

OPTIMAL POLICY FOR CARBON PRICING

*Challenging the Hotelling rule
and dissecting mitigation cost uncertainties*

KAJ-IVAR VAN DER WIJST

*Master thesis in Mathematical Sciences
Utrecht University
& PBL Netherlands Environmental Assessment Agency*

October 28, 2018

Under supervision of:

Prof. Jason Frank
UU

Dr. Andries Hof
PBL

Prof. Detlef van Vuuren
PBL



Utrecht University



Planbureau voor de Leefomgeving

Acknowledgements

This thesis is performed as an internship at PBL Netherlands Environmental Assessment Agency, in Den Haag, from February 2018 to October 2018. I would like to begin by thanking my three supervisors, prof. Jason Frank from Utrecht University, dr. Andries Hof and prof. Detlef van Vuuren, both from PBL, for their insightful comments, interesting discussions and adept supervision. I would also like to thank my colleagues at PBL for the knowledge sharing and open discussions. Next, I would like to thank the Centre for Complex Systems Studies for allowing me to spend one month in Milan for a collaboration with researchers at the European Institute on Economics and the Environment. Last, but certainly not least, I would like to thank my friends and family for their continuous support and for reminding me that life besides the thesis is also important.

Contents

1	General introduction	1
I	Optimal carbon pricing	3
2	Introduction part I	4
3	Mathematical prelude	6
3.1	Optimal Control Theory	6
3.1.1	Pontryagin's Principle	7
3.1.2	Variational approach	8
3.1.3	Bellman Equation	9
3.1.4	Hamilton-Jacobi-Bellman Equation	10
4	Hotelling and Climate Policy	12
4.1	The Hotelling Rule	12
4.2	Towards Climate Policy	13
5	The Stylised Policy Setting	15
5.1	Carbon prices and carbon taxes	15
5.2	MAC Curves	15
5.3	Dynamics	17
5.4	Costs	18
5.4.1	Economic interpretation	18
5.4.2	Example price path	19
6	Scenarios and functional forms	20
6.1	Reference	20
6.2	Learning over Time	21
6.3	Learning by Doing	21
6.4	Inertia	22
6.5	Minimum emission level	22
6.6	Combination	22
6.7	Discontinuous or convex MAC	22
7	Solving the optimal control problem	24
7.1	Analytical solution	24
7.2	Numerically	25
7.2.1	Discretising the system	25
7.2.2	Formulating the Bellman Equation	26
7.2.3	Implementing the inertia constraint	26
7.3	Computational considerations	27

7.3.1	Discretisation of objective function	27
7.3.2	Interpolation	28
7.3.3	Chebyshev interpolation	29
7.4	Calculating the Proxy Hotelling paths	32
7.5	GPU computation	32
8	Results — Optimal carbon price paths	34
8.1	Analytical results	34
8.1.1	Reference	34
8.1.2	Learning over Time	35
8.1.3	Learning by Doing	36
8.2	Numerical	36
8.2.1	Numerical confirmation of Hotelling	37
8.2.2	Scenario: Learning by Doing	37
8.2.3	Scenario: Minimum emission level	37
8.2.4	Scenario: Inertia	40
8.2.5	Scenario: Combining learning and constraints	40
8.2.6	Different MAC shapes	40
8.2.7	Costs	41
8.3	Comparison to Goulder et al	42
8.3.1	General setting	43
8.3.2	Solving the optimal control problem	43
8.3.3	Optimal carbon tax	44
8.3.4	Shadow costs or carbon tax	45
8.4	Implications and impact	45
9	Application of stylised setting: BECCS to the Future	46
9.1	Motivation	46
9.2	Method	47
9.2.1	Analytical setting	47
9.2.2	Finding analytical expressions of indicators	47
9.2.3	Convexity of the MAC	48
9.2.4	WITCH for technology specification	49
9.3	Results	50
9.3.1	Key indicators	50
9.3.2	Intergenerational equity	50
9.3.3	Comparative statics	52
9.4	Different forms of baseline and discounting	53
9.4.1	Quadratic baseline	54
9.4.2	Hyperbolic discount rate	54
9.5	Implications and impact	57
10	Wrap-up of part I	58
10.1	Conclusion	58
10.2	Discussion	59
II	Sources of uncertainty in the costs of meeting climate targets	60
11	Introduction part II	61

12 Mathematical prelude	64
12.1 Variance decomposition	64
12.1.1 The Sobol method	64
12.1.2 Towards an expression for the variance	65
12.2 Monte-Carlo Simulation method	66
12.3 Statistical considerations	67
12.3.1 The β -PERT distribution	68
12.3.2 Sampling from an arbitrary distribution	68
12.3.3 Percentiles from distributions	69
13 The model	70
13.1 Climate sensitivity	70
13.1.1 Parameter distribution	70
13.1.2 Fitting the distributions to the plume	71
13.2 Non CO ₂ contributions	74
13.3 Current global temperature	76
13.4 Costs	76
13.4.1 Comparing the three cost measures	77
13.4.2 Towards a continuous description	77
13.4.3 Calculating the statistical distribution	78
13.5 Putting it all together	79
14 Results — Uncertainty in mitigation costs	81
14.1 Model output	81
14.1.1 Carbon budgets	81
14.2 Sensitivity analysis	84
15 Wrap-up of part II	86
15.1 Conclusion	86
15.2 Discussion	87
A Derivations	91
A.1 Analytical results	91
A.1.1 Less than exponential growth of carbon tax in LBD case	91
A.1.2 Closed form statics for convex MAC and concave baseline	92
A.2 Sobol’s Variance-based Sensitivity Analysis	93
A.2.1 Zero-mean condition for decomposition in summands of different dimensions	93
A.2.2 From function decomposition to variance decomposition	94
B Additional figures	95
B.1 Optimal carbon tax in log-plot	96
B.2 Optimal carbon tax under extreme Learning by Doing	97
B.3 Cost reduction Learning by Doing using Convex MAC	98
B.4 Sign of derivatives of net-zero year and overshoot	99

Chapter 1

General introduction

The main theme of this thesis is applying mathematics to climate policy research. More specifically, this thesis encompasses two main distinct parts. In this chapter, we give a short introduction of both parts, and a description of the structure of the thesis. More detailed background information and context is given at the beginning of each respective part.

BRIEF GOALS OF THIS THESIS

Extreme weather, severe droughts, increased famine and more: climate change is affecting our daily lives, and the societal impact will strongly increase in the future unless drastic measures are taken. Since climate change is mainly driven by human emitted CO₂ gases, many economists agree that a *carbon price* should be implemented, effectively making CO₂ emissions more expensive.

Over the recent years, Integrated Assessment Models (IAMs), which aim to explore possibilities of curbing climate change effects by combining geophysical models with socio-economic models, have investigated many scenarios for reducing greenhouse gas emissions over the coming century. When calculating the most cost-effective ways to achieve a given emission reduction, the resulting carbon price paths have many different shapes, as shown on the cover of this thesis.

However, the Hotelling rule — a famous theory in economics, developed in 1931, treating the optimal management of a non-renewable natural resource — suggests that the most cost-effective carbon price path should follow a fixed exponential growth path, contradicting the results from some of the obtained IAM scenarios. In this thesis, we investigate under which conditions the optimal carbon price path follows indeed an exponential path, as predicted by the Hotelling rule, and when this path differs.

The second part of this thesis covers a different topic of climate policy: comparing the uncertainty resulting from geophysical models to the uncertainty from socio-economic models giving mitigation costs. The former models mainly answer the question what the change in global mean temperature will be as function of CO₂ emissions, while the latter models are used to assess the costs necessary to limit global CO₂ emissions to a chosen carbon budget.

Using available data from literature, we assess the uncertainties in both types of models, and combine the models by expressing the mitigation costs as function of temperature change. By calculating the contribution of the geophysical and socio-economic parts to the total uncertainty in the combined model, we investigate which field of research would benefit most from a reduced uncertainty.

The mathematical tools employed in the first part of the thesis consist mainly of techniques from Optimal Control Theory: Pontryagin's Principle for solving continuous optimal control problems, and the Bellman Equation for discretised problems. The second part of the thesis relies mostly on a combination of the fields of stochastics, Monte Carlo simulations and statistics.

INVOLVED INSTITUTIONS

This thesis is a collaboration between three institutes: Utrecht University, the Netherlands Environmental Assessment Agency (PBL) in Den Haag and the European Institute on Economics and the Environment (EIEE) in Milan.



Utrecht University

Utrecht University — During my thesis, I have primarily worked with two institutes of the University: the Mathematical Institute, home to my master programme, and the Centre for Complex Systems Studies (CCSS), an interdisciplinary group of researchers focused on solving societal problems using a complex systems approach.



Planbureau voor de Leefomgeving

PBL Netherlands Environmental Assessment Agency — PBL is the national institute for strategic policy analysis on the environment, nature and spatial planning. I performed my thesis as an internship at the international climate scenario group at the department of Climate, Air and Energy at PBL.



European Institute on Economics and the Environment (EIEE) — This newly formed research institute, based in Milan, consists of researchers previously member of the Fondazione Eni Enrico Mattei (FEEM) and the Centro Euro-Mediterraneo sui Cambiamenti Climatici (CMCC). EIEE is focused on delivering research and policy advice on environmental, energy and natural resources topics. Thanks to the CCSS travel fund, I spent one month collaborating with researchers of this institute.

Part I

Optimal carbon pricing

Chapter 2

Introduction

— *Carbon prices and optimal control*

CONTEXT

At the time of writing, prof. William D. Nordhaus had just been awarded the Nobel Prize in Economics, “*for integrating climate change into long-run macroeconomic analysis*”. The prize is a recognition for his innovative work of approaching climate change from an economic point of view, mainly by introducing the concept of a global carbon price and creating the first Integrated Assessment Model (IAM) in the 1990’s.

Many of today’s state-of-the-art IAMs use a global carbon price to assess cost-optimal ways to meet carbon emission reduction goals. This carbon tax is typically not static: a full time profile is calculated to properly assess cost-effective mitigation scenarios. Example carbon tax paths resulting from different IAMs are shown in Fig. 2.1. At first glance, all tax paths seem different. However, upon closer inspection, they can be categorised into two groups: (almost) exponentially increasing paths and paths with a different shape.

The theoretical justification for exponentially increasing carbon price paths dates back to 1931, when Harold Hotelling developed a theory on optimal management of non-renewable (also called exhaustible) natural resources (Hotelling, 1931). According to the Hotelling Rule, an owner of such a resource — like a gold mine or an oil well — should sell the resource at an exponentially increasing price in order to maximise his profits. The growth rate of the price path should equal the interest rate. By considering the climate system as a similar exhaustible resource, where the amount of greenhouse gas emissions are limited by an imposed carbon budget (Grimaud and Rougé, 2005), the Hotelling Rule would imply an exponentially increasing carbon price path, necessary to minimise the total mitigation costs.

For this reason, some IAMs got criticised for obtaining carbon price paths that significantly deviated from exponential growth. Moreover, for simple settings, the Hotelling Rule has been shown to be optimal. For instance, when considering learning by doing — where abatement costs decrease as a function of cumulative abatement, or previous effort — Goulder and Mathai (2000) argue that the Hotelling Rule is optimal.

However, limitations on the availability of abatement technologies and political and social challenges create a more complicated setting, in which the validity of the Hotelling Rule appears unclear. In this thesis, we show that already under a number of simple assumptions, the optimal carbon price deviates from the exponential path. In fact, contrary to the findings of Goulder, we prove that even under a learning-by-doing assumption the optimal price path does not follow the Hotelling path. The same holds when imposing constraints on reduction speed (inertia) or when putting a minimum on the emission level.

BOUNDARIES OF THIS THESIS

The goal of this part of the thesis is to investigate under which conditions the Hotelling rule is not optimal anymore. We therefore do not show that certain carbon price paths from Fig. 2.1 are correct, or better than others. Instead, we limit ourselves to show that the price paths deviating from the exponentially increasing paths are not necessarily wrong.




Figure 2.1: Carbon price paths resulting from a cost-effectiveness setting using different Integrated Assessment Models, in the Shared Socio-economic Pathways database (SSP2, RCP2.6 carbon prices), normalised on the initial carbon price. The MESSAGE and MESSAGE-GLOBIUM paths overlap.

To explore the limits of the Hotelling rule, we build a model that attempts to capture the most important features of climate policy. We do not claim that our model is a full Integrated Assessment Model: it is, however, an easy to interpret simplified model, with useful applications. For instance, we use it to calculate the sensitivity of key climate policy indicators to changes in discount rate. The simplicity of the model allows for intuitive explanations of the observed results.

From the mathematical perspective, the methods we propose to calculate optimal carbon paths are especially useful for simple, yet computationally hard, optimal control problems. We take full advantage of modern technology to efficiently obtain the optimal results necessary to challenge the Hotelling rule. While our numerical method could be extended to systems with higher dimensionality and complexity, the computational requirements would quickly become too expensive. However, recent developments in Machine Learning techniques already provide alternatives to overcome the curse of dimensionality.

READING GUIDE

Before introducing the simplified model we use to calculate optimal carbon prices, we start with an overview and description of the main mathematical techniques used throughout this part of the thesis in chapter 3. A reader less interested in the mathematical background can safely skip this chapter.

The mathematical prelude is followed by an introduction of the Hotelling Rule in chapter 4, in which we also describe the connection with climate policy in more detail. Next, the general model is presented in chapter 5, with an overview of the chosen scenarios and assumptions in chapter 6.

In chapter 7, we describe how we apply the Optimal Control techniques from chapter 3 to obtain optimal carbon price paths, and how this is implemented computationally. This part can also safely be skipped by a reader who is primarily interested in the results, and less in the methods.

The results, discussed in chapter 8, are split in three parts: first, we provide general analytical results on the shape of the carbon price path under for few simple scenarios. Second, all scenarios are solved using the numerical technique. Finally, in the third part, we provide an in-depth analysis of why our results deviate from the results in [Goulder and Mathai \(2000\)](#) for the learning by doing scenario.

The last chapter of part I, chapter 9, concerns the work performed in collaboration with the European Institute on Economics and the Environment (EIEE) in Milan: analysing the importance of choosing the discount factor in optimal climate policy.

Chapter 3

Mathematical prelude

To determine the optimal tax path for carbon emissions, we first need to describe the mathematical framework we will be using throughout the first part of this thesis. This is mainly focused around the field of Optimal Control Theory. First, we look at methods to solve problems set in a continuous time setting, from both a Hamiltonian perspective as well as from a Lagrangian perspective. Second, we introduce the equations used to solve discrete time problems.

3.1. OPTIMAL CONTROL THEORY

An optimal control problem always consists of a dynamical system defining the evolution of a *state variable* x , which can be influenced by a *control variable* u . Traditionally, the letter u is used, since it is the first letter of the Russian word “управление”, which means “control”. Mathematically, the dynamics of the system are described as:

$$\dot{x} = f(t, x, u), \tag{3.1}$$

where x and u are vectors in \mathbb{R}^m and \mathbb{R}^n respectively, for a given function $f : \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$. For the purpose of this thesis, we assume the system is deterministic, and therefore we do not worry about the field of Stochastic Optimal Control.

The next important ingredient in an optimal control problem is the associated cost. This can depend on the control variable (giving more gas during a rocket propulsion is more expensive), on the state variable (being far from a preset path for a rocket can also be expensive), on the time (the longer the rocket is away the more expensive) or on a combination of them. We will call these costs, for a given time $t \in [0, T]$, $R(t, x, u)$. Since we will work in an economic setting, these costs are discounted with a discount rate r . There can also be a cost $\Psi(x(T))$ depending on the end state only (for example, how far the rocket lands from its intended location).

The goal of optimal control is to find, for every time step, a value for the control variable $u(t)$ such that it minimises the total costs:

$$\underset{u(t)}{\text{minimise}} \int_0^T e^{-rt} R(t, x, u) dt + \Psi(x(T)), \quad \text{subj. to } \dot{x} = f(t, x, u) \tag{3.2}$$

Similarly, we can perform a maximisation problem where instead of minimising the costs, we maximise the reward at every time step.

While there are many techniques to solve or approximate solutions to these problems, we will focus on two of them: in section 3.1.1, we introduce Pontryagin’s Principle, which yields necessary conditions for optimality (and therefore a local optimum), followed by the Bellman Equation in section 3.1.3, which yields a necessary and sufficient condition (and therefore a global optimum). The latter solution, however, is only defined for a system which is discrete in time. We finish in section 3.1.4 with a brief discussion of the extension of the Bellman equation to continuous time, which leads to the Hamilton-Jacobi-Bellman differential equation.

3.1.1. PONTRYAGIN'S PRINCIPLE

As introduced above, the goal of an optimal control problem is to maximise (or minimise) a payoff functional $P[u]$:

$$\begin{aligned} P[u, x] &= \int_0^T e^{-rt} R(t, x(t), u(t)) dt + \Psi(x(T)), \\ &= \int_0^T q(t, x(t), u(t)) dt + \Psi(x(T)) \end{aligned}$$

where the dynamics of the systems are determined by $\dot{x}(t) = f(t, x(t), u(t))$ and $q(\cdot)$ represents the discounted payoff (or costs). Note that we consider an m -dimensional state variable $x \in \mathbb{R}^m$ and an n -dimensional control variable u taking values in the set of admissible controls $\mathcal{A} \subset \mathbb{R}^n$. Pontryagin's Principle, which we state below, requires the definition of a so-called *Hamiltonian*, familiar from applications in Classical Mechanics. Note that for notational reasons, we disregard the explicit time dependence in the function f and q . The Hamiltonian is then defined as:

$$\mathcal{H}(x, u, \lambda) = q(x, u) + \lambda \cdot f(x, u),$$

where we introduce the *costate variable* $\lambda(t) \in \mathbb{R}^m$, an extension of the Lagrange Multiplier in a static optimisation problem. With this, we can proceed to the statement of Pontryagin's Principle, as described in [Knowles \(1981\)](#).

Theorem 1 (Pontryagin's Principle). *In the optimal control problem defined in Eq. (3.2), suppose u^* is the optimal control with corresponding trajectory x^* that maximises the payoff functional $P[u]$. Then there exists a function $\lambda^*(t) \in \mathbb{R}^m$ such that the following conditions are satisfied:*

- (a) $\dot{x}^* = \nabla_{\lambda} \mathcal{H}(x^*, u^*, \lambda^*)$
- (b) $\dot{\lambda}^* = -\nabla_x \mathcal{H}(x^*, u^*, \lambda^*)$
- (c) $\mathcal{H}(x^*(t), u^*(t), \lambda^*(t)) = \max_{v \in \mathcal{A}} \mathcal{H}(x^*(t), v, \lambda^*(t))$

We refer the user interested in the formal proof and derivation of this principle to [Fleming and Rishel \(1975\)](#) and [Macki and Strauss \(1982\)](#).

Note that due to the definition of the Hamiltonian \mathcal{H} , equation (a) from the theorem simply becomes the dynamics equation $\dot{x} = f(x, u)$. Equation (b), which is called the *adjoint equation*, can also be written without vector notation:

$$\dot{\lambda}_i^*(t) = -\frac{\partial}{\partial x_i} \mathcal{H}(x^*(t), u^*(t), \lambda^*(t)), \quad \text{for } i = 1, 2, \dots, m. \quad (3.3)$$

Finally, equation (c), implies that we can also solve the following equation:

$$\nabla_u \mathcal{H}(x^*(t), u^*(t), \lambda^*(t)) = 0, \quad (3.4)$$

assuming the differentiability of the costs q and the dynamics f . The previous equation potentially yields a local instead of a global optimum, but will still turn out to be useful in our analytical computations.

LEV PONTRYAGIN 1908-1988



Lev Pontryagin was a mathematician from the Soviet Union. Despite becoming blind by the age of 14, his mother, by reading him mathematical papers out loud, empowered him to contribute to the field of topology. In 1931, he pledged to focus his work to be in line with the policies of the Communist Parties. Then, after WW2, with the advancement of missile and aerospace technology, he redirected his research towards the field of Optimal Control, designed among other things to find optimal fuel emissions paths for jet engines.¹

¹Information obtained from [Pesch and Plail \(2012\)](#) and [The Editors of Encyclopaedia Britannica \(2018\)](#).

3.1.2. VARIATIONAL APPROACH

The results obtained using Pontryagin's Principle can also be derived using a variational approach. This method uses the formalism of Calculus of Variations. The goal is identical to the Hamiltonian approach: minimise a cost functional $P[x, u]$ subject to the dynamics $\dot{x}(t) = f(t, x(t), u(t))$. Instead of using a cost term $\Psi(x(T))$ for the terminal value, we incorporate the boundary conditions in the minimisation. We now write $P[x, u]$ as:

$$P[x, u] = \int_0^T q(t, x(t), u(t)) dt$$

This method is a generalisation of the Lagrange multipliers used to find fixed equilibrium points subject to some constraints. Now, these constraints are time dependent. Therefore, we define the Lagrangian L as:

$$L(t, x, \dot{x}, u, \dot{u}, \lambda, \dot{\lambda}) = q(t, x, u) + \lambda(\dot{x} - f(t, x, u)),$$

where $\lambda(t)$ is the same quantity as the costate variable in the Pontryagin setting. Note that in this example, L does not depend on \dot{u} , but in general, there can be an explicit dependence on \dot{u} .

We now define the Lagrangian *functional* as the following integral:

$$\mathcal{L}[x, u, \lambda] = \int_0^T L(t, x(t), \dot{x}(t), u(t), \dot{u}(t), \lambda(t), \dot{\lambda}(t)) dt$$

In order to minimise this quantity, as quoted from [Courant and Hilbert \(1953\)](#), "a necessary condition for a minimum is the vanishing of the first variation". The first variations, also called functional derivatives, and denoted by $\frac{\delta \mathcal{L}}{\delta x}$, $\frac{\delta \mathcal{L}}{\delta u}$ and $\frac{\delta \mathcal{L}}{\delta \lambda}$, can be calculated using the Euler-Lagrange equations, which is a well known result from statistical physics. These equations are:

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta x} &= \frac{\partial L}{\partial x} - \frac{d}{dt} \frac{\partial L}{\partial \dot{x}}, \\ \frac{\delta \mathcal{L}}{\delta u} &= \frac{\partial L}{\partial u} - \frac{d}{dt} \frac{\partial L}{\partial \dot{u}}, \\ \frac{\delta \mathcal{L}}{\delta \lambda} &= \frac{\partial L}{\partial \lambda} - \frac{d}{dt} \frac{\partial L}{\partial \dot{\lambda}}. \end{aligned}$$

Using these, the necessary conditions for optimality are obtained by first solving $\frac{\delta \mathcal{L}}{\delta x} = 0$, giving the adjoint equation:

$$\begin{aligned} \frac{\partial L}{\partial x} - \frac{d}{dt} \frac{\partial L}{\partial \dot{x}} &= 0 \\ \Leftrightarrow \frac{\partial q}{\partial x} - \lambda(t) \frac{\partial f}{\partial x} - \dot{\lambda}(t) & \\ \Leftrightarrow \dot{\lambda}(t) &= \frac{\partial q}{\partial x} - \lambda(t) \frac{\partial f}{\partial x} \end{aligned} \quad (3.5)$$

then by solving $\frac{\delta \mathcal{L}}{\delta u} = 0$, yielding the optimality condition:

$$\begin{aligned} \frac{\partial L}{\partial u} - \frac{d}{dt} \frac{\partial L}{\partial \dot{u}} &= 0 \\ \frac{\partial q}{\partial u} - \lambda(t) \frac{\partial f}{\partial u} &= 0, \end{aligned} \quad (3.6)$$

since L does not depend on \dot{u} . Note that Eq. 3.5 and 3.6 are exactly equal to Eq. 3.3 and 3.4. We have therefore obtained the same necessary conditions as the Pontryagin's Principle using a variational approach. However, throughout this thesis, we will use the Hamiltonian setting from Pontryagin's Principle, since it is typically used for optimal control problems.

