decision time about 65% of the nodes (representing carbon stock, temperature and emissions levels) satisfy the Nash equilibrium criterion. Further, we determine that that 25% of the Stackelberg equilibria are also Nash equilibria. See Section 4.2.2 for details.



(a) Follower emissions

(b) Leader emissions



(c) Total emissions

Figure 11: Optimal control versus pollution stock when leader receives a green reward for emissions reductions. Current temperature = 1 °C, Willingness to pay = 3 utility units.

7 Concluding comments

In this paper we have examined the strategic interactions of large regions making choices about greenhouse gas emissions in the face of rising global temperatures. We have modelled optimal decisions of players in a fully dynamic, closed loop Stackelberg game and have demonstrated its numerical solution. Our modelling of the evolution of carbon stock and temperature is based on Nordhaus's Integrated Assessment Model (Nordhaus 2013). We take into account the fundamental random nature of the temperature response to atmospheric carbon levels. In fact, our analysis shows that purely random effects are likely to cause global

temperature changes of $\pm(0.5 - 1.0)$ degrees regardless of emission levels.

Assumptions regarding damages from climate change are a highly controversial part of the climate economics literature. In our model, damages from rising temperature are subtracted from benefits derived from emissions, which implies total utility tends to a very large negative number as temperature rises. This is different from the DICE model and other similar models in which damages are specified as a percentage reduction in output. As noted by Nordhaus (2013), “this limits the usefulness of this approach for catastrophic climate change.” We adopt an exponential damage function rather than the commonly used quadratic damage function. We find the latter results in very little emissions curtailment even at very high temperatures and carbon stock levels. We argue that an exponential damage function is more appropriate to account for the disastrous effects of global average temperatures above 3 °C relative to pre-industrial levels.

Our results indicate a classic tragedy of the commons whereby regions acting in their own self interest in a non-cooperative game choose higher levels of emissions and have lower total utility than would be chosen by a social planner. In the paper we examined several factors which are likely to affect the strategic interactions of players and considered their effects on carbon emissions choices and utilities.

Volatility is found to have an important effect on optimal choices of players in the game as well as the social planner. An increase in volatility increases the likelihood of high temperatures and resulting high damages. This causes players in the game to choose lower levels of emissions as does the social planner. The difference in total utility between the social planner and the game (the social planner advantage) tends to be larger for higher volatility, implying that the tragedy of the commons is exacerbated by higher volatility, or, in other words, the need for cooperative action is increased. Although the drift in long run temperature is key in climate change policy, the impact of volatility on strategic interactions of decision makers is significant.

Asymmetric damages are also found to affect the outcome of the game. When one player experiences greater harm with rising temperatures, we find that the player with higher

damages is made worse off by the response of the other player. While the player with higher damages cuts back on their emissions more aggressively, the low damage player takes advantage of this by increasing their own own emissions relative to the symmetric damage case. The social planner advantage is higher in the case of asymmetric damages versus symmetric damages, as the socially planner optimally distributes emissions across the two players.

We also examined a case where one of the players receives a psychic benefit from emissions reductions compared to a benchmark, an effect we labelled the green reward. A green reward for one player causes that player to cut back their emissions from what would have otherwise been the case. There are various responses by the player with no green reward (the brown player) ranging from no response, to increasing or decreasing emissions depending on the values of the state variable (i.e. the carbon stock and the temperature). For more moderate values of temperature and carbon stock we observe either no response on the part of the brown player or an increase in emissions. The latter effect is similar to the one observed by Wirl (2011) in a deterministic game whereby the increase in green sentiments increases the free riding of brown players.

We also observed a contrary effect, which we call the green bandwagon effect, to contrast with the green paradox. For some high values of the carbon stock, the presence of a green reward for one player causes the brown player to reduce its own emissions (relative to the case with no green reward). Our interpretation is that at high carbon stocks where disaster is on the horizon, the brown player can be assured that the green player will cut back emissions, making it worthwhile for the brown player to also reduce emissions. The green preferences of the green player give the brown player more agency to effect a change in climate outcomes.