3.1.3. BELLMAN EQUATION

Some optimisation problems are intrinsically discrete. While it is possible to translate Pontryagin's Principle to a discrete setting, it is sometimes possible to calculate global optima numerically: the principles of *dynamic programming* can be used, by applying the Bellman Equation. This yields necessary and sufficient conditions for optimality. The main idea of dynamic programming is that an optimisation problem can be split into smaller subsets. Each subset is then solved, or optimised, and combined in the larger problem. This principle will become clearer with the mathematical formulation in this section. We primarily use the notation from [Kappen \(2011\)](#).

Let x_t be the value of the state variable at time t , and u_t the corresponding control variable. The dynamics, which are now discrete, are described in general by:

$$x_{t+1} = x_t + f(t, x_t, u_t) \quad (3.7)$$

Instead of calculating the total costs with an integral, they are now defined as:

$$C_T = \sum_{t=0}^{T-1} \frac{1}{(1+r)^t} R(t, x_t, u_t) + \Psi(x_T), \quad (3.8)$$

where $R(t, x_t, u_t)$ are the costs incurred when taking action u_t at time t , and $\Psi(x)$ is the penalty depending purely on the final value of the state variable. In the Bellman Equation, the principle of dynamic programming is integrated in the so-called *value function*: this represents the remaining costs from time t to T , given that the remaining part is calculated along an optimal path. In other words, we disregard how the system got to the state x_t at time t , we purely look at the minimum amount of costs left of going from the current state x_t to the end state x_T . Let the value function $J(t, x_t)$ be defined as:

$$J(t, x_t) = \min_{u_{t:T-1}} \left(\sum_{s=t}^{T-1} \frac{1}{(1+r)^s} R(s, x_s, u_s) + \Psi(x_T) \right).$$

Note the strong resemblance to Eq. 3.8: the only difference is that these are the costs starting at time t , and not time 0 (the value function is therefore also sometimes called the *cost to go*), and that the right hand side is minimised over the controls $u_{t:T-1}$.

By splitting up the sum in the expression for $J(t, x_t)$, we obtain successively the following expressions:

$$J(t, x_t) = \min_{u_t} \left(\frac{1}{(1+r)^t} R(t, x_t, u_t) + \min_{u_{t+1:T-1}} \left(\sum_{s=t}^{T-1} \frac{1}{(1+r)^s} R(s, x_s, u_s) + \Psi(x_T) \right) \right). \quad (3.9)$$

$$= \min_{u_t} \left(\frac{1}{(1+r)^t} R(t, x_t, u_t) + J(t+1, x_{t+1}) \right). \quad (3.10)$$

By using the dynamics of the system and combining it with the terminal condition $J(T, x_T) = \Psi(x_T)$, we obtain the recursive definition of the value function:

$$\begin{cases} J(T, x_T) = \Psi(x_T) \\ J(t, x_t) = \min_{u_t} \left(\frac{1}{(1+r)^t} R(t, x_t, u_t) + J(t+1, x_t + f(t, x_t, u_t)) \right) \end{cases} \quad (3.11)$$

Due to the recursive nature of the value function, it has to be calculated backwards in time. Moreover, the value function has to be calculated for every possible value of the state variable x_t . For high dimensional problems, this quickly becomes numerically challenging due to the curse of dimensionality.

While calculating the value function will indeed give the minimal costs along an optimal trajectory by evaluating $J(0, x_0)$, where x_0 is the chosen initial condition, it does not solve the optimal control problem yet. In fact, this requires that for every time t and every value of the state variable x_t , the optimal control u_t^* which minimises the value in 3.11, is saved as well. We call this step the *backward induction* step. Once the value

function and the optimal controls $u^*(t, x_t)$ are calculated, the resulting state variable path can be calculated by starting with the initial value x_0 , and advancing it using the dynamics of Eq. 3.7 with the corresponding control $u^*(t, x_t)$. This last step is the *forward sweep*.

The process described by the Bellman Equation can be summarised in the following three steps:

1. INITIALISATION STEP: calculate $J(T, x_T) := \Psi(x_T)$.
2. BACKWARD INDUCTION: calculate recursively, for $t = T - 1, T - 2, \dots, 0$ the optimal control

$$u^*(t, x_t) := \arg \min_{u_t} \left(\frac{1}{(1+r)^t} R(t, x_t, u_t) + J(t+1, x_t + f(t, x_t, u_t)) \right)$$

and corresponding value function:

$$J(t, x_t) = \frac{1}{(1+r)^t} R(t, x_t, u^*(t, x_t)) + J(t+1, x_t + f(t, x_t, u^*(t, x_t)))$$

3. FORWARD SWEEP: calculate the optimal trajectory starting with initial value x_0 for $t = 0, \dots, T - 1$:

$$x_{t+1}^* = x_t^* + f(t, x_t^*, u^*(t, x_t^*)).$$

3.1.4. HAMILTON-JACOBI-BELLMAN EQUATION

The previous derivation is only valid for optimal control problems which are discrete in time. For continuous (in time) problems, we have already introduced Pontryagin's Principle. However, as mentioned above, this only yields necessary conditions for optimality, contrary to the Bellman Equation for discrete problems. By extending the equation to continuous time problems, we can also find sufficient conditions for them: this extension is called the Hamilton-Jacobi-Bellman equation.

The equation can be derived by starting from the Bellman Equation, specifically, the formulation for the value function in Eq. 3.10. To be slightly more general and to make the notation easier, we assume no discounting factor (it can be incorporated in the cost term $R(t, x_t, u_t)$). As a reminder, the discrete value function was:

$$J(t, x_t) = \min_{u_t} \left(R(t, x_t, u_t) + J(t+1, x_{t+1}) \right)$$

Instead of taking a time step of 1, we now take an arbitrarily small time step Δt . The above equation becomes:

$$J(t, x(t)) = \min_{u(t)} \left(\int_t^{t+\Delta t} R(s, x(s), u(s)) ds + J(t + \Delta t, x(t + \Delta t)) \right) \quad (3.12)$$

The second term on the right hand side can be expanded using a Taylor series, since we have chosen Δt to be small. It can then be written as:

$$J(t + \Delta t, x(t + \Delta t)) = J(t, x(t)) + \dot{J}(t, x(t))\Delta t + \nabla J(t, x(t)) \cdot \dot{x}(t)\Delta t + \text{higher order terms}$$

By putting this into Eq. 3.12, we obtain successively:

$$\begin{aligned} J(t, x(t)) &= \min_{u(t)} \left(\int_t^{t+\Delta t} R(s, x(s), u(s)) ds + J(t, x(t)) + \dot{J}(t, x(t))\Delta t + \nabla J(t, x(t)) \cdot \dot{x}(t)\Delta t + \text{higher order terms} \right) \\ \Leftrightarrow 0 &= \min_{u(t)} \left(\int_t^{t+\Delta t} R(s, x(s), u(s)) ds + \dot{J}(t, x(t))\Delta t + \nabla J(t, x(t)) \cdot \dot{x}(t)\Delta t + \text{higher order terms} \right) \end{aligned}$$

By dividing every term by Δt and taking the limit of $\Delta t \rightarrow 0$, we first observe that $\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int_t^{t+\Delta t} R(s, x(s), u(s)) ds = R(t, x(t), u(t))$ and that the above expression becomes:

$$\min_{u(t)} \left(R(t, x(t), u(t)) + \dot{J}(t, x(t)) + \nabla J(t, x(t)) \cdot \dot{x}(t) \right) = 0$$

Finally, by replacing \dot{x} by the dynamics of Eq. 3.1, we obtain the Hamilton-Jacobi-Bellman equation:

$$\dot{J}(t, x(t)) + \min_{u(t)} \left(R(t, x(t), u(t)) + \nabla J(t, x(t)) \cdot f(t, x(t), u(t)) \right) = 0$$

The above equation is not only very difficult to solve analytically, it is also computationally very expensive: the value function has to be approximated by a discretisation in time. Analogous to solving a differential equation numerically, the precision of the numerical result is increased by using smaller time steps, which quickly becomes infeasible for most computing engines. For this reason, we have chosen to stick with the Bellman Equation for numerical purposes, and with Pontryagin's Principle for analytical results.

Chapter 4

Hotelling and Climate Policy

Climate policy has been present in policy discussions for decades. While its main focus has shifted between greenhouse gas reductions, holes in the ozone layer, air and water quality, its earliest form concerns resource management. In fact, managing forests, lakes and natural resources has been a focus for mankind for centuries or more. In this chapter, we describe the link between resource management and modern day climate policy.

Traditionally, the study of resource management has been split in two categories: renewable resources and exhaustible resources. The first category addresses the management of, for example, forests, fish ponds or livestock, which, if exploited carefully, will never deplete. On the other hand, exhaustible resources, like a coal mine, an oil field or a set of minerals in a country, are only available in a limited quantity: every unit of current consumption implies it cannot be consumed at a later stage. For the purpose of this thesis, we limit ourselves to exhaustible resources.

To make optimal use of the exhaustible resource when owning or managing an amount of this resource, it is important to know at which rate to extract and at which price to sell it. This is traditionally solved by the Hotelling Rule, which we discuss in section 4.1. Next, we will argue in section 4.2 that we can consider the climate system as an example of an exhaustible resource.

4.1. THE HOTELLING RULE

Almost a century ago, [Hotelling \(1931\)](#) wrote an influential paper exploring the optimal extraction profile of an exhaustible resource. While its empirical evidence is limited ([Livernois, 2009](#)), his paper forms the first solid theoretical foundation for pricing a non-renewable asset. In this section, we explain the setting of the Hotelling Rule, a modern proof of its main result and an economic interpretation.

Let us consider the very simplified setting where someone owns a fixed and limited stock of an exhaustible (non-renewable) resource. At every time t , an amount $E(t)$ is extracted, which is subtracted from the stock $R(t)$. Mathematically, this becomes:

$$\dot{R}(t) = -E(t)$$

The extracted resource is then sold with a net benefit (sell price minus extraction costs) $p(t) - c(t)$. Without loss

HAROLD HOTELLING (1895-1973)



Hotelling was an American mathematical statistician. He earned his PhD in Topology, but quickly focused his research on statistics and economics. During WW2, he led a statistics research group on munition logistics and hired several Jewish-European refugees in his group. Besides Hotelling's Rule, he is famous for Hotelling's Law (the principle of minimum differentiation) and Hotelling's Model (location preference of similar companies).¹

¹Information obtained from [Arrow and Lehmann \(1979\)](#).

of generality, we assume that the extraction costs are 0. The revenue at time t is now given by the price times the extracted amount, discounted at interest rate r . The goal of the owner is to maximise the total revenue:

$$\underset{p(t)}{\text{maximise}} \int_0^{\infty} e^{-rt} p(t) E(t) dt, \quad \text{subj. to } \dot{R}(t) = -E(t)$$

Using the Pontryagin's Principle, we can solve this by analysing the corresponding Hamiltonian \mathcal{H} :

$$\mathcal{H} = e^{-rt} p(t) E(t) + \lambda(t) (-E(t)), \quad (4.1)$$

where $\lambda(t)$ is a time varying Lagrange multiplier called the *costate variable*. Economically, it represents the shadow cost of extracting one unit of resource more. The optimisation problem is then solved by solving the following system, as defined in Theorem 1, where the extraction rate $E(t)$ is the control variable, and the stock $R(t)$ the state variable:

$$\begin{cases} \frac{\partial \mathcal{H}}{\partial E(t)} = 0, \\ \frac{\partial \mathcal{H}}{\partial R(t)} = -\dot{\lambda}(t), \\ \frac{\partial \mathcal{H}}{\partial \lambda(t)} = \dot{R}(t) \end{cases} \quad (4.2)$$

From the first equation, we obtain $e^{-rt} p(t) = \lambda(t)$, while the second yields $\dot{\lambda}(t) = 0$, implying a constant costate variable, which we call p_0 . Combining these, we obtain the final expression for the optimal price $p(t)$:

$$p(t) = p_0 e^{rt} \quad (4.3)$$

In order to also calculate the extraction rate $E(t)$, we need an extra relationship combining the price with the extraction: this is the demand function. Since we are mostly interested in the price, we do not specify this function any further.

The previous derivation, while mathematically complete, does not give much intuition in its results. A clearer way of thinking about this result is by a no-arbitrage argument. First of all, the price of the extracted resource has to increase over time, since money a company earns now can be used freely, whereas money in the future has less "value". Second, the price should grow at least as fast as the interest rate, otherwise putting money on the bank would be more profitable. Finally, the price should not grow at a faster rate than the interest rate, otherwise it would be worth waiting indefinitely with extracting the resources. Combining this gives an intuition for the same results derived above.

We have now proven that in this very simple setting, the optimal price path grows exponentially, at the same rate as the interest rate.

4.2. TOWARDS CLIMATE POLICY

While the previous derivation yields clean analytical results, the Hotelling setting is designed for a traditional exhaustible resource, like an oil well or a gold mine. However, many view climate change also as an exhaustible resource problem.

During the Paris agreement in 2015, world leaders have agreed on keeping the global temperature rise well below 2 degrees Celsius. This emission goal can be translated to a certain *carbon budget* (Allen et al., 2009; Matthews et al., 2009), which is the maximum cumulative amount of CO₂ emissions for limiting temperature change to the agreed maximum global temperature rise. It is in this sense that we approach the problem: the stock of CO₂ that can be emitted is here the exhaustible resource (and not, for example, as one might think, the amount of fossil fuel available on Earth). The associated price represents then an internationally agreed upon *carbon tax*, often expressed as dollars per ton of emitted CO₂ (or equivalent). This is discussed in further detail in section 5.1.

Similarly to the Hotelling setting, we have to define the dynamics of the system (section 5.3). Whereas previously we were maximising the revenue of the exhaustible resource, we will similarly minimise the costs of abatement. The associated costs in this carbon budget setting are discussed in section 5.4.

One should note that this is a cost-effectiveness setting, in contrast to a cost-benefit analysis where besides taking into account the cost of switching to renewable energy, we also take into account the benefits of a better climate.

THE PARIS AGREEMENT

In 2015, representatives of 196 countries gathered in Paris to discuss an agreement to mitigate the effects of global climate change. Contrary to earlier conventions, an ambitious agreement was reached, in which countries pledged to limit the global temperature rise to 2.0°C above pre-industrial levels, and preferably below 1.5°C. Instead of imposing emission budgets from above, a bottom-up approach was adopted, where countries choose their own Nationally Determined Contribution (NDC), which are strengthened over time.²

²More information at <https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement>

Chapter 5

The Stylised Policy Setting

In order to assess the optimal global climate policy, we propose a simple, stylised, setting, which we describe in this chapter. In this setting, global CO₂ emissions are reduced compared to a baseline scenario by imposing a price, or tax, on CO₂ emissions, introduced in section 5.1. The higher this price, the more renewable, clean technologies become cheaper than their fossil fuel alternatives. This is quantified through the Marginal Abatement Cost Curve, defined in section 5.2. The carbon tax is then chosen such that it minimises the total abatement costs (section 5.4), while keeping the total emissions below a chosen carbon budget.

All of our results will contain emissions paths which are relative to a baseline. Since we are working in a stylised setting, we use the simplified baseline shown in Fig. 5.1, which match the main characteristics of the mean of the SSP2 scenarios, shown in Fig. 8 of [Bauer et al. \(2017\)](#). We call $B(t)$ the baseline emissions at time t .

5.1. CARBON PRICES AND CARBON TAXES

While many climate policy tools exist and are worth examining, many models used to assess optimal climate policy use a global carbon tax to obtain cost-optimal emission reductions.

Such a global carbon tax has not been implemented. There already exist various systems using a carbon tax, like the European Emissions Trading System, but the average carbon price is considered to be currently too low to incur significant abatement. Moreover, since these markets are regional, the problem of free-riding still exists. Agreeing on a global, significant, carbon tax remains a political challenge. How this could be reached is outside of the scope of this thesis: we assume that there is the political will and power to set a single, global, carbon tax. This gives the most cost-effective situation, the minimum costs needed to reach certain emission levels, in an ideal world, against which more realistic scenarios can be compared.

Given a carbon price, we have to define how much emission reduction relative to the baseline is achieved. The general idea behind this is that even though a specific source of renewable energy is more expensive than, say, coal, if the price of carbon emissions is high enough, it will become cheaper to switch to the renewable alternative. This is quantified in the form of MAC curves, which we discuss in the next section.

5.2. MAC CURVES

A transparent way to determine how much CO₂ can be abated relative to the baseline, given a carbon price, is by using Marginal Abatement Cost Curves (MAC curves). These functions give the carbon tax necessary to achieve a certain reduction potential. They are based on the idea that not every sustainable technological or societal change comes at the same price. For example, replacing incandescent light bulbs by LED lighting will reduce emissions at a very low cost, whereas directly capturing CO₂ from the atmosphere is much more expensive. A fictional MAC curve is shown in Fig. 5.2a. The transition to cleaner energy is typically split in different categories, with a corresponding reduction potential (the amount of CO₂-reduction this technology will provide)


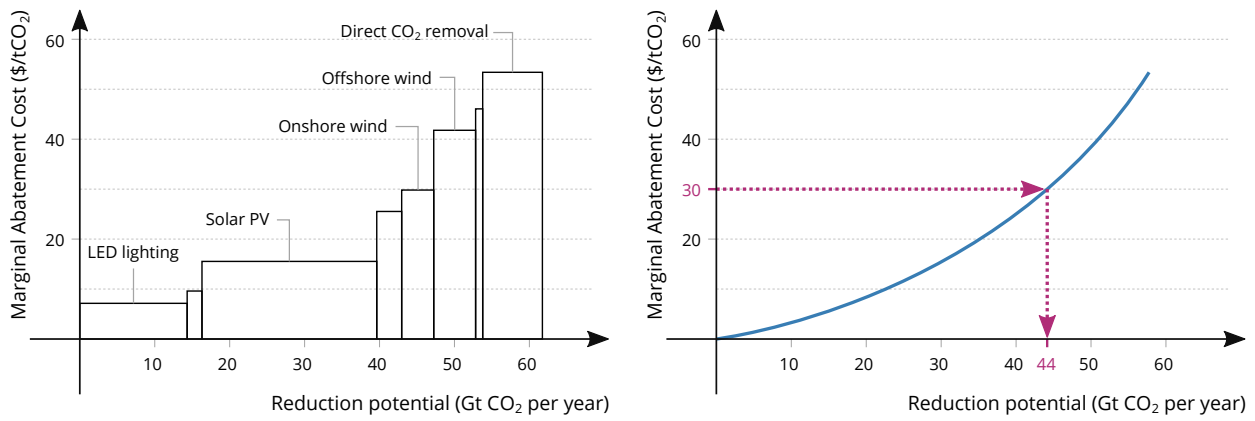


Figure 5.1: Baseline emissions in the Business-As-Usual scenario, equivalent to an emission path with constant zero carbon tax. The unit is relative to the emission level of 2015.



- (a)** The boxes represent different technological alternatives to more polluting sources. The height of the box gives the increase in cost and the width corresponds to the scale of the impact if this technology is chosen.
- (b)** Continuous approximation of the MAC curve on the left. Given a carbon tax (e.g. \$30), we can find the reduction potential (44 Gt CO₂) by calculating the inverse of the MAC curve.

Figure 5.2: Hypothetical MAC curve, giving the price required to reduce a certain amount of CO₂ emissions. In reality, they contain many more technologies, are different per region and sector and are not constant over time.

and cost. The categories are then ordered by cost. The more realistic the MAC curve, the more categories are necessary. In fact, on a global level, the categorical MAC can be well approximated by a continuous MAC, shown in Fig. 5.2b.

The MAC curves shown in Fig. 5.2 are expressed in absolute reduction levels (Gt CO₂). However, we choose to express the reduction potential as a percentage of the baseline (see discussion on the right). Therefore, the x -axis of the MAC curve should go from 0 to 1. The absolute emission reduction is then dependent on the baseline at time t .

RELATIVE VS ABSOLUTE MAC CURVE

There is some discussion on whether the MAC should be absolute or relative. In modelling work, a *relative* MAC is often used: by imposing a global carbon tax, all countries will theoretically switch to some amount of renewable technologies: with more countries (or regions, or industries), the relative reduction stays the same. The total abatement costs, however, do become larger for more countries, and should therefore be expressed in *absolute* terms (or relative to the GDP).

Mathematically, we can express the required carbon price p for a given abatement level a at time t by:

$$p = \text{MAC}(a)$$

More often, however, we need to calculate the abatement level, given a carbon price. This is done with the inverse of the MAC curve:

$$a = \text{MAC}^{-1}(p)$$

Later on, we will include various factors that modify the MAC curve over time: exogenous technological change, for example, lowers the MAC curve over time, and endogenous technological change lowers the MAC curve depending on the cumulative reduction at that time. These will be denoted as extra factors in the MAC function, and, when explicitly used, will be written as $\text{MAC}(p; \cdot)$, where \cdot stands for the extra dependent variables.

While the MAC curve describes the level of emission abatement in a given year, we still have to define the dynamics of the climate policy system we consider. This is detailed in the next section.

5.3. DYNAMICS

In the Hotelling setting, the dynamics described the change in stock of the exhaustible resource, $\dot{R}(t) = -E(t)$: the stock is reduced at exactly the extraction rate, until the stock was exhausted. In this setting, the cumulative emissions are increased at exactly the emission rate, until the carbon budget is reached:

$$\dot{CE}(t) = E(t), \quad (5.1)$$

where $CE(t)$ is the cumulative emissions from 2015 to time t and $E(t)$ the emissions at time t . The latter can be instantaneous, or yearly when we consider a discrete time setting.

The emission level $E(t)$ can be calculated explicitly using the MAC curve. We assume that when no carbon tax is set, the emissions are equal to the baseline $B(t)$. However, as we saw in the previous section, when setting a carbon tax $p(t)$, the emissions are reduced by $\text{MAC}^{-1}(p)$ relative to the baseline:

$$E(t) = B(t) \left(1 - \text{MAC}^{-1}(p(t)) \right) \quad (5.2)$$

Together with Eq. (5.1), this becomes:

$$\dot{CE}(t) = B(t) \left(1 - \text{MAC}^{-1}(p(t)) \right) \quad (5.3)$$

The exact shape and evolution of the MAC curve will be stated in more details when describing the various scenarios we consider. Before we do that, we still have to define the associated costs of choosing a certain carbon price path.

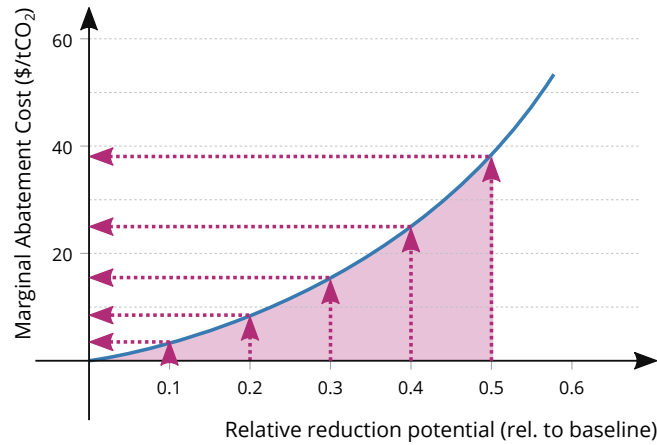


Figure 5.3: Associated costs given a reduction level. If we want to achieve a 50% reduction, then the first 10% are very cheap, the next 10% slightly higher, and so on, until the last 10% which are at the full carbon price. Hence, the costs of achieving a 50% reduction are equal to the integral of $MAC(a)$ with a going from 0 to 0.5.

5.4. COSTS

Just like in the previous section, we start by drawing the parallel with the classical Hotelling setting. The total revenue to be maximised was the selling price times the amount extracted, discounted with interest rate r , then summed (or integrated) over time.

Here the situation is very similar, with the exception that we are not maximising a revenue but minimising a cost. In fact, the costs are equal to the reduction level times the cost of reducing that much, discounted and summed (or integrated) over time. First of all, we can determine the (relative) reduction level as $a_t = MAC^{-1}(p)$. However, not the whole reduction a_t was effectuated at cost p . As can be seen in Fig. 5.3, the first part of the reduction can be achieved at a much lower cost than the last part of the desired reduction. Therefore, the total costs of achieving reduction a_t at time t are:

$$R(t, p(t)) = B(t) \left(\int_0^{a_t} MAC(a) da \right) = B(t) \left(\int_0^{MAC^{-1}(p(t))} MAC(a) da \right),$$

where we multiply by $B(t)$ to transform the relative MAC to absolute costs.

The total costs we try to minimise are therefore given by:

$$C = \int_0^T e^{-rt} B(t) \left(\int_0^{MAC^{-1}(p(t))} MAC_t(a) da \right) dt. \quad (5.4)$$

When calculating optimal carbon price paths, we will always use a discount rate $r = 5\%$, which is typically used in Integrated Assessment Modelling. In chapter 9, we discuss the importance of this factor, along with the implications of changing it to slightly different values.

5.4.1. ECONOMIC INTERPRETATION

The costs associated with the carbon price can be viewed from two perspectives. From a macro-economic perspective, the costs $R(t)$ are exactly equal to the loss of economic efficiency due to the use of renewable energy, instead of cheaper fossil fuels, for example. This is money that is lost to the economy, which we therefore want to minimise.

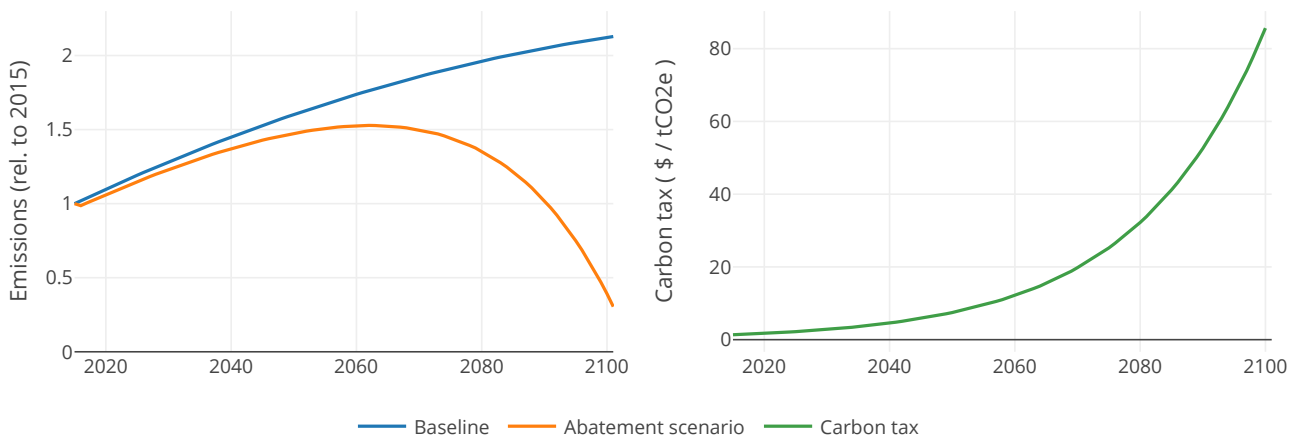


Figure 5.4: Given an example price path $p(t)$ (on the right), we can calculate the corresponding emission path (orange, on the left). This emission path can be compared to the baseline (in blue).

On the other hand, from an industry perspective, there is also a cost associated to the carbon tax itself. For every ton of CO₂ a company emits, they have to pay a carbon tax $p(t)$. This would suggest that on top of the price $R(t)$, we would need a term $p(t) \cdot E(t)$, which is the total paid in taxes. However, this last term is money that can directly be reinvested in the economy, in, for example, R&D in new sustainable technology, in infrastructure or education. Therefore, in a perfectly efficient economy, this amount is not seen as a cost, but merely as a redistribution of wealth.

5.4.2. EXAMPLE PRICE PATH

By using Eq. (5.3) and Eq. (5.4), we can calculate the emission path and associated total cost, given a price path $p(t)$, for every $t \in [0, T]$. In Fig. 5.4, we show the emission path corresponding to an exponentially increasing price path, calculated using a linear MAC curve. In the next chapter, we dive in more detail into these results, as well as when varying different assumptions.

We have now discussed all the ingredients needed to describe the system. While the expressions for the costs are more complicated than in the Hotelling system due to the non-constant MAC curves, the idea behind this description is very similar to the idea behind the Hotelling setting. We now proceed with describing a few simple scenarios in chapter 6 and solving the associated optimal control problems analytically or numerically in chapter 8.

Chapter 6

Scenarios and functional forms

To determine the (non)optimality of the Hotelling rule, we have designed a number of scenarios that investigate various forms of technological progress and emission constraints. This chapter provides functional forms and concrete descriptions of two types of technological learning (learning over time and learning by doing), as well as two types of constraints (inertia and a minimum emission level).

6.1. REFERENCE

The first and simplest scenario we analyse is what we consider as the reference. We assume that the Marginal Abatement Cost curve, defining the costs required to reduce a unit of CO₂ emissions, is linear and constant over time. In the setting described above we have only used the inverse of the MAC — giving the emission reduction as function of a chosen carbon tax — and the area under the MAC — giving the abatement costs associated with such an emission reduction. Without loss of generality, we have assumed that a carbon price of \$100/tCO₂ achieves a 100% emission reduction. In mathematical terms, this gives us:

$$MAC^{-1}(p) = p/100$$
$$\int_0^{a_{end}} MAC(a)da = 100 \cdot \frac{1}{2} \cdot a_{end}^2$$

Note that there is no restriction on the carbon tax p , except that it be non-negative. This means that for high enough carbon tax, the emission reduction can actually exceed 100%, leading to net-negative emissions. While seemingly counter-intuitive, technologies for Carbon Dioxide Removal (CDR, see explanation box) are becoming increasingly advanced. The Intergovernmental Panel on Climate Change (IPCC), an organisation of independent experts in climate change, has created a large number of scenarios predicting possible pathways to reaching the goals set during the Paris Agreement. Many of these predict extensive usage of CDR towards the end of the century. For this reason, and for its mathematical simplicity, we include this possibility in our reference scenario.

NEGATIVE EMISSIONS

Carbon Dioxide Removal (CDR) techniques are getting more and more attention. They mainly rely on two distinct techniques: Direct Air Capture (DAC) and Bio-Energy with Carbon Capture and Storage (BECCS). The former are huge filters sucking out CO₂ directly from the air. Due to the relatively low concentration of CO₂ in the atmosphere (less than 0.05% of the air), this technique is still costly and requires high amounts of energy. The latter consists of burning biomass (like trees and plants, which throughout their life, have absorbed high amounts of CO₂) and capturing the CO₂ resulting from combustion. While being a promising technology, a few major issues are still unsolved, like the effect of biomass plantations on food production and biodiversity, or the safe storage of the captured CO₂ in, for example, old gas fields.

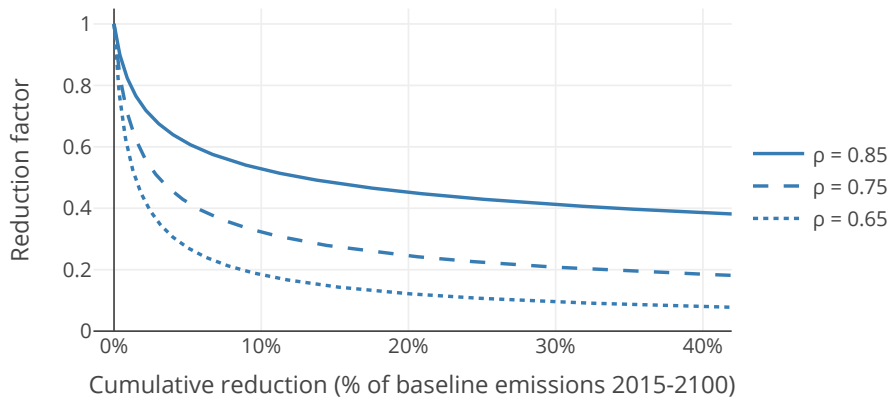


Figure 6.1: Reduction factor in the Learning By Doing setting for various progress ratios.

6.2. LEARNING OVER TIME

In the previous scenario, the MAC was constant over time. However, it is more realistic to assume that the abatement costs are now more expensive than in 50 years time, thanks to evolving technology. This motivates the introduction of a learning factor, which depends purely on time. This situation is also called exogenous technological growth: the abatement technologies just become cheaper because of “outside” factors.

Let the MAC in the beginning of the simulation be equal to the previously defined linear $MAC(a)$. Now, since it depends on time, we write explicitly $MAC(a; t)$. The learning over time is modelled as:

$$MAC(a; t) = MAC(a) \left(\frac{1}{(1 + \text{rate})^t} \right), \quad (6.1)$$

where the parameter rate takes values between 0 and 1. This means that the inverse MAC will be divided by the learning factor, and the integral of the MAC, because of its linearity, will be multiplied by the factor.

There are various ways to calibrate the rate parameter. First of all, this learning over time can be used to model improvements in the efficiency of energy use: this is the so called Autonomous Energy Efficiency Improvement index (AEEI) (see Löschel (2002) for more information). In the DICE model, a famous model used for abatement options using climate damage (Nordhaus, 1993), this autonomous learning rate is set to 2.5% per 5 years. Another way of looking at learning over time is as being an approximation of learning by doing, which we will discuss in the next section. In this sense, the rate parameter is larger, matching the reductions in abatement cost from a learning by doing scenario. This implies a value of rate between 2 and 3% per year.

6.3. LEARNING BY DOING

While technological advancement over time is already more realistic than a static MAC curve, it is argued that technological advancements happen more endogenously (Kahouli-Brahmi, 2008): this is more commonly known as Learning by Doing, or Induced Technological Change (Goulder and Mathai, 2000).

We use the functional form of learning by doing described in Van Vuuren (2007). The cost is lowered proportional to the cumulative reductions, to a certain power $-\pi$, where π is the learning rate. In our notation, the cumulative reductions are equal to the cumulative baseline $B_C(t) = \int_0^t B(s) ds$ minus the cumulative emissions, $CE(t)$. This leads to a reduction factor of:

$$\text{factor} = \alpha \left(B_C(t) - CE(t) + 1 \right)^{-\pi}.$$

The added 1 in the above expression simply comes from the fact that we want the factor to be equal to 1 when there has been no reduction at all.

It is sometimes more intuitive to express the learning rate π in terms of the progress ratio, ρ . This represents how much lower the reduction costs are when the cumulative reductions are doubled. Mathematically, we can then say that $\rho = 2^{-\pi}$. Studies have shown that this ratio lies between 0.65 and 0.95 (median 0.82) (Van Vuuren, 2007). For simplicity, we have chosen a value of $\alpha = 1$. In Fig. 6.1 we have shown the reduction factor for various values of ρ .

Bringing everything together gives us the following MAC curve:

$$\text{MAC}(a; t, \text{CE}(t)) = \text{MAC}(a) \left(B_C(t) - \text{CE}(t) + 1 \right)^{\log_2(\rho)}.$$

6.4. INERTIA

In the previous scenarios, especially with learning over time and learning by doing, the emission reduction starts very gradually, but becomes very steep after a while, as shown in more detail in chapter 8. This means that during a very short period, a huge amount of CO₂ emission reduction must be achieved. This is considered to be very challenging: for economical, political and societal reasons, such a quick transition could very well be impossible (Ha-Duong et al., 1997; Li and Strachan, 2017). In other words, there is an inertia associated to changes in the climate energy system.

We model this inertia by imposing a maximum reduction rate on the emission profile, essentially limiting the derivative of $E(t)$. Even if we choose a very steeply increasing carbon tax, the emission levels will not go down faster than the chosen inertia rate. The implementation of this brings along a few technical difficulties, which are discussed in section 7.2.3.

6.5. MINIMUM EMISSION LEVEL

Putting a limitation on the reduction speed still assumes that the emissions can become as low as needed. However, because of, for instance, hard to abate emissions and limited amounts of negative emission potentials (Gough et al., 2018), we also include a parameter to control the minimum emission level such that the resulting emission path cannot go below this level. In this scenario, we can for example model the effect of not allowing any net-negative emissions by setting the minimum level to zero.

6.6. COMBINATION

This scenario investigates the combined effect of learning by doing, inertia and imposing a minimum emission level.

6.7. DISCONTINUOUS OR CONVEX MAC

Finally, we analyse the influence of a different shape of the MAC by using a discontinuous MAC (an example of which is shown in Fig. 6.2) and a convex MAC (also shown in the same figure). The discontinuous MAC is an approximation of the linear MAC using a step function, where the number of steps used can be configured using the parameter numDiscont.

For the convex MAC, we use a cubic function: $\text{MAC}(a) = 200a^3$. Note that we use a factor 200 instead of the factor 100 for the linear MAC. While this could be chosen arbitrarily, we have chosen this number such that the area under the MAC when abating 100% is equal for the linear MAC as for the cubic MAC.

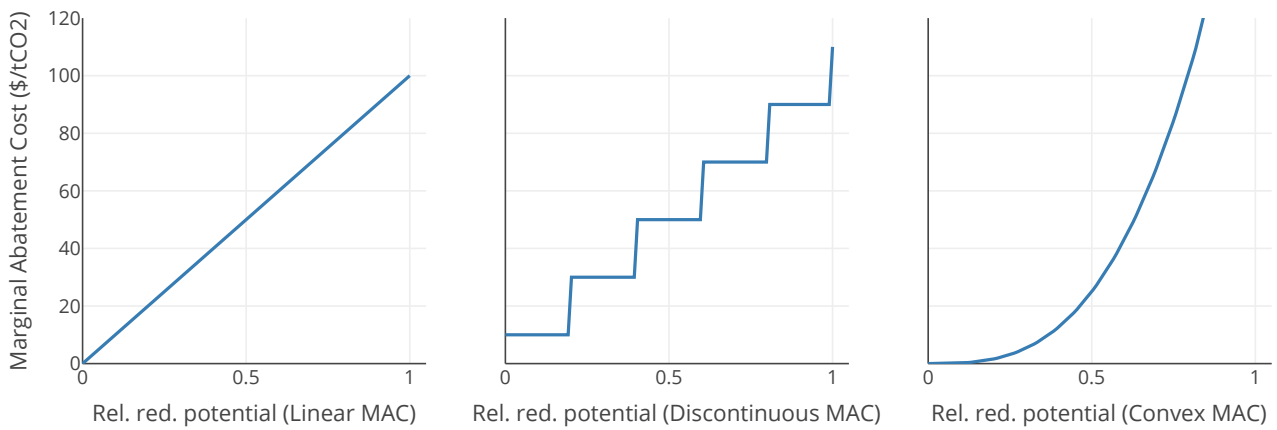


Figure 6.2: MAC Shapes. We use three different shapes of the MAC at time $t = 0$: linear, discontinuous and convex. The discontinuous MAC has a parameter, `numDiscont`, which gives the number of discontinuities, spread out evenly.

Now that we have described the setting in which we work and the scenarios we consider, we proceed to explain the solving techniques we implement to obtain our results: the optimal carbon tax paths.

Chapter 7

Solving the optimal control problem

The main goal of this thesis is to calculate the optimal carbon taxes under various scenarios and assumptions. In this chapter, we detail how we implement and calculate such optimal carbon tax paths. A reader less interested in these more technical details can safely skip to the next chapter, where we present our results.

In chapter 5, we introduced the setting in which we calculate such price paths. This can be summarised in the following optimal control problem:

$$\begin{aligned} & \underset{p(t)}{\text{minimise}} \int_0^\infty e^{-rt} B(t) \left(\int_0^{\text{MAC}^{-1}(p(t))} \text{MAC}(a) da \right) dt \\ & \text{subj. to } \dot{\text{CE}}(t) = B(t) \left(1 - \text{MAC}^{-1}(p(t)) \right), \\ & \text{CE}(T) = \text{CB}, \end{aligned} \tag{7.1}$$

where, as a reminder, $B(t)$ is the baseline emission scenario, $p(t)$ the carbon tax, $\text{CE}(t)$ the cumulative emissions and CB the carbon budget. In this setting, the carbon tax $p(t)$ is the *control variable*, which controls the *state variable* $\text{CE}(t)$, the cumulative emissions. We will first solve this analytically for a few simple scenarios. However, as we add more assumptions to the model (for example, a more complicated functional form for the MAC curve or more limitations on emission levels), this becomes unfeasible: we then resort to numerical global optimisation techniques.

STATE VARIABLE: CUMULATIVE OR INSTANTAN. EMISSIONS?

While it is clear that since the carbon tax is the main policy tool, it is also the *control variable* in this optimal control problem. There are more options for the *state variable*: both the cumulative $\text{CE}(t)$ and instantaneous $E(t)$ emissions are eligible choices. However, since the extra constraint is that the cumulative emissions are bounded in the year 2100, this constraint is much easier to implement as the terminal value of $\text{CE}(T)$ than as a constraint on the *integral* of the values $E(t)$.

7.1. ANALYTICAL SOLUTION

In order to obtain the analytical results, we will analyse the necessary conditions given by Pontryagin's Principle. More theoretical background, as well as the details of the theorem, are given in chapter 3. It has to be kept in mind that these results are indeed *necessary*, but not always sufficient. We start this analysis with the most general setting. In the next sections, we will assume more concrete functional forms.

To start, we need to define the (present-time) Hamiltonian:

$$\mathcal{H}(\text{CE}(t), \lambda(t), p(t), t) := -e^{-rt} B(t) \left(\int_0^{\text{MAC}^{-1}(p(t))} \text{MAC}(a) da \right) - \lambda(t) B(t) \left(1 - \text{MAC}^{-1}(p(t)) \right), \tag{7.2}$$

where $\lambda(t)$ is the shadow cost of emissions at time t (or, more mathematically, the costate variable of the optimal control problem). Since higher emission levels are actually a cost to society, and not a benefit, $\lambda(t)$ is actually the negative of the costate variable, to give a better intuition of $\lambda(t)$. This has to be taken into account when setting up the adjoint equation, which now becomes $-\dot{\lambda}(t) = -\partial\mathcal{H}(\cdot)/\partial\text{CE}(t)$.

A first order necessary condition for optimality, given by Pontryagin's Principle, is that the value of the optimal carbon tax $p^*(t)$ at time t has to be an extremal point of the Hamiltonian (here, we assume the maximum is an internal point of the set of possible values of $p(t)$):

$$\begin{aligned} \frac{\partial\mathcal{H}(\text{CE}(t), \lambda(t), p(t), t)}{\partial p(t)} &= 0 \\ \Leftrightarrow -e^{-rt}B(t)\text{MAC}(\text{MAC}^{-1}(p(t)))\frac{\partial}{\partial p(t)}\text{MAC}^{-1}(p(t)) + \lambda(t)B(t)\frac{\partial}{\partial p(t)}\text{MAC}^{-1}(p(t)) &= 0 \end{aligned}$$

where we used the Fundamental Theorem of Calculus together with the chain rule. Note that $\text{MAC}(\text{MAC}^{-1}(p(t))) = p(t)$, which allows us to simplify the equation to:

$$\Leftrightarrow B(t) \left(\frac{\partial}{\partial p(t)} \text{MAC}^{-1}(p(t)) \right) \left(-e^{-rt}p(t) + \lambda(t) \right) = 0$$

Assuming now that $\partial\text{MAC}^{-1}(p(t))/\partial p(t) \neq 0$, and that $B(t) \neq 0$ as well, we can conclude that:

$$p(t) = \lambda(t)e^{rt} \quad (7.3)$$

This is a significant result: depending on the sign of $\dot{\lambda}(t)$, we can directly conclude whether $p(t)$ grows at exponential rate r (which would confirm the Hotelling rule), faster (when $\lambda(t) > 0$) or slower (when $\lambda(t) < 0$) than exponential.

7.2. NUMERICALLY

While we are able to infer some analytical results for the carbon path and the emission profile, more complicated scenarios become too difficult: we solve these numerically. Specifically, we implement the Bellman Equation (see chapter 3 for more details). This has as added advantage that we find the global optimum, instead of risking to end up in a local optimum.

The Bellman equation treats problems with discrete timing. We therefore have to reframe the previous setting to a discrete setting. This actually also makes economic sense, since this implies that the carbon tax can only be changed, as we implement, every year, instead of continuously at every time.

7.2.1. DISCRETISING THE SYSTEM

The first step is to discretise the dynamics of equation 5.3. This becomes:

$$\text{CE}_{t+1} = \text{CE}_t + B(t) \left(1 - \text{MAC}^{-1}(p_t) \right) \quad (7.4)$$

Moreover, the costs from equation 5.4 should also be discretised. Instead of using a discount factor of e^{-rt} , we should then use $(1+r)^{-t}$. The costs become:

$$\begin{aligned} C_T &= \sum_{t=0}^{T-1} \frac{1}{(1+r)^t} B(t) \left(\int_0^{\text{MAC}^{-1}(p_t)} \text{MAC}(a) da \right) + \Psi(\text{CE}_T) \\ &= \sum_{t=0}^{T-1} \frac{1}{(1+r)^t} R(t, p_t) + \Psi(\text{CE}_T), \end{aligned}$$

where the carbon budget constraint is taken care of by introducing a penalty term $\Psi(\text{CE}_T)$. We use:

$$\Psi(x) = \mathbb{1}_{\{x > \text{CB}\}} \cdot u \cdot (x - \text{CB})^3,$$

where CB is the carbon budget, $\mathbb{1}_a$ is the indicator function, which is equal to 1 if a is true, and 0 otherwise, and u is a parameter to scale the penalty (we use $u = 5000$ for best numerical stability). Note that other forms can also be used, like step functions, as long as the penalty is large enough.

It now again becomes clear why we consider the cumulative emissions, CE_t , as the state variable, instead of E_t , since otherwise the carbon budget constraint would be very difficult to enforce, as Ψ would depend on the sum $\sum_{t=0}^T E_t$ instead of just the single variable CE_T .

7.2.2. FORMULATING THE BELLMAN EQUATION

As described in chapter 3, the Bellman equation relies on the concept of the *value function*, $J(t, x)$, giving the costs that remain from time t to time T assuming that an optimal trajectory is followed. In our case, the value function is equal to:

$$J(t, \text{CE}_t) = \min_{p_t:T-1} \left(\sum_{s=t}^{T-1} \frac{1}{(1+r)^s} R(s, p_s) + \Psi(\text{CE}_T) \right),$$

which can be rewritten recursively as:

$$\begin{aligned} J(T, \text{CE}_T) &= \Psi(\text{CE}_T) \\ J(t, \text{CE}_t) &= \min_{p_t} \left(\frac{1}{(1+r)^t} R(t, p_t) + J(t+1, \text{CE}_{t+1}) \right) \end{aligned} \quad (7.5)$$

By using the dynamics of the system (Eq. 7.4), we can also express CE_{t+1} in terms of CE_t :

$$J(t, \text{CE}_t) = \min_{p_t} \left(\frac{1}{(1+r)^t} R(t, p_t) + J \left(t+1, \text{CE}_t + B(t) \left(1 - \text{MAC}^{-1}(p_t) \right) \right) \right).$$

While this equation looks very difficult, it is actually rather intuitive: the remaining costs at time t are equal to the minimum of the costs in this current step (the first term) plus the remaining costs from $t+1$ until the end (the second term).