These results provide some novel insight into regions' strategic behaviour regarding emissions choices. In future work we hope to explore the development of green preferences, and in particular how preferences of one player might be linked with preferences of the other player.

Appendices

A Numerical methods

A.1 Advancing the solution from $t_{m+1}^- \rightarrow t_m^+$

Since we solve the PDEs backwards in time, it is convenient to define $\tau = T - t$ and use the definition

$$\begin{aligned}\hat{V}_p(e_1, e_2, x_i, s, \tau) &= V_p(e_1, e_2, x_i, s, T - \tau) \\ \hat{\pi}_p(e_1, e_2, x_i, s, \tau) &= \pi_p(e_1, e_2, x_i, s, T - \tau).\end{aligned}\quad (31)$$

We rewrite Equation ((11)) in terms of backwards time $\tau = T - t$

$$\begin{aligned}\frac{\partial \hat{V}_p}{\partial \tau} &= \hat{\mathcal{L}}\hat{V}_p + \hat{\pi}_p + [(e_1 + e_2) + \rho(\bar{S} - s)]\frac{\partial \hat{V}_p}{\partial s} \\ \hat{\mathcal{L}}\hat{V}_p &\equiv \frac{(\sigma)^2}{2}\frac{\partial^2 \hat{V}_p}{\partial x^2} + \eta(\bar{X} - x)\frac{\partial \hat{V}_p}{\partial x} - r\hat{V}_p.\end{aligned}\quad (32)$$

Defining the Lagrangian derivative

$$\frac{D\hat{V}_p}{D\tau} \equiv \frac{\partial \hat{V}_p}{\partial \tau} + \left(\frac{ds}{d\tau}\right)\frac{\partial \hat{V}_p}{\partial s}, \quad (33)$$

then Equation (32) becomes

$$\frac{D\hat{V}_p}{D\tau} = \hat{\mathcal{L}}\hat{V}_p + \pi_p \quad (34)$$

$$\frac{ds}{d\tau} = -[(e_1 + e_2) + \rho(\bar{S} - s)]. \quad (35)$$

Integrating Equation (35) from τ to $\tau - \Delta\tau$ gives

$$s_{\tau-\Delta\tau} = s_\tau \exp(-\rho\Delta\tau) + \bar{S}(1 - \exp(-\rho\Delta\tau)) + \left(\frac{e_1 + e_2}{\rho}\right)(1 - \exp(-\rho\Delta\tau)). \quad (36)$$

We now use a semi-Lagrangian timestepping method to discretize Equation (32) in backwards time τ . We use a fully implicit method as described in Chen & Forsyth (2007).

$$\begin{aligned} \hat{V}_p(e_1, e_2, x, s_\tau, \tau) &= (\Delta\tau)\hat{\mathcal{L}}\hat{V}_p(e_1, e_2, x, s_\tau, \tau) \\ &\quad + (\Delta\tau)\pi_p(e_1, e_2, x, s_\tau, \tau) + \hat{V}_p(e_1, e_2, x, s_{\tau-\Delta\tau}, \tau - \Delta\tau) . \end{aligned} \quad (37)$$

Equation (37) now represents a set of decoupled one-dimensional PDEs in the variable x , with (e_1, e_2, s) as parameters. We use a finite difference method with forward, backward, central differencing to discretize the $\hat{\mathcal{L}}$ operator, to ensure a positive coefficient method. (See Forsyth & Labahn (2007/2008) for details.) Linear interpolation is used to determine $\hat{V}_p(e_1, e_2, x, s_{\tau-\Delta\tau}, \tau - \Delta\tau)$. We discretize in the x direction using an unequally spaced grid with n_x nodes and in the S direction using n_s nodes. Between the time interval t_{m+1}^-, t_m^+ we use n_τ equally spaced time steps. We use a coarse grid with $(n_\tau, n_x, n_s) = (2, 27, 21)$. We repeated the computations with a fine grid doubling the number of nodes in each direction to verify that the results are sufficiently accurate for our purposes.