Because of the recursive nature of this equation, we have to calculate the value function J backwards in time. The optimisation happens in three steps:

- First, calculate $J(T, \text{CE}) = \Psi(\text{CE})$ for every value of CE.
- Second, recursively for $t = T-1, T-2, \dots, 0$, calculate the optimal price p_t^* and corresponding $J(t, \text{CE}_t)$ using $J(t+1, \text{CE}_{t+1})$, again for every CE_t .
- Third, calculate the emission path using the dynamics of the system of Eq. (7.4) and the optimal prices $p_t^*(\text{CE}_t)$.

Computational considerations and precise description of interpolation methods and possible values for the state variable CE_t are detailed in section 7.3.

7.2.3. IMPLEMENTING THE INERTIA CONSTRAINT

Some special attention needs to be paid to the inertia scenario. As mentioned before, the inertia is modelled as a constraint on the derivative of the emission level $E(t)$. In the discrete setting, this can be translated to the following constraint:

$$E_t - E_{t+1} \leq M,$$

where M is the maximum reduction parameter. A problem arises directly when formulating this condition: we use the cumulative emissions CE_t as state variable, and do not keep track of the instantaneous (yearly) emissions E_t . The reason for this was that we had to be able to enforce the carbon budget, which could only be achieved using $CE(t)$ as state variable. Here, however, knowing what the cumulative emissions at time t are does not give us information about the current emission level E_t .

A TWO-DIMENSIONAL STATE SPACE

To deal with this we use both E_t and CE_t as state variables. As a consequence, we have a two dimensional state variable $x_t = (E_t, CE_t)$. The dynamics of the system, previously described by Eq. (7.4), now become:

$$\begin{cases} E_{t+1} = B(t)(1 - \text{MAC}^{-1}(p_t)) \\ CE_{t+1} = CE_t + E_{t+1} \end{cases} \quad (7.6)$$

Increasing the dimension of the state variable has an impact on the calculation of the value function. Where before it sufficed to calculate $J(t, CE)$ for every possible value of CE , we now have to calculate it for every possible combination of CE and E :

$$J(t, CE_t, E_t) = \min_{p_t} \left(\frac{1}{(1+r)^t} R(t, p_t) + J(t+1, CE_{t+1}, E_{t+1}) \right),$$

where we can fill in the values for E_{t+1} and CE_{t+1} from Eq. (7.6).

IMPLEMENTING THE INERTIA CONSTRAINT

The new formulation proposed above allows us to formulate constraints and scenarios which not only depend on time and the cumulative emissions CE_t , but also on the emissions E_t . This is required for our formulation of inertia: $E_t - E_{t+1} \leq M$, or, equivalently, $E_{t+1} \geq E_t - M$. This is implemented in the dynamics of the system. Eq. (7.6) now becomes:

$$\begin{cases} E_{t+1} = \max \left(B(t)(1 - \text{MAC}^{-1}(p_t)), E_t - M \right) \\ CE_{t+1} = CE_t + E_{t+1} \end{cases} \quad (7.7)$$

One can also model inertia as a constraint on the relative difference between two consecutive emission levels, instead of on the absolute difference $E_t - E_{t+1}$ as we just did. The constraint then becomes:

$$\begin{cases} E_{t+1} = \max \left(B(t)(1 - \text{MAC}^{-1}(p_t)), E_t(1 - M') \right) \\ CE_{t+1} = CE_t + E_{t+1} \end{cases} \quad (7.8)$$

where we now use a relative maximum reduction parameter M' , which has values between 0 and 1.

7.3. COMPUTATIONAL CONSIDERATIONS

As explained before, all of the numerical results in this thesis have been obtained by solving the Bellman Equation. This equation solves *discrete* optimal control problems. However, while our model uses indeed a discrete time variable (the carbon tax is updated yearly), the state variable is continuous: there is no well defined discrete set of (cumulative) CO₂ emission levels. In order to still be able to use the Bellman Equation, we need some form of discretisation, which we discuss in section 7.3.1. Discretising a continuous variable directly brings with it the question of interpolation, addressed in sections 7.3.2 and 7.3.3.

7.3.1. DISCRETISATION OF OBJECTIVE FUNCTION

The Bellman Equation uses the value function, which, for our climate policy setting, is defined by Eq. 7.5:

$$J(t, CE_t) = \min_{p_t} \left(\frac{1}{(1+r)^t} R(t, p_t) + J(t+1, CE_{t+1}) \right),$$

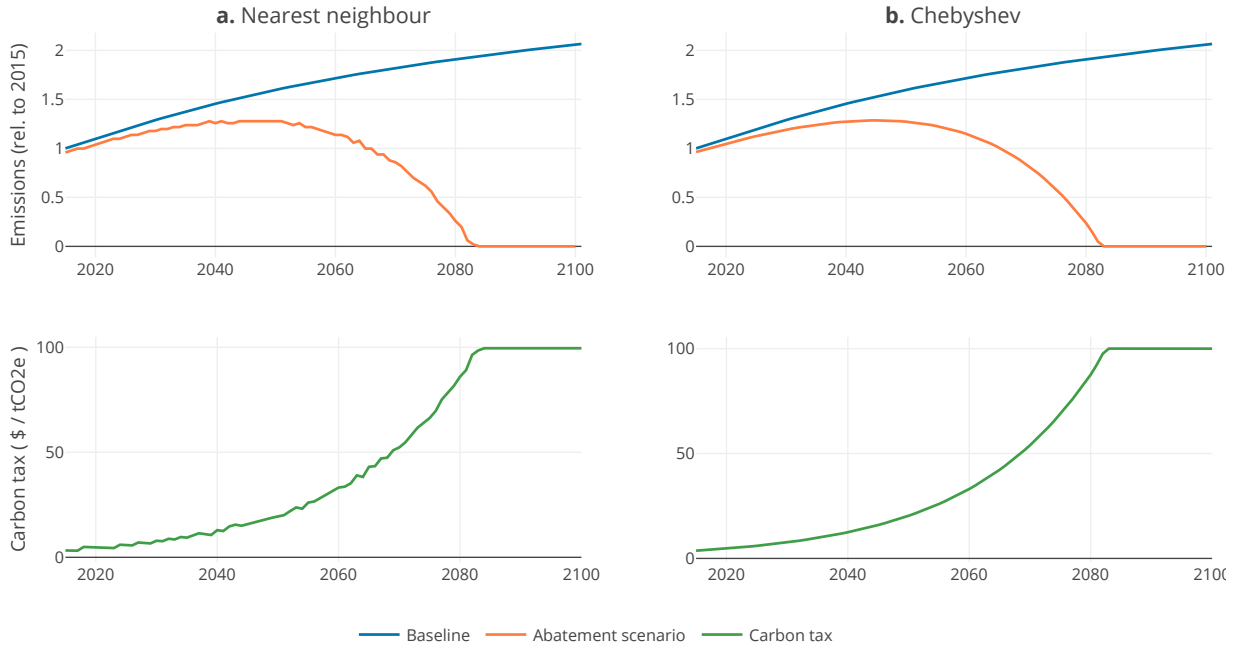


Figure 7.1: Comparison for a test run with equal run times, using a Nearest Neighbour interpolation (**a**) and Chebyshev polynomial interpolation (**b**).

where the next step of the state variable CE_{t+1} was defined as:

$$CE_{t+1} = CE_t + B(t) \left(1 - MAC^{-1}(p_t)\right). \quad (7.9)$$

To solve the Bellman Equation, a backward induction is performed: first the value of $J(T, CE_T)$ is calculated for every possible CE_T , then, using $J(T, CE_T)$, the value of $J(T-1, CE_{T-1})$ is calculated, and so on. In order to do this numerically, we discretise the (continuous) variable CE_t into a set of N equally spaced values between the minimum possible value $CE_{t,min}$ and its maximum possible value $CE_{t,max}$. In practice, we bound the cumulative emissions by the cumulative baseline emissions, since even with a zero carbon tax, the cumulative emissions will never exceed the baseline. Therefore,

$$CE_{t,max} = B_C(t),$$

and as lower bound we simply use $CE_{t,min} = 0$.

For scenarios in which we need a two dimensional state space, with the included (yearly) emission level E_t , we discretise this dimension similarly in M steps, with lower bound 0 and upper bound $B(t)$, the baseline level of that year. The state space at time t is then a two dimensional matrix with $N \times M$ elements. Note that these dimensions can also be time dependent.

While this seems to work fine at first, there is one problem: when calculating $J(t, CE_t)$ for a given value CE_t , we need the value of $J(t+1, CE_{t+1})$, where CE_{t+1} is given by Eq. 7.9. It is however very possible that CE_{t+1} does not coincide exactly with one of the discretised values between $CE_{t+1,min}$ and $CE_{t+1,max}$. For this reason, we need to interpolate the value function to be able to evaluate it for every value of the state variable.

7.3.2. INTERPOLATION

The easiest way to deal with this, is by using nearest-neighbour interpolation. Here, we simply use the value of $J(t+1, CE_{t+1})$ (and corresponding optimal carbon tax $u^*(t+1, CE_{t+1})$) for the value of the state variable closes to CE_{t+1} . The resulting optimal carbon tax is, however, very imprecise: when using $N = 5000$, the resulting

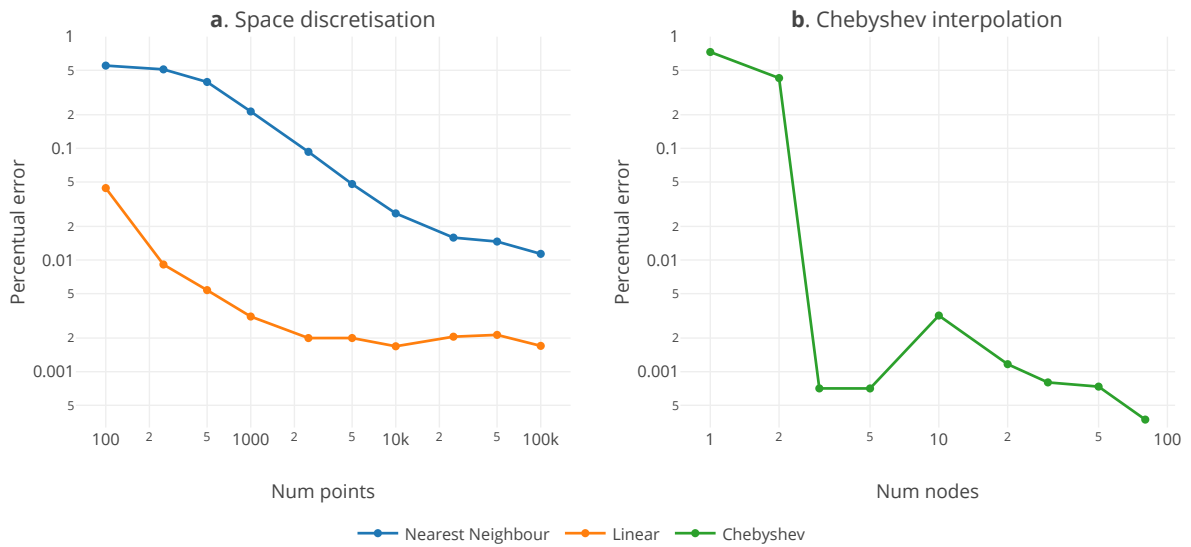


Figure 7.2: Error in resulting optimal price path compared to reference run. For the space discretisation, the number of points in which the state variable is discretised is varied (a), for the Chebyshev interpolation method the number of Chebyshev nodes (b). Note that both (a) and (b) are log-log plots.

optimal tax and emission paths, shown in Fig. 7.1 (a), is still very bumpy, especially compared to the results of the same setting, but using a higher order interpolation method (which we will explain in the next section). We can quantify this “bumpiness” by defining the percentual error as the relative absolute difference between a calculated carbon tax path and a reference path. The reference path is simply calculated by using a very high number of discretisation points. This percentual error for the nearest neighbour method is shown in Fig. 7.2 (a). While the percentual error does indeed decrease for higher N , it still takes over 100 000 points to obtain a numerical error of less than 1%.


To reduce this numerical error, a higher order interpolation method is needed. Throughout this thesis, we have used linear interpolation as interpolation method. For the two-dimensional state variable, this is extended to bilinear interpolation. The percentual error for this method is also shown in Fig. 7.2 (a). Clearly, the numerical error as function of N goes down much faster than with nearest neighbour interpolation. It does seem to stagnate around $N = 2000$ points. Then, the error is bounded from below by other numerical errors (like the discretisation of the control variable, which we have kept constant in this figure). Moreover, as we have optimised these algorithms for the GPU (Graphical Processing Unit), we use single-float precision, which has a limited accuracy as well.

While (bi)linear interpolation already yields much more accurate results, we have also investigated the use of polynomial-types of interpolation method. Specifically, we have implemented Chebyshev interpolation, treated in the next section.

7.3.3. CHEBYSHEV INTERPOLATION

The key idea behind Chebyshev interpolation is to find a suitable linear combination of so-called *Chebyshev polynomials* that best approximates a given function. This function is then sampled at certain points: for best convergence and accuracy, these sample points are chosen to be the *Chebyshev nodes*. The coefficients resulting from the optimal linear combination are then saved, and used later on to approximate the original function at any point in its interpolation domain.

Let us now translate this to a more mathematical definition. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the real valued function which



(a) Chebyshev polynomials.



(b) Chebyshev nodes.

Figure 7.3: Representation of the Chebyshev polynomials and nodes using $n = 8$ nodes.

we wish to approximate over a domain $[a, b] \subset \mathbb{R}$. Then, the approximated function \hat{f} is defined as:

$$\hat{f}(x) = \sum_{j=1}^n c_j \phi_j(x),$$

where $\phi_j(x)$ is the j -th basis function (in our case, this will be the Chebyshev polynomial, which we will define later on), and c_j its corresponding coefficient. To define the coefficients c_j , we insist that \hat{f} interpolates f at the points x_i :

$$\hat{f}(x_i) = \sum_{j=1}^n c_j \phi_j(x_i) = f(x_i), \quad i = 1, \dots, n$$

As proposed by [Chebyshev \(1854\)](#), we use the Chebyshev polynomial for our basis function, defined recursively as:

$$\begin{aligned} \phi_0(x) &= 1 \\ \phi_1(x) &= x \\ \phi_{j+1}(x) &= 2x\phi_j(x) - \phi_{j-1}(x) \end{aligned}$$

Then, as mentioned above, we still need to define the interpolation points, called here the Chebyshev nodes:

$$x_i = \frac{1}{2} \left[(a+b) - (b-a) \cos \left(\frac{\pi}{n} \left(i - \frac{1}{2} \right) \right) \right]$$

Finally, given the sampled values $f(x_i)$, the coefficients c_j are calculated by solving the set of equations $\hat{f}(x_i) = f(x_i)$ for all i . If this is not possible directly (when, for example, more sample points are used than Chebyshev nodes), a least-square fit is used.

After this brief introduction to Chebyshev interpolation in general, we proceed to a description on how this can be implemented to solve the Bellman Equation with a continuous state variable.

IMPLEMENTATION IN BELLMAN EQUATION

As mentioned during the description of linear interpolation, the main problem of using a continuous state variable in the Bellman equation is the approximation of the value function, Eq. 7.5. We now approximate the value function $J(t, CE_t)$ by its Chebyshev approximation $\hat{J}(t, CE_t)$:

$$\hat{J}(t, x) = \sum_{j=1}^n c_{t,j} \phi_{t,j}(x),$$

such that the Bellman Equation becomes:

$$\sum_{j=1}^n c_{t,j} \phi_{t,j}(CE_t) = \min_{p_t} \left(\frac{1}{(1+r)^t} R(t, p_t) + \sum_{j=1}^n c_{t+1,j} \phi_{t+1,j}(CE_{t+1}) \right),$$

where we first calculate the right hand side for the Chebyshev nodes $CE_{t,i}$, which are then used to fit the coefficients $c_{t,j}$ on the left hand side of the above equation.

The last step in the implementation is to define and calculate the bounds a and b of the interpolation domain. While we could simply use $a = 0$ and $b = \gamma \cdot CB$ (with γ large enough) such that all possible values of the cumulative emissions CE_t will be in this domain, this induces a few problems. First of all, especially for low t , these bounds are way too broad: it is highly unlikely that the cumulative emissions should already have reached the full carbon budget after only a few years. This would decrease the accuracy or efficiency of the algorithm. The second problem is that if these bounds are too loose, these values of CE_t will never be reached, regardless of the carbon tax path chosen: consequently, this brings with it discontinuities, or, at the very least, a discontinuous derivative of the value function. While this is in principle not a big problem, the accuracy of Chebyshev approximation goes down dramatically for non-differentiable target functions. For these reasons, we choose tighter bounds. In fact, at any time t , the only useful values for the cumulative emissions are those that are not too small nor too large to reach the set carbon budget at the final time T . Specifically, the lower bound is the value of CE_t such that even when no carbon tax is imposed, the cumulative emissions at time T do not exceed the carbon budget. The upper bound, similarly, is the value of CE_t such that when choosing the highest possible carbon tax for the years $t+1$ to T , the carbon budget constraint is still met. Note that this last bound only makes sense when we do indeed have a maximum carbon tax, which is the case for the minimum emission level scenario, where we limit the amount of negative emissions. Both bounds are calculated using binary search, which converges quickly to the desired levels of cumulative emissions.

This method of interpolation has been implemented, as shown in an example run in Fig. 7.1 (b). Clearly, the resulting carbon tax and emission paths are much smoother than the results of nearest neighbour interpolation using the same amount of computing time. This is directly the main advantage of this method: for only a small number (10 or 20) of Chebyshev nodes, the results are numerically very robust. In fact, the interpolation error compared to a reference run with 150 nodes goes down very fast for increasing number of nodes, as shown in Fig. 7.2 (b). Note that in this figure, already a very big drop in error is achieved for 3 and 5 nodes: these are probably the result of lucky fit of the price path, which can apparently already be described very well with a third order polynomial. However, if we disregard these points, a more exponential relationship can be observed in interpolation error reduction for increased number of Chebyshev nodes.

While this method allowed for a great reduction in computing time and an increase in accuracy, this higher order interpolation is only suitable for smooth scenarios: a discontinuous MAC will be much harder to approximate, as well as scenarios with inertia (with two dimensional state variables) and other constraints. For this reason, we have chosen to calculate all of our results with the (bi-)linear interpolation method, which is a good compromise between accuracy, speed and flexibility.

7.4. CALCULATING THE PROXY HOTELLING PATHS

While calculating the optimal carbon tax path ensures that it is in fact the most cost effective, the goal of this thesis is to compare those paths to the corresponding Hotelling path. Throughout this thesis, we call this the proxy Hotelling path: this is the carbon tax path which grows at the same rate as the interest rate (exponentially with 5% yearly growth) and which limits the emissions to the chosen carbon budget.

One of the reasons that Hotelling paths are still so popular in current IAMs is a purely computational one: calculating the optimal Hotelling path is much easier than calculating the full optimal carbon tax path. In fact, the former only requires computing the correct initial carbon tax: all subsequent tax levels are determined through the exponential growth.

We use binary search to find this initial carbon tax, implemented as follows. Let $CE_T(p_0)$ be the total emissions resulting from a Hotelling carbon tax path starting at p_0 and growing at the same rate as the interest rate, calculated using equation 7.4:

$$CE_T(p_0) = \sum_{t=0}^T B(t) \left(1 - \text{MAC}^{-1}(p_0(1+r)^t)\right)$$

Besides, let $p_{0,low}$ and $p_{0,high}$ be an initial lower and upper bound of the possible initial carbon taxes. These bounds do not have to be strict, since binary search converges very quickly (exponentially).

The binary search algorithm then relies on calculating the midpoint between $p_{0,low}$ and $p_{0,high}$:

$$CE_T \left(\frac{1}{2}(p_{0,low} + p_{0,high}) \right)$$

If this quantity is lower than the carbon budget, it means that the initial carbon tax was too high. In the next iteration, we therefore choose the new upper bound to be the current midpoint:

$$p_{0,high} \leftarrow \frac{1}{2}(p_{0,low} + p_{0,high})$$

and similarly, if the obtained cumulative emissions exceed the carbon budget, the initial tax was too low: we therefore update the lower bound to the current midpoint. This process is repeated until the difference between $p_{0,high}$ and $p_{0,low}$ is small enough. The chosen initial carbon tax is then:

$$p_0^* = \frac{1}{2}(p_{0,low} + p_{0,high})$$

From p_0^* , we directly have the full carbon tax path, and can easily calculate the resulting emission path using Eq. 7.4 again.

7.5. GPU COMPUTATION

The method described in 7.2 requires a lot of calculations. Specifically, the backward induction step when solving the Bellman Equation requires the calculation of the value function $J(t, CE_t)$ for all possible values of CE_t , repeatedly for all values of t from $t = T$ to $t = 0$. For high enough accuracy, we discretise the possible values of CE_t in 10 000 steps. This number becomes even much higher when using the inertia constraint, where we have to use a two-dimensional state space.

While the values for time step t require the calculations from step $t + 1$, all calculations *within* one time step can be performed in parallel. This is an ideal task for a Graphical Processing Unit, GPU. These pieces of hardware are designed to provide a smooth experience when playing a game or watching a movie on a PC or laptop, where the value of every pixel has to be calculated simultaneously. This is done thanks to thousands of processors which calculate in parallel. While every one of these processors is slower than the traditional CPU, the number of processors (called nodes) is so high, that the total computing power is much higher.

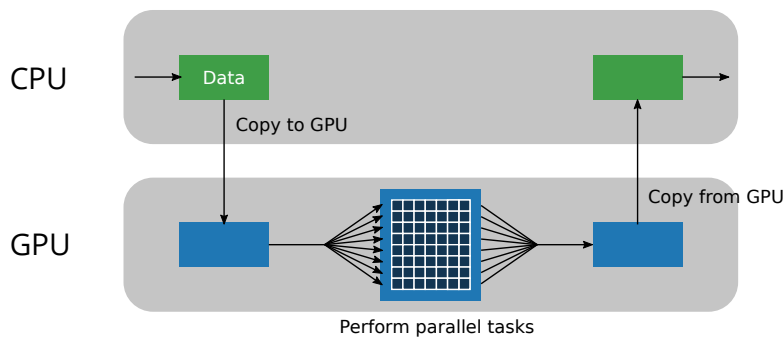


Figure 7.4: When processing data on a GPU, the data has to be first transferred to GPU memory. Then, a kernel function is mapped on all elements of the input data. The resulting output is then transferred back to the CPU/RAM, where it can be used again.

In the last few years, the use of GPU's for general and scientific computation has become increasingly popular, and is now known as GPGPU: General Purpose GPU computation. All our numerical implementations of the Bellman Equation also make extensive use of the GPU. This speeds up the computations dramatically. On a normal desktop, this brings with it already a 100x speed up.

The code becomes slightly more complicated, however, mainly due to the fact that the GPU and the rest of the computer do not share the same memory. As shown in Fig. 7.4, processing data on a GPU requires first transferring all required data to the GPU before it can be used by it. Then, a function, the *kernel*, is mapped on all elements of the input data. For example, for all values of CE_t , the kernel will calculate the value function $J(t, CE_t)$. Once all values have been calculated, the resulting data is transferred back to the CPU and the traditional memory hardware (the RAM).

Chapter 8

Results — Optimal carbon price paths

By using the mathematical and computational techniques described in the previous chapter, we now proceed to actually computing the optimal carbon tax paths, and comparing them to the Hotelling paths. First, we do this analytically for a limited amount of scenarios, then we compute the optimal paths for every scenario from chapter 6 numerically.

8.1. ANALYTICAL RESULTS

An elegant way to obtain results for the optimal carbon tax path is through analytical computations. In section 7.1, we have derived that in a general setting with optional technological growth and general MAC curve, the carbon tax path is given by:

$$p(t) = \lambda(t)e^{rt},$$

where $\lambda(t)$ represents the costate variable in the Hamiltonian of the system. In order to obtain closed form analytical results, we will need to assume a functional form for the baseline and the MAC. In this chapter, we use a quadratic baseline:

$$B(t) = b_0 + b_1t + b_2t^2, \quad (8.1)$$

and a power law Marginal Abatement Curve:

$$\text{MAC}(a) = \beta_0 a^{\beta_1}, \quad \beta_0, \beta_1 > 0 \quad (8.2)$$

The parameters can be calibrated on IPCC predictions as needed. Typically, we will use a concave baseline ($b_1 > 0, b_2 < 0$) and convex MAC ($\beta_1 > 1$), but any other combination of parameters is valid. In some situations we will use a simplification by taking a constant baseline ($b_1 = b_2 = 0$) and a linear MAC ($\beta_1 = 1$), this will be stated explicitly.

We now proceed to give analytical results for the Reference, Learning over Time and Learning by Doing scenarios.

8.1.1. REFERENCE

The easiest scenario to consider analytically is one without technological learning, without inertia, and *with* the possibility to have negative emissions. Despite the lack of learning and inertia, we keep everything else as general as possible: we use a general MAC form, and a general baseline $B(t)$.

Under these assumptions, the adjoint equation, posed by Pontryagin's Principle, becomes:

$$\begin{aligned} -\dot{\lambda}(t) &= -\frac{\partial \mathcal{H}(\cdot)}{\partial \text{CE}(t)} \\ &\Leftrightarrow \dot{\lambda}(t) = 0, \end{aligned}$$

since the Hamiltonian does not depend on the state variable $CE(t)$. In other words, $\lambda(t)$ is constant. In fact, combining it with Eq. (7.3), we see that it is equal to the initial carbon tax at time $t = 0$, which we call p_0 :

$$p(t) = p_0 e^{rt}. \quad (8.3)$$

This means that in the reference scenario, the Hotelling rule is still valid. This is not surprising: this scenario is also an almost exact translation of the original Hotelling setting to the climate policy setting. Different carbon budgets and MAC shapes have no effect on the shape of the tax path: they will only influence the initial carbon tax.

CALCULATING THE INITIAL CARBON TAX p_0

In order to calculate the initial carbon tax p_0 , we have to use the final condition on the cumulative emissions $CE(T) = CB$. We choose to express the carbon budget CB as a fraction $\alpha \in (0, 1)$ of the cumulative baseline emissions: $CB = \alpha \int_0^T B(t) dt$. The cumulative emissions, given by the integral of Eq. (5.3), should be equal to:

$$\begin{aligned} \int_0^T B(t)(1 - \text{MAC}^{-1}(p(t))) dt &= \alpha \int_0^T B(t) dt \\ \Leftrightarrow \int_0^T B(t)(1 - \alpha - \text{MAC}^{-1}(p_0 e^{rt})) dt &= 0 \end{aligned} \quad (8.4)$$

Without assuming any functional form of the baseline or the MAC curve, this expression cannot be simplified further. However, when assuming a simple linear MAC ($\text{MAC}(a) = \beta_0 a$) and constant baseline ($B(t) = b_0$), the above equation simplifies to:

$$\Leftrightarrow \int_0^T b_0 \left(1 - \alpha - \frac{p_0 e^{rt}}{\beta_0}\right) dt = 0,$$

which has as solution for p_0 :

$$p_0 = \frac{T(1 - \alpha)r\beta_0}{e^{rT} - 1}$$

CLOSED FORM INITIAL CARBON TAX FOR CONVEX MAC AND CONCAVE BASELINE

It is possible to make the above results more realistic and still obtain closed form analytical results. We now assume the quadratic form for the baseline and power law form for the MAC as defined in Eq. 8.1 and 8.2. After a few steps detailed in Appendix A.1.2, we obtain the following form for the initial carbon price p_0 :

$$p_0 = \beta_0 \left(\frac{(1 - \alpha) \left(b_0 T + \frac{1}{2} b_1 T^2 + \frac{1}{3} b_2 T^3 \right)}{\left(\frac{b_0 \beta_1}{r} - \frac{b_1 \beta_1^2}{r^2} + \frac{2b_2 \beta_1^3}{r^3} \right) (e^{rT/\beta_1} - 1) + e^{rT/\beta_1} \left(b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1 T^2}{r} - \frac{2\beta_1^2 T}{r^2} \right) \right)} \right)^{\beta_1} \quad (8.5)$$

8.1.2. LEARNING OVER TIME

The previous scenario confirmed the Hotelling rule under the no-technological-learning assumption. This scenario relaxes this assumption by now considering learning over time, also called exogenous technological growth.

We again keep the most general form where possible: a general baseline $B(t)$ and a general MAC curve $\text{MAC}(a; t)$ which now depends explicitly on time. The exact time dependency is, as we will show, irrelevant for the carbon tax path. We do, however, assume that $\frac{\partial \text{MAC}(a; t)}{\partial t} < 0$, to account for the fact that time makes renewable technologies cheaper, and not more expensive.

In fact, the results from the Reference scenario concerning the *shape* of the carbon price path are directly applicable here: the Hamiltonian still does not depend on the cumulative emissions explicitly, which means that

the adjoint variable $\lambda(t)$ is still constant. Therefore, the exponential form given in Eq. 8.3 is still valid, meaning that the carbon tax grows at the same rate as without any technological learning. This confirms the Hotelling rule for this scenario.

Learning over Time does influence the carbon tax, but only through a reduction of the *initial* carbon tax, p_0 , since all subsequent reductions will only become cheaper.

8.1.3. LEARNING BY DOING

The above results are not applicable any more for the learning by doing component of technological learning (also called endogenous technological learning, or induced technological change). In this section, we provide an analytical proof that the optimal carbon tax grows at a strictly *lower than exponential* rate when we include learning by doing in the MAC curve. For simplicity, we disregard learning over time. We start with the most general definition of learning by doing, given in section 6.3: the MAC curve now depends on the cumulative abatement, $CA(t)$, which was defined as the cumulative baseline emissions minus the cumulative emissions:

$$CA(t) = \left(\int_0^t B(s) ds \right) - CE(t)$$

The MAC curve is now, in its most general form, given by $MAC(a; CA(t)) := MAC(a; t, CE(t))$, where we assume that

$$\frac{\partial MAC(a; CA(t))}{\partial CA(t)} < 0$$

for all a , meaning that higher cumulative emissions reduce the abatement costs. We can now calculate the adjoint equation $-\dot{\lambda}(t) = -\partial \mathcal{H}(\cdot) / \partial CE(t)$, which now, contrary to the previous section, is not zero any more:

$$\begin{aligned} \dot{\lambda}(t) &= \frac{\partial \mathcal{H}(\cdot)}{\partial CE(t)} \\ &= -e^{-rt} B(t) \left[\frac{\partial}{\partial CE(t)} \left(\int_0^{MAC^{-1}(p(t); CA(t))} MAC(a; CA(t)) da \right) \right] + \lambda(t) B(t) \frac{\partial}{\partial CE(t)} MAC^{-1}(p(t); CA(t)) \end{aligned}$$

It can be shown (see Appendix A.1.1) that this is equal to:

$$\dot{\lambda}(t) = e^{-rt} B(t) \int_0^{MAC^{-1}(p(t); CA(t))} \left(\frac{\partial}{\partial CA(t)} MAC(a; CA(t)) \right) da \quad (8.6)$$

By assumption, $\frac{\partial}{\partial CA(t)} MAC(a; CA(t))$ was strictly negative for all a , which means that in the most general Learning by Doing case,

$$\dot{\lambda}(t) < 0,$$

or, in other words, that the carbon tax grows strictly *less fast* than the exponential Hotelling price path.

This is in contradiction to the results obtained by [Goulder and Mathai \(2000\)](#), which say that even when considering Learning by Doing, the optimal carbon tax should follow an exponential path with the same rate as the discount rate. This is discussed in more detail in section 8.3.

While we have proved that the Hotelling rule does not apply any more when considering learning by doing, we cannot infer any conclusions analytically on how large the difference is between the optimal path and the Hotelling path. For this, we need numerical results, the topic of the following section.

8.2. NUMERICAL

Numerical optimisation allows us to analyse a wide range of scenarios, introduced in chapter 6. To obtain the optimal carbon tax path and the corresponding emission path, we use the Bellman Equation, described in detail in section 7.2.

For every scenario listed below, unless stated otherwise, we use a carbon budget of 30% of the baseline cumulative emissions. This means that over the period 2015-2100, 70% of the business as usual emissions need to be reduced. This can also occur through net negative emissions. We first discuss the *shape* of the carbon tax and emission paths, and compare the optimal carbon tax path with the proxy price path. As explained in section 7.4, this is the price path which is taken to be a Hotelling path (exponentially increasing with 5% a year), and where we calculate the initial carbon tax such that the carbon budget is reached. By comparing the optimal carbon tax path (shown in green in the results below), and the proxy tax path (shown in gray), the necessity of the optimisation becomes clear.

This difference between optimal and Hotelling path is quantified further in section 8.2.7, when we look more closely at the total costs associated to the emission paths.

8.2.1. NUMERICAL CONFIRMATION OF HOTELLING

The first scenario we calculate numerically is also the simplest: we assume a MAC which is constant over time, and no additional constraints in the emission reduction path. As shown in the previous section, we can prove analytically that the Hotelling rule is still valid. In fact, this is confirmed by our numerical results, as shown in Fig. 8.1a: the carbon tax path follows an exponential curve growing at rate 5%, which is the discount rate we use. There is indeed no visible difference between the proxy price path (gray) and the optimal carbon tax path (green).

The same can be said about the Learning over Time scenario, shown in Fig. 8.1b. However, a higher learning rate implies that the carbon tax is lower, since the abatement costs become lower over time. Interestingly, when one knows beforehand how strong the exogenous technological learning is, still leads to a Hotelling path. While in Fig. 8.1b the tax path for the higher learning rate of rate = 0.05 is much lower than the tax path with rate = 0.03, the growth rates of the tax paths are the same: 5% per year. This result was also proved analytically in the previous section, and is shown more clearly when plotting the carbon tax on a log-scale, shown in the Appendix in Fig. B.1.

While the shape of the tax path is unchanged compared to the no-learning scenario, the emission path is changed: in fact, for higher learning rates, the emission reductions are pushed to the future, since then the abatement costs will be lower anyway. This takes quite extreme forms: the emission paths in Fig. 8.1b go all the way to -6, which means that by 2100, we should remove in total 6 times the amount of CO₂ that we are currently emitting in 2015. The technological possibility of this is strongly disputed.

8.2.2. SCENARIO: LEARNING BY DOING

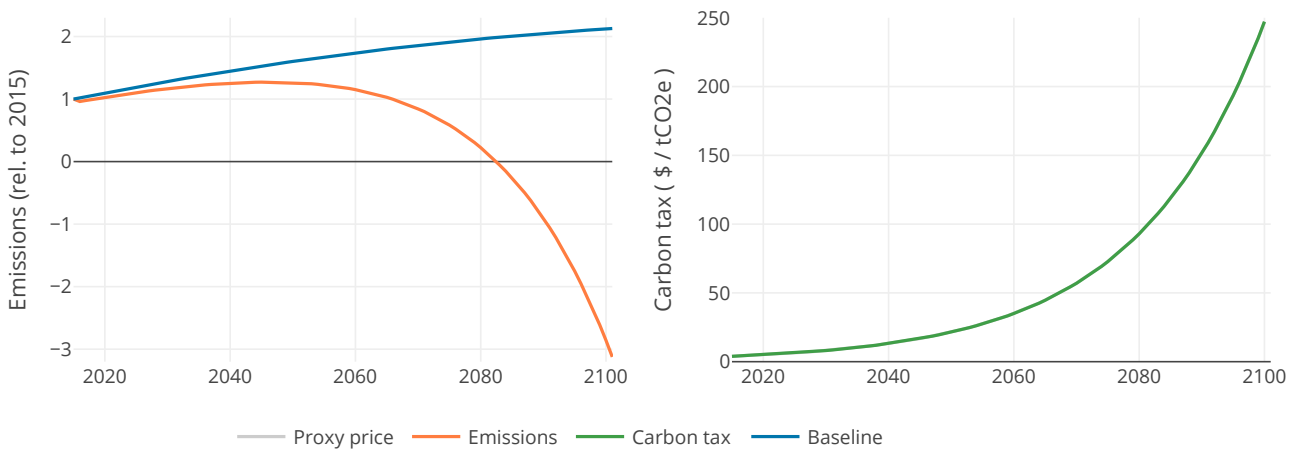
The Learning by Doing scenario implements a more realistic view of technological growth: marginal abatement costs don't just go down over time by themselves, they are now assumed to be reduced as a function of cumulative abatement, or, in other words, as a function of total knowledge. As proved in the previous section, the growth rate of the carbon tax level should be strictly slower than the Hotelling path.

Indeed, as shown in Fig. 8.1c, the carbon tax paths differ significantly from the corresponding proxy Hotelling paths. This effect becomes stronger for higher progress ratios, which is the parameter determining how large the LBD effect is. As can be seen on a log-scale in Fig. B.1, the optimal growth rate is clearly lower than for the Hotelling path.


For extremely high progress ratios, the optimal carbon tax path can even decrease after some time, as shown in Appendix B.2.

8.2.3. SCENARIO: MINIMUM EMISSION LEVEL


In the previous scenario, emission levels still become strongly negative towards the end of the century. For technical and societal reasons, this is considered to be unrealistic. This is addressed in the Minimum emission level scenario, shown in Fig. 8.1d. Indeed, the emission paths are limited at a predefined emission level L , in our case chosen to be 0.25 and -0.25. The optimal carbon tax path follows exactly the proxy Hotelling path, until the point minimum emission level is reached. Since a higher carbon tax will not infer a higher abatement, the carbon tax will stay constant.



(a) Default scenario: linear MAC, no learning.




(b) Learning Over Time: linear MAC, cost reduction purely depending on time.




(c) Learning By Doing scenario: linear MAC, cost reduction depending on cumulative abatement.

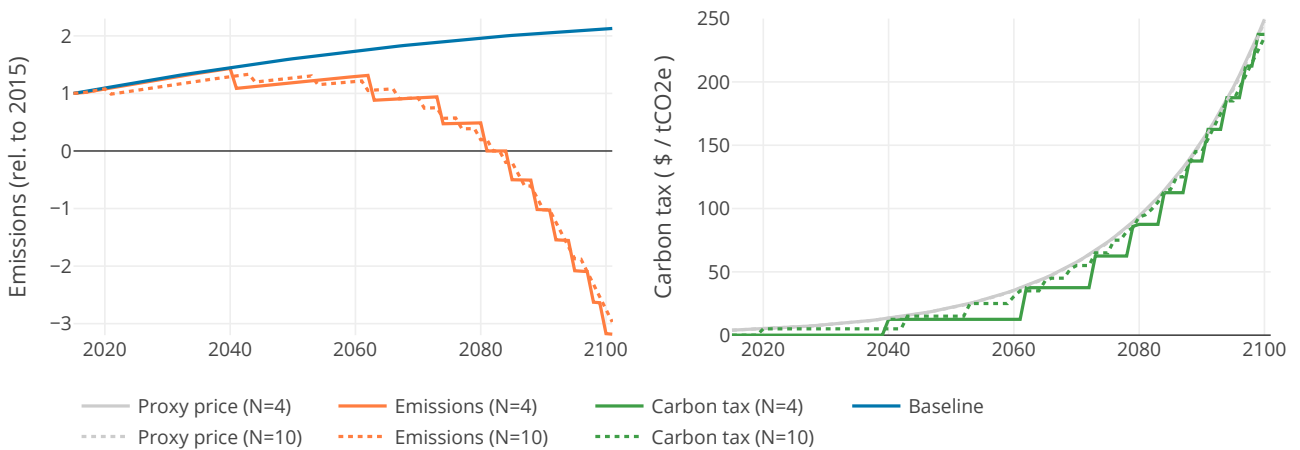
Figure 8.1: For the carbon budget $\alpha = 0.3$, optimal carbon tax paths (right) and corresponding emission paths (left). In gray, the (proxy) Hotelling paths are shown.



(d) Minimum emission level scenario: constraint on the emission level.




(e) Inertia scenario: constraint on the reduction rate.



(f) Discontinuous MAC scenario: no learning or extra constraints.

Figure 8.1: (CONTINUED) For the carbon budget $\alpha = 0.3$, optimal carbon tax paths (right) and corresponding emission paths (left). In gray, the (proxy) Hotelling paths are shown.



(g) Convex MAC scenario: instead of linear MAC, we use a cubic MAC. No learning or extra constraints.

Figure 8.1: (CONTINUED) For the carbon budget $\alpha = 0.3$, optimal carbon tax paths (right) and corresponding emission paths (left). In gray, the (proxy) Hotelling paths are shown.

It should be noted that a sustained constant *positive* emission level might not be optimal on longer time scales, since atmospheric CO_2 levels will continue to rise after the year 2100, therefore exceeding the carbon budget.

8.2.4. SCENARIO: INERTIA

The second sort of constraint we implement is a constraint on the emission reduction rate: the Inertia scenario. The optimal carbon tax is now made of two parts: the first part follows an exponential growth as predicted by the Hotelling rule, but with a lower initial carbon tax, which is followed by a part where the inertia constraint is binding. The emissions are reduced at their maximal reduction rate, resulting in a linear emission path. The corresponding necessary carbon tax is then slightly concave, as a result of the concavity of the baseline. This is shown in Fig. 8.1e. The stronger the inertia constraint, the earlier this transition occurs.

8.2.5. SCENARIO: COMBINING LEARNING AND CONSTRAINTS

In this last scenario, we combine learning by doing with the inertia constraint and a minimum emission level. The optimal carbon tax, shown in Fig. 8.2 shows three regimes: the first, unconstrained, part yields a carbon tax increasing slightly more slowly than the Hotelling path due to learning by doing. In the second part, the emission path is constrained by the inertia: the tax path shows a similar shape as in the Inertia experiment. Finally, when the emissions are limited by the no-negative-emissions limit, the carbon tax decreases due to technological learning, which was not present in the Min. Emissions level scenario (Fig. 8.1d). Instead of varying parameter values, we show the effect of a change in carbon budget: we use a budget of 20%, 30% and 40% of baseline emissions. This doesn't have any significant effect on the shape of the tax pathway, but rather on the timing of the different regimes.

8.2.6. DIFFERENT MAC SHAPES

Additionally to changing the technological learning methods and using different constraints, we also investigate the role of the shape of the MAC. While for all previous results, we have used a linear MAC, we show the resulting emission and carbon tax paths for a discontinuous and a convex MAC.

When using a discontinuous MAC, shown in Fig. 6.2, the carbon tax path also becomes, unsurprisingly, discontinuous, as shown in Fig. 8.1f. However, the carbon tax path still follows the Hotelling path (albeit discontinuously).

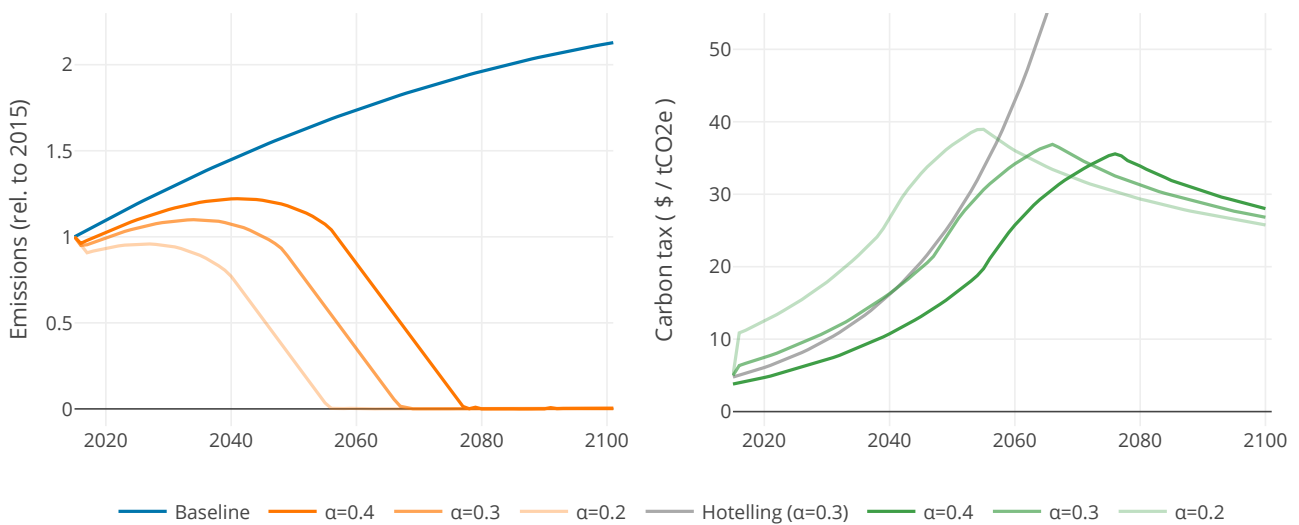


Figure 8.2: Combined scenario: by combining learning by doing, inertia and a minimum emission level, the optimal carbon tax reflects the three domains. Instead of varying parameters, we vary the carbon budget here.

Changing the MAC to a convex (in this case, cubic) function, also shown in Fig. 6.2, the optimal carbon tax path is equal to the proxy Hotelling path: this was also shown in the analytical part in section 8.1, where the optimal carbon tax path follows a Hotelling path independently of the shape of the MAC, as long as there was no learning by doing or extra constraints. The resulting emission path is shown in Fig. 8.1g.

8.2.7. COSTS

When discussing the results of the optimal carbon tax and emission paths, we have until now focused on the *shape* of the carbon tax. We now analyse the difference in *costs* between the optimal and Hotelling paths. By definition, the total costs, as described in section 5.4, for the optimal tax path should be always equal or smaller than the Hotelling path. In this section, we show how much smaller they are.

First of all, for the default and the learning over time scenarios, the costs are equal, since the carbon tax paths are equal.

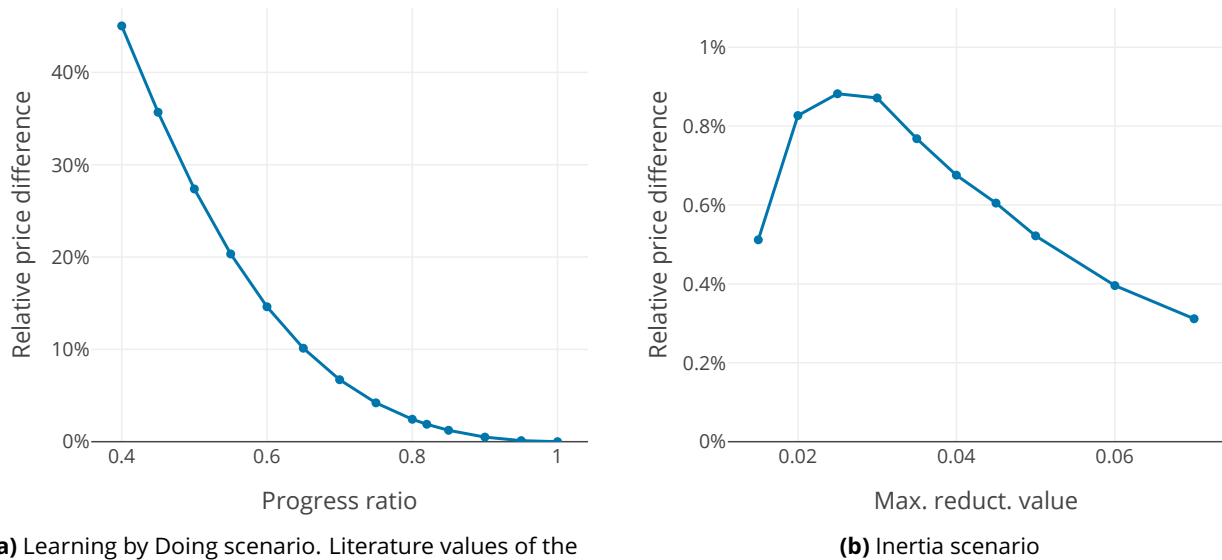
MINIMUM EMISSION LEVEL

In the minimum emission level scenario, the optimal carbon tax paths differs from the proxy Hotelling path (see Fig. 8.1d). However, since they are equal until the emission level constraint comes in, the total costs are equal to the Hotelling costs. In fact, as shown in Eq. (5.4), the costs only depend on the effective abatement, resulting from the carbon tax, and not on the carbon tax itself. Since the effective emission paths from the optimal tax path and the Hotelling path are equal, the total costs are also equal.