A.2 Advancing the solution from $t_m^+ \rightarrow t_m^-$

We model the possible emission levels as four discrete states for each of e_1, e_2 , which gives 16 possible combinations of (e_1, e_2) . We then determine the optimal controls using the methods described in Section 4.2.1. We use exhaustive search (among the finite number of possible states for (e_1, e_2)) to determine the optimal policies. This is, of course, guaranteed to obtain the optimal solution.

References

- Ackerman, F., Stanton, E. A. & Bueno, R. (2013), ‘Epstein-Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy?’, *Environmental and Resource Economics* **56**(1), 73–84.
- Amarala, S. (2015), Monotone numerical methods for nonlinear systems and second order partial differential equations, PhD thesis, University of Waterloo, Waterloo, Ontario, Canada.
- Barcena-Ruiz, J. C. (2006), ‘Environmental taxes and first-mover advantages’, *Environmental and Resource Economics* **35**, 19–39.
- Bednar-Friedl, B. (2012), ‘Climate policy targets in emerging and industrialized economies: the influence of technological differences, environmental preferences and propensity to save’, *Empirica* **39**, 191–215.
- Bressan, A. (2011), ‘Noncooperative Differential Games’, *Milan Journal of Mathematics* **79**(2), 357–427.
- Bressan, A. & Shen, W. (2004), ‘Semi-cooperative strategies for differential games’, *International Journal of Game Theory* **32**(4), 561–593.
- Cacace, S., Cristiani, E. & Falcone, M. (2013), Numerical approximation of Nash equilibria for a class of non-cooperative differential games, in L. Petrosjan & V. Mazalov, eds, ‘Game Theory and Applications’, Vol. 16, Nova Science Publishers.
- 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.
- Chesney, M., Lasserre, P. & Troja, B. (2017), ‘Mitigating global warming: a real options approach’, *Annals of operations research* **255**(1-2), 465–506.
- Clean Energy Canada (2015), How to adopt a winning carbon price. Initiative of the Centre for Dialogue, Simon Fraser University, Vancouver Canada; Retrieved November

- 27, 2015 at <http://cleanenergycanada.org/wp-content/uploads/2015/02/Clean-Energy-Canada-How-to-Adopt-a-Winning-Carbon-Price-2015.pdf>.
- Crost, B. & Traeger, C. P. (2014), ‘Optimal CO2 mitigation under damage risk valuation’, *Nature Climate Change* **4**(7), 631–636.
- Dixit, A. & Pindyck, R. (1994), *Investment Under Uncertainty*, Princeton University Press.
- Dockner, E. J., Jorgensen, S., Long, N. V. & Sorger, G. (2000), *Differential games in economics and management science*, Cambridge University Press.
- Dockner, E. J. & Long, N. V. (1993), ‘International pollution control: Cooperative versus noncooperative strategies’, *Journal of Environmental Economics and Management* **25**, 13–29.
- Dockner, E., VanLong, N. & Sorger, G. (1996), ‘Analysis of Nash equilibria in a class of capital accumulation games’, *Journal of Economic Dynamics & Control* **20**(6-7), 1209–1235.
- Forsyth, P. & Labahn, G. (2007/2008), ‘Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance’, *Journal of Computational Finance* **11**(2), 1–44.
- Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2014), ‘Optimal taxes on fossil fuel in general equilibrium’, *Econometrica* **82**(1), 41–88.
- Hambel, C., Kraft, H. & Schwartz, E. (2017), Optimal carbon abatement in a stochastic equilibrium model with climate change. NBER Working Paper Series.
- Harris, C., Howison, S. & Sircar, R. (2010), ‘Games with exhaustible resources’, *SIAM Journal of Applied Mathematics* **70**(7), 2556–2581.
- Insley, M. & Forsyth, P. (2018), Climate games: Who’s on first? What’s on second? Working Paper.