LEARNING BY DOING

When considering learning by doing, the optimal path results in a 1.9% lower total discounted costs for a progress ratio of 0.82, all the way to a 10.1% reduction for the higher end of the progress ratio spectrum at 0.65. As shown in Fig. 8.3a, this trend continues for even higher progress ratios.

To test the sensitivity of the cost reduction in the learning by doing scenario to the shape of the MAC, we performed the same calculations using a convex MAC. The results are shown in Appendix B.3. While the cost difference is less than when using a linear MAC, the general behaviour is similar: the smaller the progress ratio, the stronger the price difference between optimal and proxy costs.



(a) Learning by Doing scenario. Literature values of the progress ratio are between $\rho = 0.65$ and $\rho = 0.95$, with most likely value $\rho = 0.82$.

(b) Inertia scenario

Figure 8.3: Costs reduction using optimal carbon tax path compared to the corresponding proxy Hotelling path, for the Learning by Doing and the Inertia scenarios. Note that the range of the y-axes are different in the left and right figures.

INERTIA

The costs for the Inertia scenario are also lower than the corresponding Hotelling paths. However, due to the constraint imposed by the inertia, there is less freedom in influencing the abatement path by choosing a different carbon tax path. The resulting cost reduction is therefore slightly less than 1%. This becomes lower for weaker inertia constraints, since the tax path comes closer to the Hotelling path, and also becomes lower for very stringent constraints, since the freedom in choosing the path is severely reduced. The relative price differences between optimal and Hotelling paths are shown in Fig. 8.3b.

COMBINATION

The cost reduction in the Combination scenario is limited by the cost reduction in the Learning by Doing part and by the reduction from the Inertia part. However, since there is less freedom in choosing the carbon tax path because of the added minimum emission constraint, the price differences will be smaller than in the two previous cases.

8.3. COMPARISON TO GOULDER ET AL

One of the reasons to investigate under which settings the Hotelling rule still applies goes back to results obtained by [Goulder and Mathai \(2000\)](#). In their paper, it is shown that when considering induced technological change, and, in particular, learning by doing, the optimal carbon tax path follows an exponential path with rate equal to the sum of the interest rate and the natural rate of removal of CO_2 from the atmosphere, therefore confirming the Hotelling rule. In this section, we will explore their method, provide a more mathematically rigorous derivation of their results and attempt to explain why we obtain significantly different results.

8.3.1. GENERAL SETTING

First of all, the authors consider an optimal control problem, in which they optimise an abatement path $A(t)$ which reduces the CO₂ concentration compared to a baseline level, while minimising the associated costs. The control variable is therefore $A(t)$, and the state variable the concentration $S(t)$. This concentration is governed by the equation:

$$\dot{S}(t) = -\delta S(t) + E^0(t) - A(t), \quad (8.7)$$

where δ is the natural rate of removal of CO₂ from the atmosphere and $E^0(t)$ the baseline emission rate at time t . The learning by doing element is modelled by introducing an accumulated knowledge variable $H(t)$ which grows at rate α (accounting for exogenous technological change, which we call learning over time) and depending on the abatement level $A(t)$, accounting for the learning by doing:

$$\dot{H}(t) = \alpha H(t) + \Psi(A(t), H(t)), \quad (8.8)$$

where the knowledge growth function Ψ is chosen such that higher abatement induces more knowledge ($\Psi_A(\cdot) > 0$, where the subscript denotes the partial derivative) and the more knowledge there is, the more difficult it becomes to learn ($\Psi_H(\cdot) < 0$).

Finally, the associated costs of abating an amount A given a knowledge level H are given by $C(A, H)$. This is the *abatement cost curve*, which has the following properties: $C_A(\cdot) > 0$, $C_{AA}(\cdot) > 0$, $C_H(\cdot) < 0$ and $C_{AH}(\cdot) < 0$.

This setting is actually very similar to ours, with a few minor differences: we consider absolute amounts of CO₂ instead of concentrations, we use the carbon tax as control variable instead of the abatement level (but, as we have seen, there is a bijective relationship between the two through the MAC: $p = \text{MAC}(a)$) and we incorporate knowledge accumulation directly in the MAC instead of through a separate knowledge function H . This last difference will turn out to be essential. Before we analyse this in more depth, we first solve the problem using Pontryagin's Principle.

8.3.2. SOLVING THE OPTIMAL CONTROL PROBLEM

By combining the equations in the previous section, we obtain the full optimal control problem:

$$\begin{aligned} & \underset{A(t)}{\text{minimise}} \int_0^\infty e^{-rt} C(A(t), H(t)) dt \\ & \text{subj. to } \dot{S}(t) = -\delta S(t) + E^0(t) - A(t), \\ & \quad \dot{H}(t) = \alpha H(t) + \Psi(A(t), H(t)), \\ & \quad S(t) \leq \bar{S} \text{ for } t \geq T. \end{aligned}$$

Here, we also added the constraint that the concentration should have dropped below a certain threshold after some time T . We focus on the time *before* T .

To solve the above problem, we will apply Pontryagin's Principle, introduced in section 3.1.1. We start by defining the Hamiltonian:

$$\begin{aligned} \mathcal{H}(s, h, \lambda_1, \lambda_2, a) & := r(s, h, a) + \mathbf{f}(s, h, a) \cdot (\lambda_1, \lambda_2)^T \\ & = -e^{-rt} C(a, h) + \lambda_1(-\delta s + E^0 - a) + \lambda_2(\alpha h + \Psi(a, h)) \end{aligned}$$

Here, s , h , λ_1 , λ_2 and a are functions from \mathbb{R} to \mathbb{R} . Note that we included a minus sign in front of the cost function to account for the fact that we are minimising the costs, or maximising the negative of the costs. It is customary in economics (Hanley et al., 1997) to transform the Hamiltonian to the *current value* Hamiltonian. Since the Pontryagin's Principle is stated using the *present value* Hamiltonian (the regular Hamiltonian), we have to rewrite \mathcal{H} . We start by factoring out e^{-rt} :

$$\mathcal{H}(s, h, \lambda_1, \lambda_2, a) = e^{-rt} \left(-C(a, h) + \lambda_1 e^{rt} (-\delta s + E^0 - a) + \lambda_2 e^{rt} (\alpha h + \Psi(a, h)) \right)$$

Now, by introducing the new costate variables $\tau(t) := -\lambda_1(t)e^{rt}$ and $\mu(t) := \lambda_2(t)e^{rt}$, we arrive at an expression similar to the Hamiltonian introduced in [Goulder and Mathai \(2000\)](#):

$$\begin{aligned}\mathcal{H} &= e^{-rt} \left(-C(a, h) - \tau(-\delta s + E^0 - a) + \mu(\alpha h + \Psi(a, h)) \right) \\ &= e^{-rt} \widehat{\mathcal{H}},\end{aligned}$$

where $\widehat{\mathcal{H}}$ is the current value Hamiltonian.

ADJOINT EQUATIONS

Let us now assume that $A^*(\cdot)$ is optimal for Eq. (8.7) and Eq. (8.8), and $S^*(\cdot)$ and $H^*(\cdot)$ the corresponding trajectories. Then, from Theorem (1), we know that there exists a function $\lambda^* : [0, \infty) \rightarrow \mathbb{R}^2$ (where $\lambda^* = (\lambda_1^*, \lambda_2^*)^T$), such that, at least for all $t < T$:

$$\dot{\lambda}^*(t) = -\nabla_{(s, h)} \mathcal{H}(S^*(t), H^*(t), \lambda_1^*(t), \lambda_2^*(t), A^*(t)),$$

or in terms of the individual components:

$$\begin{cases} \dot{\lambda}_1^*(t) = -\frac{\partial}{\partial s} \mathcal{H}(S^*(t), H^*(t), \lambda_1^*(t), \lambda_2^*(t), A^*(t)) \\ \dot{\lambda}_2^*(t) = -\frac{\partial}{\partial h} \mathcal{H}(S^*(t), H^*(t), \lambda_1^*(t), \lambda_2^*(t), A^*(t)) \end{cases}$$

$$\Leftrightarrow \begin{cases} \dot{\lambda}_1^*(t) = e^{-rt} \lambda_1^*(t) e^{rt} \delta \\ \quad = \delta \lambda_1^*(t) \\ \dot{\lambda}_2^*(t) = e^{-rt} C_h(A^*(t), H^*(t)) + \lambda_2^*(t) (\alpha + \Psi_h(A^*(t), H^*(t))) \end{cases}$$

By substituting τ and μ , and noting that $\dot{\lambda}_1^*(t) = (-\dot{\tau}(t) + r\tau(t))e^{-rt}$, we obtain first:

$$\dot{\tau}^*(t) = (r + \delta)\tau^* \tag{8.9}$$

and second:

$$\dot{\mu}^*(t) = r\mu^*(t) + C_h(A^*(t), H^*(t)) + \mu^*(t) (\alpha + \Psi_h(A^*(t), H^*(t)))$$

OPTIMALITY CONDITION

Theorem (1) also states that under the same assumptions that $A^*(\cdot)$ is optimal, we have:

$$\mathcal{H}(S^*(t), H^*(t), \lambda_1^*(t), \lambda_2^*(t), A^*(t)) = \max_{a \in \mathcal{A}} \mathcal{H}(S^*(t), H^*(t), \lambda_1^*(t), \lambda_2^*(t), a),$$

where \mathcal{A} is the admissible set of controls. This implies another necessary condition: $\frac{\partial \mathcal{H}}{\partial a} \Big|_{a=A^*(\cdot)} = 0$:

$$\begin{aligned} -e^{-rt} C_a(A^*(t), H^*(t)) + \tau(t) e^{-rt} + \mu(t) e^{-rt} \Psi_a(A^*(t), H^*(t)) &= 0 \\ \Leftrightarrow -C_a(A^*(t), H^*(t)) + \tau(t) + \mu(t) \Psi_a(A^*(t), H^*(t)) &= 0 \end{aligned} \tag{8.10}$$

8.3.3. OPTIMAL CARBON TAX

In the previous derivations, we have only considered the optimal abatement levels $A^*(t)$. However, the main tool available to policy makers is the carbon tax. [Goulder and Mathai](#) argue that τ , as defined above, is equal to “the shadow value of a small additional amount of CO_2 ”, since it is the costate variable associated with the change in CO_2 amount, Eq. (8.7). In fact, in an ideal competitive economy, the carbon tax should be equal to the shadow

cost of additional CO₂ emission (which is equivalent to saying that it is equal to the marginal benefit of one extra unit of abatement).

When disregarding knowledge growth ($\Psi(\cdot) = 0$), we see indeed from Eq. (8.10) that $C_a(\cdot) = \tau(t)$, which means that τ , defined as the shadow cost of extra CO₂ emissions, is equal to the marginal abatement cost $C_a(\cdot)$, which in our definition is equal to the carbon tax, since $p = \text{MAC}(a)$.

We can combine this knowledge with the fact that from Eq. (8.9), we see that τ grows exponentially at rate $r + \delta$: $\tau^*(t) = \tau^*(0)e^{(r+\delta)t}$. If we assume that τ^* is equal to the optimal carbon tax, then this result would indeed confirm the Hotelling rule for the learning by doing setting.

This result contradicts our analytical and numerical findings, that learning by doing does indeed influence the growth rate and shape of the carbon tax. In the next section, we argue why we believe τ does in fact not represent carbon tax in general in a learning by doing setting.

8.3.4. SHADOW COSTS OR CARBON TAX

In a competitive economy, the optimal carbon tax should indeed be equal to the marginal benefit of abating one unit more. However, while abating one unit more indeed yields a benefit τ , the associated costs are also reduced, since the available knowledge H will be increased. This is not taken into account when saying that the carbon tax is equal to $\tau^*(t)$. In fact, the shadow cost of CO₂ does indeed depend in a non trivial way on the cost function C . This is only hidden when introducing an extra state variable H , when part of the true shadow cost of CO₂ is taken into the corresponding costate variable μ .

For this reason, throughout this thesis, we will always define the carbon tax as begin equal to the marginal abatement cost, and therefore avoiding the ambiguous shadow cost definition.

8.4. IMPLICATIONS AND IMPACT

In section 8.1, we proved the optimality of an exponentially increasing carbon price path when there is no technological learning, or under learning over time (exogenous technological growth), confirming the Hotelling rule. These results were supported by numerical results in section 8.2.

However, more realistic scenarios yield different carbon tax paths than predicted by the Hotelling rule. In fact, including learning by doing or an inertia constraint makes the Hotelling path strictly suboptimal. Although the price difference between optimal and Hotelling paths is rather small for most realistic parameter ranges (at most about 5% cost reduction), the main outcome of this chapter is the change in *shape* of the optimal path, compared to the corresponding Hotelling path.

The changes in the *shape* of the price path originate directly from the fact that by definition, the shape of a Hotelling path is independent of the conditions and imposed constraints. For instance, in the Minimum Emission Level scenario (Fig. 8.1d), the Hotelling price still increases when the emissions have reached their minimum level, essentially making the price increase ineffective. Since the societal costs are calculated as function of the achieved emission reduction, and not the carbon price, an ineffectively high carbon price has no effect on the total costs. However, the Hotelling path gives a wrong perception of the future climate policy burden — a problem that is solved when optimising the full carbon price path.

To conclude our analysis of the Hotelling rule, using a Hotelling carbon price path is only valid in a few simple settings. In general, minimising the full carbon price path is required to give an accurate and realistic representation of the required mitigation effort over time.

Chapter 9

Application of stylised setting: BECCS to the Future

Being able to calculate the optimal carbon tax and associated emission path in a stylised setting has already proved to be useful, as shown in the previous chapter. However, the model we have developed lends itself to uses beyond purely calculating the shape of the optimal carbon tax: in this chapter, we show a practical application of the simple model. Specifically, we use it to analyse the sensitivity of various policy indicators to the choice of the discount rate.

This part of the thesis concerns the work I did in collaboration with a research group in Milan, then part of FEEM, now part of the European Institute on Economics and the Environment EIEE.

9.1. MOTIVATION

In economics, the value of a future asset is always *discounted*: the underlying principle is that a given asset is worth less to us in the future than if we were to have it right now. The main reasons for this are generally economic growth (future generations are expected to be richer than we) and uncertainty premiums.

When designing a climate policy, mitigation efforts have to be distributed over time. Just like with monetary assets, abatement burden is discounted: the further costs for climate action occur in the future, the less weight is given to them. Historically, the global average growth rate is typically around 2 or 3%. However, in most climate policy Integrated Assessment Models (IAMs), which are used in guiding the national and international climate policy debates, the discount rate is taken to be around 5%¹.

The discount rate has already been the subject of many discussions and research in the context of cost-benefit analysis for climate policy. In these sort of analyses, the costs of reducing emissions are weighted against the benefits to society over a longer period of time (typically until 2100) of living in a world with less CO₂ in the atmosphere. In this context, it is clear that the discount rate is a key factor in the outcomes: a higher discount rate favours short term monetary gains by burning fossil fuels, while a lower discount rate gives more importance to long-term health and societal benefits of a cleaner atmosphere.

Since the Paris Agreement in 2015, focus has shifted from cost-benefit analysis to cost-effectiveness analysis: a common, global goal was agreed upon (not exceeding 2°C), and the question now resides primarily in *how* this goal can be reached in a cost-effective way.

Moreover, as mentioned in chapter 6, recently, the possible use of negative emission technologies is becoming more and more important. While promising, its full potential is still unclear. However, in most calculated scenarios in which the Paris agreement is reached in 2100, there comes a moment at which negative emissions become so important that the *net* emissions become fully negative, meaning that there is more CO₂ captured from the atmosphere than is being released. Because of this importance, some of the policy indicators explicitly mentioned in the Paris Agreement concern the amount of net negative emissions necessary in an abatement

¹See http://iamcdocumentation.eu/index.php/IAMC_wiki for model documentation and discount rate choices

scenario.

In this chapter, we investigate the sensitivity of key policy indicators like the initial carbon tax, the year at which the net emissions per year become fully negative and the amount of net-negative emissions, to changes in the discount rate.

9.2. METHOD

In order to calculate the above mentioned sensitivity, we use our stylised model described in 5 to derive the necessary policy indicators analytically. We also compare the results from the simple model to numerical results from a much more detailed IAM, called WITCH, developed in Milan. This model gives a much more detailed specification of which technologies are available and being used in different regions all over the world.

9.2.1. ANALYTICAL SETTING

The analytical model, as a reminder, is set in a cost-effective setting: it minimises the costs associated to an emission path, necessary to reach a chosen carbon budget. The amount of emission reduction per year are determined through the Marginal Abatement Cost (MAC) curve: the higher the carbon tax, the more emissions are reduced per year. In mathematical terms, the model is summarised in Eq. (7.1). When calculating the optimal carbon tax and emission path, we first have to specify the shape of the baseline $B(t)$ and the MAC curve $MAC(a)$.

In order to reduce the complexity of the analytical results, we use a constant baseline $B(t) = b_0$, and, just like in the analytical results in chapter 8, a power-law MAC: $MAC(a) = \beta_0 a^{\beta_1}$. For simplicity, we also disregard any technological learning factor, and therefore keep the MAC constant over time.

The results are obtained as in section 7.1: we calculate the Hamiltonian, and solve the equations given by Pontryagin's Principle.

9.2.2. FINDING ANALYTICAL EXPRESSIONS OF INDICATORS

Since in the analytical setting we disregard any technological growth, we have already shown in section 8.1.1 that the optimal carbon tax path follows the Hotelling rule: given the initial carbon tax, the full tax path is known by exponential increase at the discount rate. This allows us to find analytical expressions for the three key policy indicators described above:

- The initial carbon tax p_0 ,
- The year at which the emissions become net-negative t^* ,
- The relative overshoot OS.

INITIAL CARBON TAX

We have already derived the initial carbon tax for a concave baseline and convex MAC in Eq. (8.5). This can be simplified since we only consider a constant baseline here by taking $b_1 = b_2 = 0$. The expression for p_0 now becomes:

$$p_0 = \beta_0 \left(\frac{(1 - \alpha)rT}{\beta_1(e^{rT} - 1)} \right)^{\beta_1} \quad (9.1)$$

Since this expression still depends on β_0 , which was the price needed to abate 100% of emissions, we actually consider the relative initial carbon tax p_0/β_0 .

ZERO-CARBON EMISSION YEAR

In order to calculate the year at which emissions become net negative, we have to solve the equation $E(t^*) = 0$. In principle, there could be more than one solution to this equation. However, since we assumed that the baseline is constant and the price path increasing, we know that the emission path is strictly decreasing. This

means that once the instantaneous emission level becomes zero, all subsequent emission levels will be net negative.

By using the expression for $E(t)$ defined in Eq. (5.2) and filling in the Hotelling price path $p(t) = p_0 e^{rt}$, we obtain the following equation:

$$\begin{aligned} B(t^*)(1 - \text{MAC}^{-1}(p(t^*))) &= 0 \\ \Leftrightarrow \text{MAC}^{-1}(p(t^*)) &= 1 \\ \Leftrightarrow \text{MAC}(\text{MAC}^{-1}(p_0 e^{rt^*})) &= \text{MAC}(1) \end{aligned}$$

Finally, from this, we derive the general expression for the the zero-carbon emission year:

$$t^* = \frac{1}{r} \ln \left(\frac{\text{MAC}(1)}{p_0} \right) \quad (9.2)$$

where $\text{MAC}(1)$ represents the marginal costs of abating 100% of baseline emissions. Substituting in the expression for p_0 previously found, as well as the convex MAC, the formula becomes:

$$t^* = \frac{\beta_1}{r} \ln \left(\frac{\beta_1 (e^{rT/\beta_1} - 1)}{T(1 - \alpha)r} \right). \quad (9.3)$$

For large enough carbon budget α , this expression can become larger than T : in this case, there is no moment in the considered time span at which the emissions are net negative. However, if the relative carbon budget α satisfies $\alpha < 1 + \frac{(e^{-rT/\beta_1} - 1)\beta_1}{rT}$, we can be sure that t^* is indeed smaller than T .

RELATIVE OVERSHOOT

When in fact there is a period in which emissions are net negative, it is important to calculate the *overshoot*. This gives the total amount of net negative emissions, or, in other words, how much the carbon budget is exceeded before bringing down the cumulative emissions due to the net negative emissions. The relative overshoot OS is calculated as:

$$\text{OS} = \frac{\int_{t^*}^T E(t) dt}{\text{CB}},$$

where CB is the carbon budget. Due to the constant baseline, this budget is equal to $\text{CB} = \alpha b_0 T$. It can be shown (see Appendix A.1.2) that this is equal to:

$$\text{OS} = \frac{1 - \alpha}{\alpha(e^{rT/\beta_1} - 1)} - 1 + \frac{\beta_1}{\alpha r T} \left(\ln \left(\frac{\beta_1 (e^{rT/\beta_1} - 1)}{r T (1 - \alpha)} \right) - 1 \right). \quad (9.4)$$

It should be noted that neither p_0 , t^* nor OS depends on the baseline emission level b_0 .

9.2.3. CONVEXITY OF THE MAC

As mentioned before, the above expressions are obtained using a power-law MAC curve:

$$\text{MAC}(a) = \beta_0 a^{\beta_1},$$

where β_0 and β_1 are positive values. To compare the analytical results to the numerical results from the IAM, we need to find realistic values for these parameters, and specifically for the convexity of the MAC, β_1 , since the parameter β_0 is not present any more in most of the expressions above.

We assume that the convexity of the MAC is between quadratic and cubic, namely using $\beta_1 = 2.8$, as this is the value used in the often used Integrated Assessment Model DICE (Nordhaus, 1993). However, it is also possible to calibrate this parameter using the overshoot values calculated in a number of scenarios using various

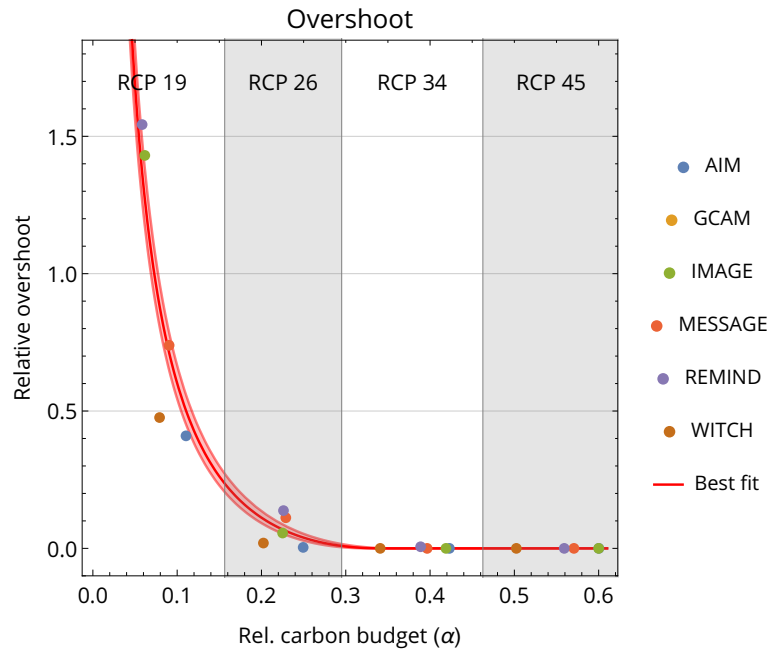


Figure 9.1: Calibration of the MAC convexity parameter β_1 using the overshoot value of various SSP2 scenario's (optimal policy scenarios, SPA0), for different models. The SSP values are fitted to the value of the overshoot derived from our analytical model, shown together with its 90% confidence interval. The discount factor of each scenario is assumed to be 5%.

Integrated Assessment Models, in the so-called Shared Socio-economic Pathways (SSP) database. Specifically, we use the model runs using the SSP2 scenario in an optimal policy context (SPA0), for all available RCPs. We assume that all these model runs have used a discount rate of 5%. First, we convert the RCPs into relative carbon budgets, such that the theoretical overshoot can be fitted through these points. This is shown in Fig. 9.1. Linear regression yields a MAC convexity of $\beta_{1,\text{high}} = 4.6 \pm 0.3$, which we use as upper bound for the MAC convexity. As lower bound, we assume that the MAC is quadratic: $\beta_{1,\text{low}} = 2$, to still implement the convexity of the MAC.

9.2.4. WITCH FOR TECHNOLOGY SPECIFICATION

Since our model is a simplification of reality, it is important to compare its results with a more realistic model. Since this research was performed in collaboration with the European Institute on Economics and the Environment (EIEE, formerly part of the Fondazione Eni Enrico Mattei) in Milan, we use the WITCH IAM developed by researchers of this institute (Bosetti et al., 2009; Emmerling et al., 2016). WITCH is a macroeconomic model used to explore abatement options and scenarios using a wide variety of technological and economic tools.

The WITCH scenarios used in this part are generated in a non-cooperative setting: every country maximises its own welfare, without imposing global cooperation treaties. Every country (or region) does take economic and technological development of other regions into account, thereby iteratively maximising its welfare. This optimisation has one imposed constraint: a chosen carbon budget cannot be exceeded.

Regarding the discount rate, WITCH uses a slightly more elaborate discounting mechanism than simply imposing a fixed rate. In fact, it uses a Ramsey Growth model (Ramsey, 1928), often used for long-term utility maximisation across several generations. It does not take into account market imperfections or market shocks (like crashes). The discount rate is then given by the Ramsey Rule:

$$r = \delta + \eta g_c,$$

where δ is the rate of pure time preference (sometimes called the “impatience term”), g_c the consumption growth rate and η a measure of elasticity of inter-temporal substitutions, which can be interpreted as the inequity aversion parameter over time. While the consumption growth can be calibrated to the average economic growth during the 21st century to $g_c = 1.6\%$, the other two parameters can be varied to effectively obtain different discount rates.

The WITCH model runs are performed with a variety of negative emission technologies: either no Carbon Capture and Storage (CCS) at all, or only BECCS (Bio-Energy with CCS), or BECCS and DAC (Direct Air Capture). However, we focus on the last scenarios, with all CCS technologies available, since this best replicates the analytical model. In fact, CCS technologies are captured in the analytical model by allowing the MAC curve to be defined for higher abatement levels than 100%, at an increasing price due to the convexity of the MAC.

Throughout this chapter, we will vary the (effective) discount rate from 1% to 8%, and use carbon budgets varying from 400 GtCO₂ to 1600 GtCO₂. Since our analytical model only uses relative carbon budgets α , we can calculate α by dividing the absolute carbon budgets by the cumulative emissions of the baseline scenario, which we choose to be 6000 GtCO₂.

9.3. RESULTS

9.3.1. KEY INDICATORS

The value of the analytical expressions of the policy indicators, derived in the previous section, are shown in Fig. 9.2, top row. Interestingly, there is a strong influence of the discount rate on all three policy indicators. For a carbon budget of 400 GtCO₂ (corresponding to the 1.5°C temperature target), the optimal initial carbon tax is doubled for every reduction of the discount rate of 1.5 percentage point. Moreover, by only reducing the discount rate from 5% to 4%, the overshoot goes from 2 to 1.5 times the total carbon budget, all the way to only 1 time the carbon budget for $r = 3\%$. These results are similar for less stringent carbon budgets, although less extreme. The zero-carbon emission year is also moved towards the future when decreasing the discount rate, especially for the higher end of the carbon budgets.

Similar results are obtained from the WITCH model: the discount rate has a strong influence on all three indicators. The general result is that by reducing the discount rate, the need for biomass CCS is strongly reduced, and its timing pushed towards the future: hence the informal name of this chapter, BECCS to the Future.

The previous results can be explained by looking at the analytical model in more detail. In this setting, the total, discounted, costs are minimised. Due to this discounting, costs occurring towards the end of the century are taken less into account in the minimisation process. As a result, the typically more expensive CCS technologies are then relatively cheap, in our present-time view. However, when reducing the discount rate, future costs appear less cheap: it becomes favourable to spread the abatement burden more evenly over the century, in order to avoid the expensive high levels of negative emissions required to reach the carbon budget.

These effects become better visible when showing the emission paths for various discount rates. In Fig. 9.3, it is clear that for higher discount rates, the total net-negative amount is much higher than for lower discount rate. As a consequence, to still reach the carbon budget, emission reductions are much more important in the beginning of the century, therefore requiring a higher initial carbon tax.

This figure also shows why, for stringent carbon budgets, the zero emission year is almost constant for different discount rates: while the reduced overshoot would bring this year more to the future, the total emission path goes down to reach the budget, bringing t^* forward.

9.3.2. INTERGENERATIONAL EQUITY

Due to the uncertain nature of negative emission technologies, obtaining optimal scenarios requiring less CCS is already quite significant. However, there is another reason why it is important not to take the traditional

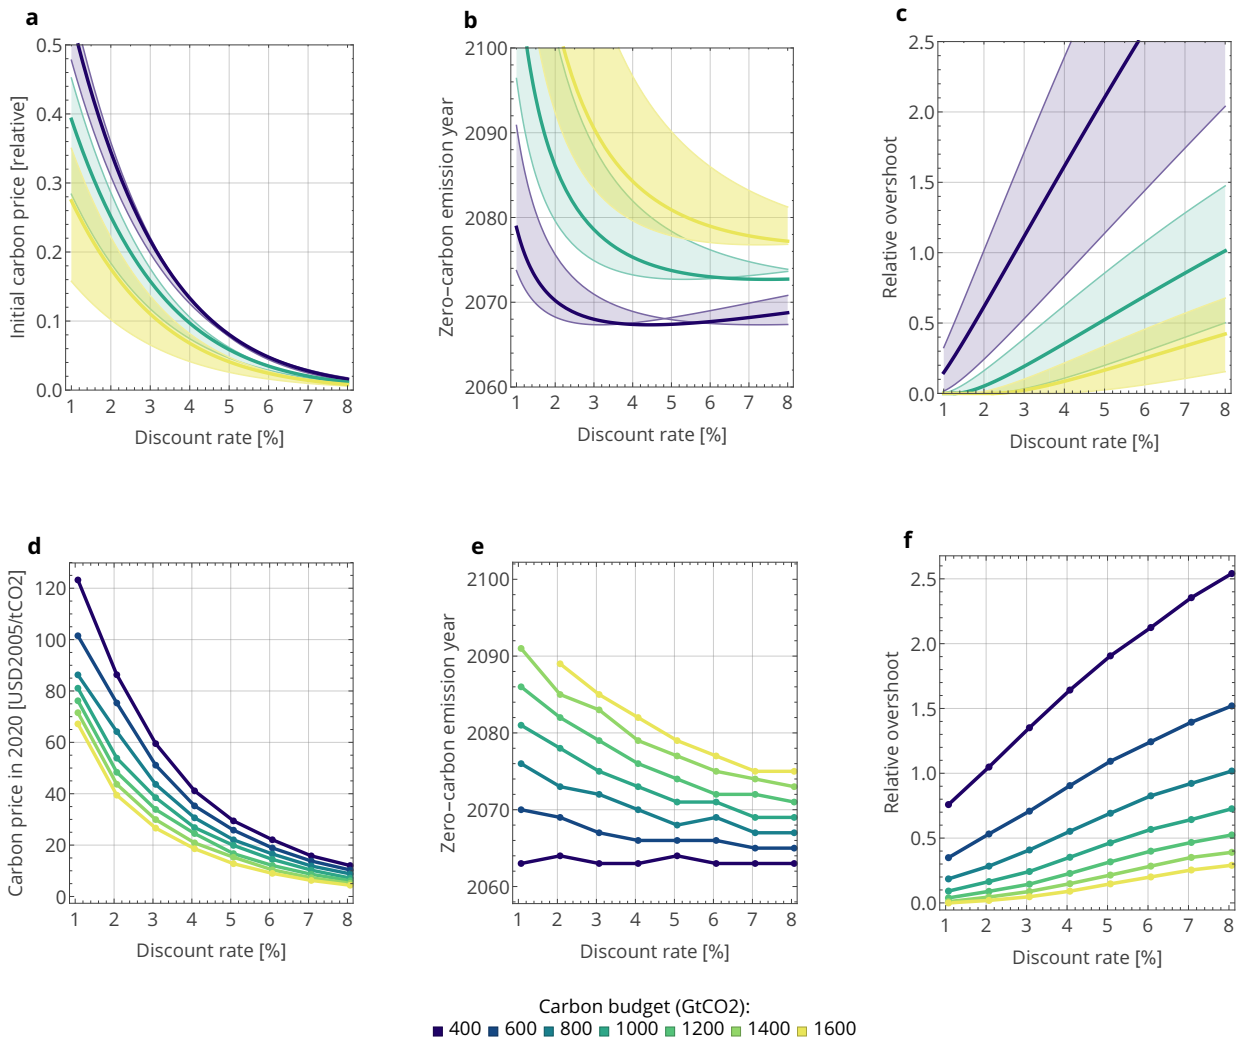


Figure 9.2: Key policy indicators as function of discount rate. The top row represents the results from the analytical model, whereas the bottom row is calculated using the WITCH model. The shaded areas in the analytical results represent the uncertainty range of the convexity of the MAC, with β_1 ranging from $\beta_{1,high} = 4.6$ to $\beta_{1,low} = 2$. The thick solid lines in the top row are calculated using $\beta_1 = 2.8$.

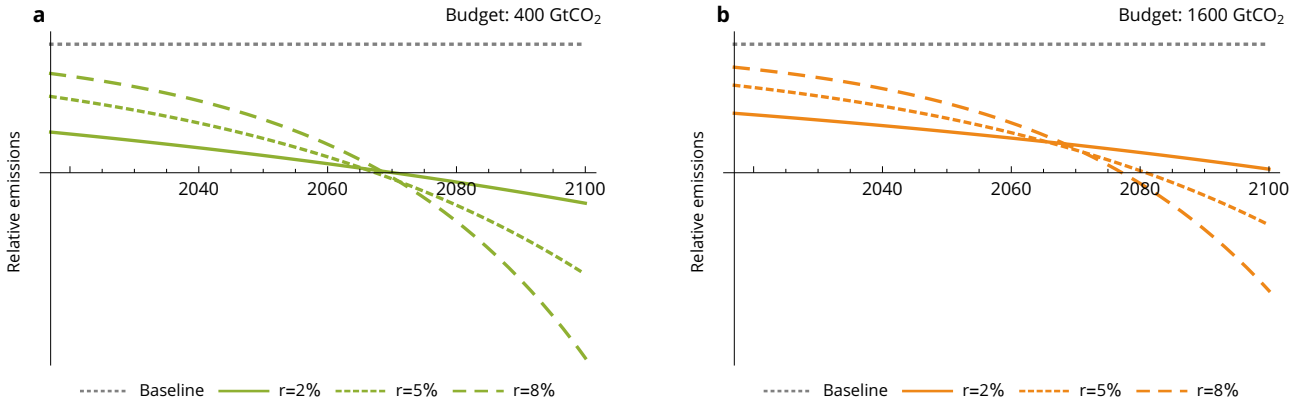


Figure 9.3: Optimal emission paths calculated using the analytical model for three discount rates. Left: a stringent carbon budget of 400 GtCO₂ has been used, compared to a looser budget of 1600 GtCO₂ on the right.

5% discount rate for granted: intergenerational equity. In fact, by the very nature of discounting, we give less weight to costs occurring in the future. This makes sense, since future generations are expected to be richer due to general growth of GDP, but the 5% discount rate implies an even larger burden on future generations than can be accounted for by growing GDP.

To quantify the intergenerational equity, we calculate the abatement costs resulting from the climate policy for three generations, assumed to be living for 30 years: those living in 2020 (2020-2049), those living in 2050 (2050-2079) and those living in 2080 (2080-2109). For each of these periods, the total undiscounted costs are calculated, relative to the GDP, to account for general GDP growth. In perfect intergenerational equity, these costs should be constant for all three generations.

However, as seen in Fig. 9.4, changing the discount rate has a large effect on this intergenerational equity. For a carbon budget of 1000 GtCO₂ and a discount rate of 5%, the costs for the generation living in 2050 are 2.3 times higher as for those living in 2020, and even 3.9 times higher for the 2080 generation. When using 2% as discount rate, these numbers become much lower: the 2050 generation has to pay 1.17 times more than the 2020 generation, and the 2080 generation 1.25 times more.

These results are calculated using the WITCH IAM. They cannot directly be calculated using our analytical model, since we do not include GDP growth in the model.

9.3.3. COMPARATIVE STATICS

In Fig. 9.2, it already becomes clear that the key indicators do strongly depend on the discount rate. We can quantify this by calculating the derivative of p_0 , t^* and OS, given by Eq. 9.1, 9.3 and 9.4. These derivatives are equal to:

$$\begin{aligned} \frac{\partial p_0}{\partial r} &= \frac{-\beta\gamma \left(\beta + e^{\frac{rT}{\beta}} (rT - \beta) \right)}{(1 - \alpha)r^2T} \left(\frac{(1 - \alpha)rT}{\beta \left(e^{\frac{rT}{\beta}} - 1 \right)} \right)^{\beta+1} \\ \frac{\partial t^*}{\partial r} &= -\frac{1}{r^2}\beta - \beta \log \left(\frac{\beta \left(e^{\frac{rT}{\beta}} - 1 \right)}{(1 - \alpha)rT} \right) + rT \left(\frac{1}{e^{\frac{rT}{\beta}} - 1} + 1 \right) \\ \frac{\partial OS}{\partial r} &= \frac{\beta \left(\log((\alpha - 1)(-r)T) - \log \left(\beta \left(e^{\frac{rT}{\beta}} - 1 \right) \right) \right)}{\alpha r^2 T} + \frac{e^{\frac{rT}{\beta}} \left((\alpha - 1)rT + \beta \left(e^{\frac{rT}{\beta}} - 1 \right) \right)}{\alpha \beta r \left(e^{\frac{rT}{\beta}} - 1 \right)^2} \end{aligned}$$

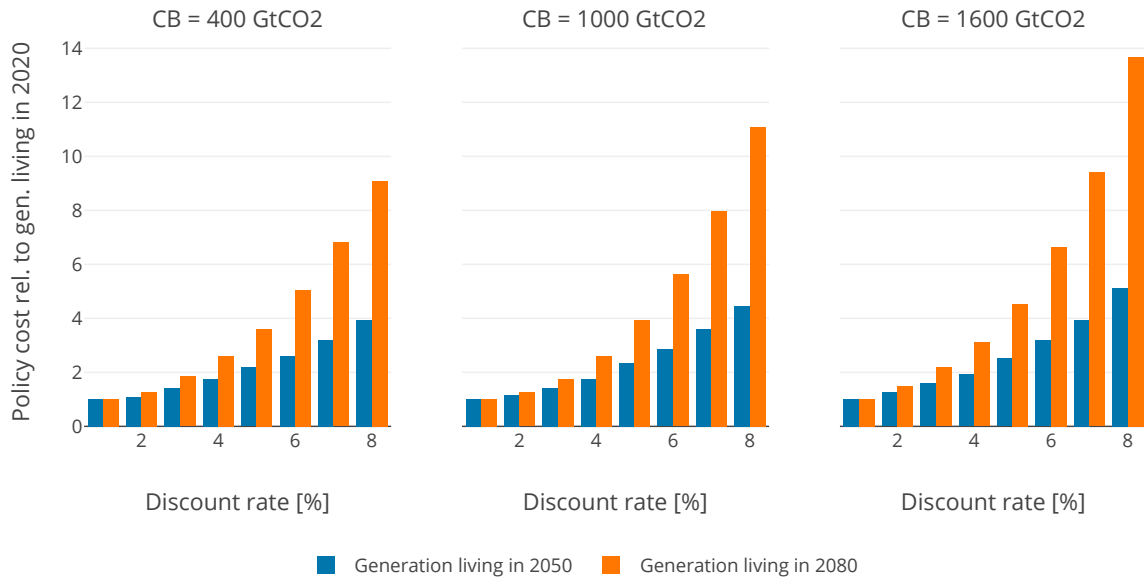


Figure 9.4: Intergenerational equity: the total abatement costs are broken down into three generations, centered around 2020, 2050 and 2080. Here, we compare the costs of the generations living in 2050 and 2080 to those living in 2020, for different discount rates and carbon budgets. These results are calculated using the WITCH IAM.

While the sign of $\frac{\partial p_0}{\partial r}$ is always negative, this is more ambiguous for the two other quantities. However, for most values of the discount rate r and the carbon budget α , $\frac{\partial t^*}{\partial r}$ is negative (except for very high discount rates and very low α), and $\frac{\partial OS}{\partial r}$ is positive (except for very low discount rates and very high α), as shown in Appendix B.4.

By evaluating the above derivatives around 5%, we obtain comparative statics of our analytical model, which give an idea of how much a quantity is changed when increasing the discount rate by 1 percentage point. These values are given in Table 9.1.

Quantity	400 GtCO ₂	1000 GtCO ₂	1600 GtCO ₂	unit (per percentage point diff. in r)
p_0	-51%	-51%	-51%	% change
t^*	0.2	-1.1	-2.5	years
OS	0.48	0.17	0.08	relative to carbon budget

Table 9.1: Comparative statics with respect to the discount rate r .

9.4. DIFFERENT FORMS OF BASELINE AND DISCOUNTING

All previous results have been calculated using a constant baseline. While this strongly simplifies the analytical results, a concave baseline is considered to be much more realistic. In this section, we analyse the effect of using this more realistic baseline. Moreover, we calculate the effect of using hyperbolic discounting, instead of exponential discounting, in section 9.4.2.

9.4.1. QUADRATIC BASELINE

We now assume a quadratic form for the baseline, and keep the same power law form for the MAC as before:

$$B(t) := b_0 + b_1 t + b_2 t^2 \quad (9.5)$$

Typically, we will use a concave baseline ($b_1 > 0, b_2 < 0$), but any other combination of parameters is valid.

Just like in the constant baseline case, we start by evaluating the initial carbon tax p_0 . This was already calculated in Eq. (8.5):

$$p_0 = \beta_0 \left(\frac{(1 - \alpha) (b_0 T + \frac{1}{2} b_1 T^2 + \frac{1}{3} b_2 T^3)}{\left(\frac{b_0 \beta_1}{r} - \frac{b_1 \beta_1^2}{r^2} + \frac{2b_2 \beta_1^3}{r^3} \right) (e^{rT/\beta_1} - 1) + e^{rT/\beta_1} \left(b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1}{r} T^2 - \frac{2\beta_1^2}{r^2} T \right) \right)} \right)^{\beta_1}$$

In a similar fashion, we obtain an expression for the time when net emissions are negative, t^* , by filling the previously found initial carbon tax into Eq. 9.2 and using $\text{MAC}(1) = \beta_0$ (details can be found in Appendix A.1.2):

$$t^* = \frac{\beta_1}{r} \ln \left(\frac{\left(\frac{b_0 \beta_1}{r} - \frac{b_1 \beta_1^2}{r^2} + \frac{2b_2 \beta_1^3}{r^3} \right) (e^{rT/\beta_1} - 1) + e^{rT/\beta_1} \left(b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1}{r} T^2 - \frac{2\beta_1^2}{r^2} T \right) \right)}{(1 - \alpha) (b_0 T + \frac{1}{2} b_1 T^2 + \frac{1}{3} b_2 T^3)} \right) \quad (9.6)$$

A quick test to assess the validity of this solution is by using $b_1 = b_2 = 0$, in which case we should be back in the constant baseline scenario. Indeed, by substituting these parameters in Eq. 8.5 and 9.6, we do obtain the expressions from Eq. 9.1 and 9.3.

Like in the constant baseline case, we can also calculate the closed form expression for the overshoot. However, this expression becomes very large and difficult to interpret, which is why we only show it in detail in Appendix A.1.2.

By using the same parameters as in section 9.3, we can again plot the sensitivity of p_0 , t^* and OS as function of the discount rate. This is shown for three carbon budgets in Fig. 9.5. These results are very similar to those shown in Fig. 9.2, obtained using a constant baseline. For this reason, we conclude that for this application, we can stick to the linear baseline, since the analytical expressions are much simpler.

9.4.2. HYPERBOLIC DISCOUNT RATE

During the last decades, many debates have emerged on the validity of exponential discounting (Dasgupta, 2008). An often cited alternative is *hyperbolic discounting*. In its simplest form, the discount factor $D(t)$ is given by:

$$D(t) = \frac{1}{1 + rt}. \quad (9.7)$$

In this section, we will again calculate analytically the carbon tax path, the moment of net negative emissions, and the fraction between net positive and net negative emissions, using this new hyperbolic discount factor. First of all, the optimal control problem now becomes:

$$\begin{aligned} & \underset{p(t)}{\text{minimise}} \int_0^\infty \frac{1}{1 + rt} b(t) \left(\int_0^{\text{MAC}^{-1}(p(t); \cdot)} \text{MAC}(a; \cdot) da \right) dt \\ & \text{subj. to } \dot{\text{CE}}(t) = B(t) \left(1 - \text{MAC}^{-1}(p(t); \cdot) \right), \\ & \text{CE}(T) = \text{CB} \end{aligned}$$

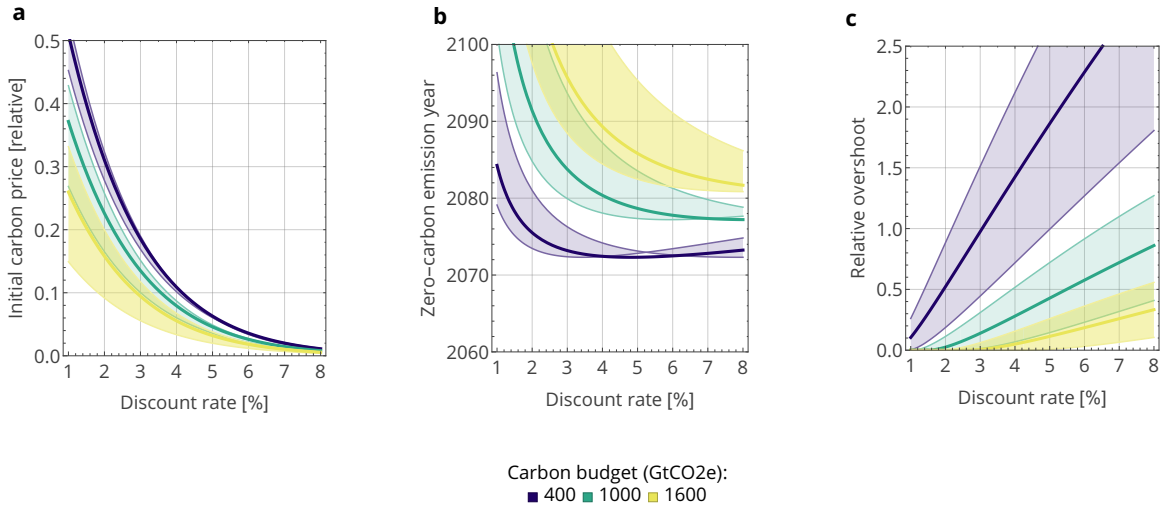


Figure 9.5: Key policy indicators as function of discount rate, now using a **concave baseline**, instead of the linear one used to make Fig. 9.2.