- Kelly, D. & Kolstad, C. (1999), ‘Bayesian learning, growth, and pollution’, *Journal of Economic Dynamics & Control* **23**(4), 491–518.
- Kossey, A., Peszko, G., Oppermann, K., Prytz, N., Klein, N., Blok, K., Lam, L., Wong, L. & Borkent, B. (2015), ‘State and trends of carbon pricing 2015’. retrieved from <http://documents.worldbank.org/curated/en/636161467995665933/State-and-trends-of-carbon-pricing-2015>.
- Leach, A. (2007), ‘The climate change learning curve’, *Journal of Economic Dynamics and Control* **31**, 1728–1752.
- Ledvina, A. & Sircar, R. (2011), ‘Dynamic Bertrand oligopoly’, *Applied Mathematics and Optimization* **63**(1), 11–44.
- Lemoine, D. & Traeger, C. (2014), ‘Watch your step: optimal policy in a tipping climate’, *American Economic Journal: Economic Policy* **6**(2), 137–166.
- List, J. A. & Mason, C. F. (2001), ‘Optimal institutional arrangements for transboundary pollutants in a second-best world: Evidence from a differential game with asymmetric players’, *Journal of Environmental Economics and Management* **42**, 277–296.
- Long, N. V. (2010), *A Survey of Dynamic Games in Economics*, World Scientific Publishing Company.
- Ludkovski, M. & Sircar, R. (2012), ‘Exploration and exhaustibility in dynamic Cournot games’, *European Journal of Applied Mathematics* **23**(3), 343–372.
- Ludkovski, M. & Sircar, R. (2015), Game theoretic models for energy production, *in* R. A’id, M. Ludkovski & R. Sircar, eds, ‘Commodities, Energy and Environmental Finance’, Springer, Berlin.
- Ludkovski, M. & Yang, X. (2015), Dynamic cournot models for production of exhaustible commodities under stochastic demand, *in* R. A’id, M. Ludkovski & R. Sircar, eds, ‘Commodities, Energy and Environmental Finance’, Springer.

- Nkuiya, B. (2015), ‘Transboundary pollution game with potential shift in damages’, *Journal of Environmental Economics and Management* **72**, 1–14.
- Nordhaus, W. (2013), Integrated economic and climate modeling, *in* P. B. Dixon & D. W. Jorgenson, eds, ‘Handbook of Computable General Equilibrium Modeling, First Edition’, Vol. 1, Elsevier, chapter 16, pp. 1069–1131.
- Nordhaus, W. & Sztorc, P. (2013), Dice 2013r: Introduction and user’s manual, Technical report.
- Pindyck, R. S. (2013), ‘Climate change policy: What do the models tell us?’, *Journal of Economic Literature* **51**, 860–872.
- Salo, S. & Tahvonen, O. (2001), ‘Oligopoly equilibria in nonrenewable resource markets’, *Journal of Economic Dynamics & Control* **25**(5), 671–702.
- Traeger, C. (2014), ‘A 4-stated dice: Quantitatively addressing uncertainty effects in climate change’, *Environmental and Resource Economics* **59**(2), 1–37.
- Urpelainen, J. (2009), ‘Explaining the Schwarzenegger phenomenon: Local frontrunners in climate policy’, *Global Environmental Politics* **9**, 82–105.
- van der Ploeg, F. (1987), ‘Inefficiency of credible strategies in oligopolistic resource markets with uncertainty’, *Journal of Economic Dynamics & Control* **11**(1), 123–145.
- Weitzman, M. L. (2012), ‘GHG targets as insurance against catastrophic climate damages’, *Journal of Public Economic Theory* **14**, 221–244.
- Williams, R. C. (2012), ‘Growing state-federal conflicts in environmental policy: The role of market-based regulation’, *Journal of Public Economics* **96**, 1092–1099.
- Wirl, F. (2008), ‘Tragedy of the commons in a stochastic game of a stock externality’, *Journal of Public Economic Theory* **10**(1), 99–124.

Wirl, F. (2011), 'Global Warming with Green and Brown Consumers', *Scandinavian Journal of Economics* **113**(4, SI), 866–884.

Xepapadeas, A. (1998), 'Policy adoption rules and global warming - theoretical and empirical considerations', *Environmental & Resource Economics* **11**(3-4), 635–646. 1st World Congress of Environmental and Resource Economists, Venice, Italy, June 25-27, 1998.

Zagonari, F. (1998), 'International pollution problems: Unilateral initiatives by environmental groups in one country', *Journal of Environmental Economics and Management* **36**(1), 46–69.