From the optimality condition given by Pontryagin's principle, we find a very similar expression as with exponential discounting:

$$\begin{aligned}
 -\frac{1}{1+rt}p(t) + \lambda(t) &= 0 \\
 \Leftrightarrow p(t) &= (1+rt)\lambda(t)
 \end{aligned} \tag{9.8}$$

The major difference now is that in cases where the shadow costs $\lambda(t)$ are constant, the carbon tax is not exponential any more, but linear. In fact, this result holds for any kind of continuous differentiable discounting function: when $\lambda(t)$ is constant, the carbon tax is proportional to the multiplicative inverse of the discounting function. If we assume that there is no learning by doing, then again the Hamiltonian will be independent of the state variable $CE(t)$, which means that $\dot{\lambda}(t) = 0$. This implies a linear optimal carbon tax.

We now derive analytical expressions for p_0 , t^* and OS as previously defined. For simplicity, we assume a constant baseline.

In general, we can calculate the initial carbon tax by solving Eq. (7.3), but now instead of using $p(t) = p_0 e^{rt}$, we use $p(t) = p_0(1+rt)$. By assuming that the baseline is given by $B(t) = b_0$, and the MAC by $MAC(a) = \beta_0 a^{\beta_1}$, the expression for the initial carbon tax is now:

$$p_0 = \beta_0 \left(\frac{rT(1-\alpha)(1+\beta_1)}{((1+rT)^{1+1/\beta_1} - 1)\beta_1} \right)^{\beta_1}$$

Next, we can calculate the initial time of going negative. To find this, we have to solve the equation:

$$\begin{aligned}
 B(t^*)(1 - MAC^{-1}(p(t^*))) &= 0 \\
 \Leftrightarrow MAC^{-1}(p(t^*)) &= 1 \\
 \Leftrightarrow t^* &= \frac{MAC(1) - p_0}{p_0 r}
 \end{aligned} \tag{9.9}$$

By using the expression for p_0 previously found, we can obtain the closed form expression for t^* .

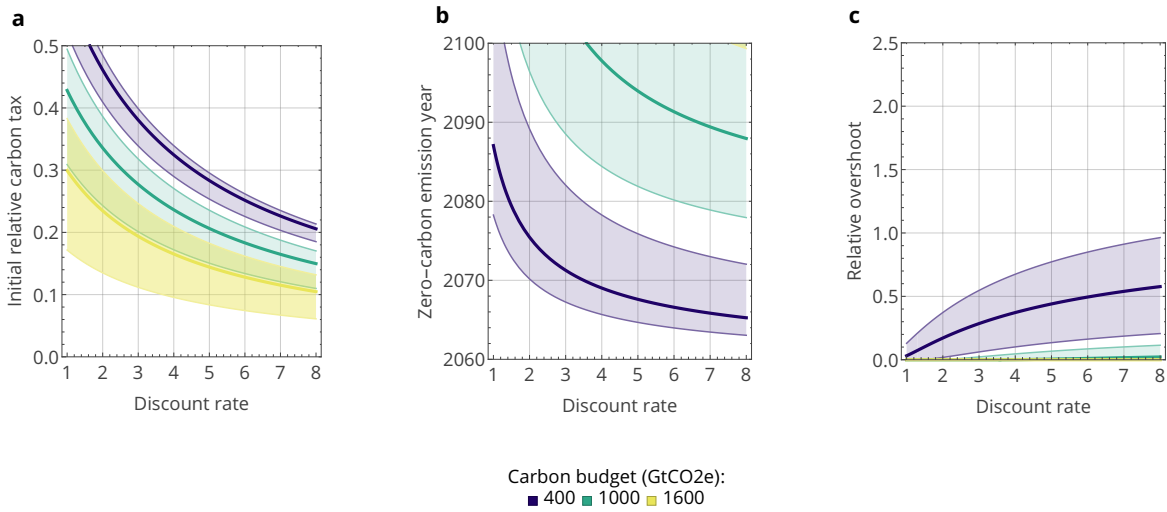


Figure 9.6: Key policy indicators as function of discount rate, using **hyperbolic discounting** instead of exponential discounting used to make Fig. 9.2.

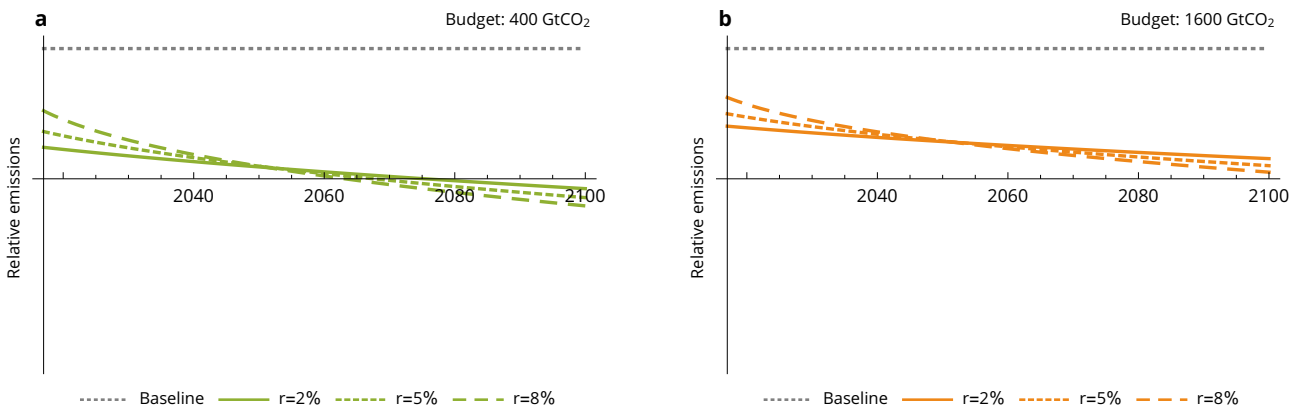


Figure 9.7: Optimal emission path using hyperbolic discounting, using a MAC convexity of $\beta_1 = 2.8$ and two carbon budgets: 400 GtCO₂ and 1600 GtCO₂.

Similarly, the overshoot becomes becomes, after a few algebraic simplifications:

$$OS = \frac{1}{\alpha T} \left(\frac{(1 - \alpha)T}{(rT + 1)^{-\frac{1}{\beta_1} - 1} - 1} - \frac{1}{r(\beta_1 + 1)} \left(\frac{\beta_1 \left((rT + 1)^{\frac{1}{\beta_1} + 1} - 1 \right)}{(1 - \alpha)(\beta_1 + 1)rT} \right)^{\beta_1} + \frac{1}{r} + T \right)$$

These expressions are again evaluated for different carbon budgets α as function of the discount rate, shown in Fig. 9.6. Contrary to the convex baseline case, there are now significant differences with the original results in Fig. 9.2. Especially the zero-carbon emission year occurs much later than for an identical r and α when using exponential discounting. This is also reflected in the overshoot, which is much lower, and even zero for the 1600 GtCO₂ carbon budget.

The reason for this effect becomes clear when looking at the optimal emission paths for different discount rates and carbon budgets, shown in Fig. 9.7. In fact, the emission paths are much flatter. While hyperbolic discounting is almost equal to exponential discounting on the short term, the difference becomes much larger

on the long term. That is also the reason why hyperbolic discounting is preferred by some economists: long term investments are discounted much less. In the context of climate policy, future costs weigh much more than when using exponential discounting.

In this sense, while the key policy indicators are hardly sensitive to the discount rate r in hyperbolic discounting, purely the fact that hyperbolic discounting is used has a huge effect on optimal carbon taxes and emission paths.

9.5. IMPLICATIONS AND IMPACT

The three policy indicators discussed in this chapter, namely the initial carbon price, the zero-carbon emission year and the budget overshoot, are indeed sensitive to changes in the discount rate. An adjusted discount rate has large societal and policy implications, especially for the initial carbon price and the relative overshoot. In our simple model, changing the discount rate to 3% would require an overshoot only half as big as when using a 5% discount rate, which is important in the light of the current uncertainties in the feasibility of large scale negative emissions. Such a lower discount rate directly implies a shift of mitigation effort closer to the present, as shown by a doubling in initial carbon tax. Bringing mitigation effort forward also increases intergenerational equity: even when taking GDP growth into account, the cost burden is then better spread over the generations.

These results were obtained using a simplified, stylised analytical model. However, the results are in good accordance with the more detailed WITCH IAM. One could then question the necessity of the analytical results, if larger IAMs already exist. A big advantage of a simple model is that it is easier to understand and to interpret. The model could be more easily explained to non-experts, like policy makers and politicians.

In conclusion, since a 2% or 3% discount rate is closer to true projected GDP growth, and given the societal effects just described, we argue that the typical choice of a 5% discount rate in IAMs should be reconsidered to reduce the projected required amount of negative emissions and increase intergenerational equity.

Chapter 10

Wrap-up of part I

10.1. CONCLUSION

In this thesis, we have analysed optimal carbon pricing by creating a simple model. This model contains three main steps: first, the amount of emission reductions obtained when choosing a carbon price is calculated through the inverse of the Marginal Abatement Cost curve (MAC). Second, the costs associated with a carbon price path are defined as the integral over time of the discounted area under the MAC. These total costs are then minimised while still keeping the total emissions below a given carbon budget constraint. Third, the minimisation is subject to a last constraint, the dynamics of the system: the emissions are reduced relative to the baseline.

Using this simple model, we found that the validity of the Hotelling rule is limited to the No Learning and Learning over Time scenarios. More generally, we have shown analytically that a Hotelling price path is indeed optimal for any scenario without constraints on the emission paths and where the MAC does not depend on the cumulative emissions, as long as the MAC is strictly increasing. In all the other scenarios we considered, the optimal path deviates from the Hotelling path: in these cases, the optimal carbon price grew less fast than the exponential path.

Another observation that follows from the optimal carbon price paths we calculated is that IAMs reporting a non-exponentially growing carbon price can not be criticised purely because the carbon price path is not exponential. To be more precise, a carbon price path obtained in a cost-effectiveness setting does not necessarily have to be exponentially increasing. In fact, for most realistic models, we should expect a carbon price path that does indeed deviate from the traditional Hotelling path.

However, we also found that the cost difference between optimal price paths and the corresponding proxy Hotelling paths were typically small. IAMs might therefore opt to still use proxy Hotelling paths, since these are computationally easier to calculate, as they only require the optimisation of the initial carbon price, and not of the whole price path. On the other hand, the differences in the shape of the optimal price paths compared to the proxy Hotelling paths were large. This might also have a large societal impact, since the optimal paths better reflect the true mitigation costs: the optimal carbon price is never higher than the amount of emission reductions would require.

Our stylised model is not only useful to analyse the limits of the Hotelling rule. In fact, we used the model to calculate the sensitivity of three policy indicators on changes of the discount rate. The required amount of net negative emissions, also called the overshoot, is strongly dependent on the discount rate. Adopting a 3% discount rate would require only half the overshoot compared to a scenario calculated with the often used 5% per year.

We therefore believe that BECCS to the Future is the appropriate name for chapter 9: by choosing a lower discount rate, mitigation efforts are shifted towards the present, leading to a smaller overshoot. The need for large scale technologies like BECCS to obtain net negative emissions to bring the cumulative emissions below

the chosen carbon budget level is then shifted to the future. Moreover, an extra argument to adopt a lower discount rate in IAMs is the improved intergenerational equity.

Finally, the stylised model developed in this thesis has the potential to make complex concepts in climate policy easier to explain and to understand.

10.2. DISCUSSION

The simplicity of the model, which makes it easier to understand, could also be considered a weakness. In fact, while we have shown that for some policy indicators the results match rather well with those from a complex IAM, it is unclear if this is true for other indicators and quantities. Moreover, we only consider a cost-effectiveness setting with a fixed carbon budget. A next step would be to incorporate damages of climate change in our analysis, transforming it into a cost-benefit analysis. Other extensions would be using separate sectors with different learning rates and inertia constraints, or adapting the model to different regions.

Next, the use of a fixed end time (taken to be the year 2100 in this thesis) is also criticised by some. Using a backstop or making more realistic assumptions on the global economy might influence the results we found.

From a methodological perspective, the use of the Bellman equation raises questions on the scalability of the technique. In fact, we already found that using two state variables reaches the limit for sufficiently accurate results on an average desktop GPU. Extra dimensions in the state variable increase the required memory and calculation time exponentially. However, new techniques using machine learning to solve high dimensional Bellman equations are being developed, which could solve the so called Curse of Dimensionality.

A different approach would be to use cheaper optimisation routines. The Bellman equation can be considered a rather naive method, since for every possible value of the state variable, at every time step, the value function has to be calculated. Clever heuristics or well chosen approximations might yield similar results at a lower computational cost.

Part II

Sources of uncertainty in the costs of meeting climate targets

Chapter 11

Introduction

— *Dissecting mitigation cost uncertainties*

The second part of this thesis consists of a different application of mathematics in climate policy. Building upon the work by master student Stijn Marsman, we create a simple model that combines geophysical models with socio-economic mitigation cost estimates. Using this model, we calculate the contribution of each model component to the total uncertainty.

The Intergovernmental Panel on Climate Change (IPCC) is the international body of the United Nations in charge of providing governments with a scientific basis to base climate change policy upon. The IPCC is a collaboration of scientists from all over the world, combining expertise of the geophysical processes leading to climate change and of the socio-economic impact of adaptation and mitigation options. Its efforts to “disseminate greater knowledge about man-made climate change” were awarded with the Nobel Peace Prize in 2007. Over the last decades, the IPCC has bundled all newly available scientific research on climate change in comprehensive Assessment Report (AR). The First Assessment Report was published in 1990, the Fifth, and most recent, AR in 2014.

To produce the IPCC report, literature published regarding climate change is reviewed and joined by a large number of experts from around the world and from many different backgrounds. For instance, a group of experts on the geophysical processes of climate change joined forces to better understand the physical effect of extra CO₂ in the atmosphere, and its relation to global temperature increases. Simultaneously, policy experts have attempted to assess the global economic costs needed to keep atmospheric CO₂ concentrations below safe levels.

The geophysical processes are synthesised in an iconic graph (Fig. 11.1), visualising the relation between global temperature change and increases in atmospheric CO₂. This relation, when approximated by a linear function, is called the Transient Climate Response to Emissions (TCRE). The TCRE is closely related to the *climate sensitivity*, which gives the long term temperature change due to radiative forcing (see explanation box on this page). Different large scale geophysical climate models, each ran with different parameter values, have been used to produce the pink plume in Fig. 11.1, and is therefore a good representation of the uncertainty in climate sensitivity.

THE CHEMICAL PROCESSES BEHIND GLOBAL WARMING

The Earth gets almost all of its energy from radiative heat emitted from the Sun. When this energy hits the Earth, part of it is directly reflected back into space (29%), part of it hits the Earth's surface. Depending on the type of land cover that this radiation hits, some of the energy is reflected from the surface back into the atmosphere, some is absorbed and emitted as black body radiation. Global warming occurs from greenhouse gases: the energy with long wave lengths emitted from the Earth is re-absorbed by greenhouse gas molecules, and emitted in all directions, so also back towards the Earth.

The amount of absorption depends on the type of greenhouse gas, and on its concentration in the atmosphere. While CO₂ and other greenhouse gases have a net positive effect on radiation absorption, other particles, like aerosols could offer negative radiative forcing, by directly reflecting more energy back into space.




Figure 11.1: Global mean temperature change compared to pre-industrial times, as function of cumulative CO₂ emissions, as synthesised in the 2014 IPCC Assessment Report (IPCC, 2014). The pink plume represents the 90% confidence range across all climate model runs used to create this figure.

The IPCC figure (Fig. 11.1) contains two plumes: the pink plume, giving the total temperature increase as function of CO₂ emissions, and the gray plume, which gives the temperature change resulting purely from the forcing from CO₂ gases. The difference between the two plumes originates from the effect of non-CO₂ gases, like methane, nitrous oxide and fluorinated gases. The extra effect from these gases remains uncertain for two reasons: first, like with CO₂ gases, the contribution to the radiative forcing is uncertain. The second reason is related to political and social choices: the amount of emitted non-CO₂ gases does not necessarily have to be directly correlated with the amount of CO₂ emissions. Technological advancements or policies targeting specifically non-CO₂ gases might result in a different contribution to global temperature change.

While the pink plume does include best estimates of the influence of non-CO₂ gases, it does not account for the aforementioned uncertainties. In this thesis, we assess the uncertainty emanating from non-CO₂ gases separately.

To capture the geophysical and social processes that give the relation between CO₂ emissions and temperature change, we create a simple linear model:

$$T(\text{CO}_2) = T_{2010} + \text{TCRE} \cdot \text{CO}_2 + \sigma_{\text{non-CO}_2} \tag{11.1}$$

The TCRE term captures the distribution of the climate sensitivity, and is described in section 13.1. The uncertainty originating from non-CO₂ gas emissions is modelled through the $\sigma_{\text{non-CO}_2}$ term, and detailed in section 13.2. The temperature increase $T(\text{CO}_2)$ as function of cumulative CO₂ emissions is expressed relative to the temperature at a chosen moment. Since the mitigation costs are expressed relative to the year 2010, we also

choose T_{2010} as temperature baseline. This value is described in section 13.3.

Besides the geophysical uncertainties, we assess the uncertainty in mitigation costs as described in current literature. In section 13.4, we quantify the mitigation costs, given a carbon budget, by the function

$$\text{costs}(\text{CO}_2, p), \quad (11.2)$$

where p is a random variable capturing the uncertainty in mitigation costs. By combining Eq. (11.1) and Eq. (11.2), the mitigation costs can be expressed in terms of a temperature goal:

$$\text{costs}(T) = \text{costs}\left(\frac{T - T_{2010} - \sigma_{\text{non-CO}_2}}{\text{TCRE}}, p\right)$$

This equation is essentially an Integrated Assessment Model, as we combine a geophysical model (Eq. 11.1) with an economic model (Eq. 11.2), albeit a simple one.

Our main motivation to build this IAM is to investigate the importance of the uncertainty of each component — the geophysical parts TCRE and T_{2010} , the non-CO₂ contribution $\sigma_{\text{non-CO}_2}$ and the socio-economic part, captured in the p cost term. We describe the mathematical background for calculating partial variances in chapter 12 and apply those techniques in chapter 14. An extra result obtained from our simple model is a quantification of the carbon budgets leading to a temperature goal with a chosen probability. These carbon budgets are also discussed in chapter 14.

Chapter 12

Mathematical prelude

In many applications in industry, statistics and modelling, it is important to determine the sensitivity of a model output with respect to its respective input variables. This is useful to determine which components need refining in a manufacturing process, to focus data retrieval efforts on certain parts of a business process, or to test the robustness of an econometric model. In this part of the thesis, we apply sensitivity analysis to analyse the sources of uncertainty resulting from the input parameters of the model, which we define in chapter 13. First, we present the necessary mathematical tools.

12.1. VARIANCE DECOMPOSITION

While there are many methods to determine the contribution to uncertainty, we use a probabilistic method in which we decompose the total variance of the model output into terms accounting for the individual variance of every input variable, as well as their interactions.

To formalise this setting, we assume that the output Y of a model depends on n independent random variables $\mathbf{X} = (X_1, \dots, X_n)^T$, where each is uniformly distributed between 0 and 1:

$$\begin{aligned} Y &= f(\mathbf{X}) \\ &= f(X_1, \dots, X_n), \end{aligned}$$

where $f : K^n \rightarrow \mathbb{R}$ is assumed to yield scalar output (although the definition can be generalised to vector output), and K^n is the unit hypercube of dimension n :

$$K^n = \{\mathbf{x} \mid 0 \leq x_i \leq 1, i = 1, \dots, n\}$$

Note that we do not incur any loss of generality when assuming a uniform distribution for the input random variables, since any desired distribution can be obtained by using the inverse cumulative distribution transformation, which we discuss in section 12.3.2.

12.1.1. THE SOBOL METHOD

A commonly used method of determining the contribution to uncertainty is by using the variance-based sensitivity analysis developed by the Russian mathematician [Sobol \(1990\)](#). He proved that for any function that is square integrable over K^n , there exists a *decomposition in summands of different dimensions*:

$$f(\mathbf{x}) = f_0 + \sum_i f_i(x_i) + \sum_i \sum_{j>i} f_{ij}(x_i, x_j) + \dots + f_{12\dots n}(x_1, \dots, x_n). \quad (12.1)$$

Note that since there is always a finite number of terms in this sum, it is not an expansion in series, but rather a decomposition in terms. Moreover, Sobol required all terms except f_0 to have a zero mean among all their

dimensions: for any subset of parameter indices i_1, \dots, i_s , the function f_{i_1, \dots, i_s} must satisfy:

$$\int_0^1 f_{i_1, \dots, i_s}(x_{i_1}, \dots, x_{i_s}) dx_k = 0, \quad k = 1, \dots, i_s.$$

These functions can be found (see, again, Sobol (1990)) such that all the terms are orthogonal. In mathematical terms, given two non-identical subset of parameter indices i_1, \dots, i_s and j_1, \dots, j_t , the following statement has to hold:

$$\int_{K^{s+t}} \left(f_{i_1, \dots, i_s}(x_{i_1}, \dots, x_{i_s}) \right) \left(f_{j_1, \dots, j_t}(x_{j_1}, \dots, x_{j_t}) \right) dx_{i_1} \dots dx_{i_s} dx_{j_1} \dots dx_{j_t} = 0$$

We now give a definition of the decomposition terms that satisfy the conditions above. It is useful to remember that we are working in a *probabilistic* setting: for this reason, we will decompose the model function $f(\mathbf{X})$ using conditional expectations of the output random variable Y :

$$\begin{aligned} f_0 &= \mathbb{E}[Y], \\ f_i(X_i) &= \mathbb{E}[Y|X_i] - f_0, \\ f_{ij}(X_i, X_j) &= \mathbb{E}[Y|X_i, X_j] - f_i(X_i) - f_j(X_j) - f_0, \end{aligned} \tag{12.2}$$

and so on. With this definition, the function f_i gives the effect of varying *only* the parameter X_i . Similarly, f_{ij} gives only the effect of the direct interactions between parameters X_i and X_j . As shown in Appendix A.2.1, these functions indeed satisfy the zero-mean condition.

12.1.2. TOWARDS AN EXPRESSION FOR THE VARIANCE

While this presents us with a decomposition of the model function f , we still have no information on the individual contributions of the input parameters to the total variance. This becomes clearer when we look at the mean square difference between f and the expectation value of Y . Since we have assumed that f was square integrable, the *variance* of Y is in fact calculated as:

$$\text{Var}(Y) = \int_{K^n} f^2(\mathbf{X}) d\mathbf{X} - f_0^2$$

The right hand side can be expanded using the decomposition of Eq. (12.1). Due to the orthogonality of the terms, all the cross terms in this expansion will be zero, leaving only the pure square terms:

$$\text{Var}(Y) = \sum_{k=1}^n \sum_{i_1 < \dots < i_k}^d \int_{K^n} f_{i_1, \dots, i_k}^2(X_{i_1}, \dots, X_{i_k}) d\mathbf{X} - f_0^2$$

The beauty of this derivation is that the above sum is actually a sum of (partial) variances. In fact, as shown in Appendix A.2.2, we can rewrite the previous equation as:

$$\Leftrightarrow \text{Var}(Y) = \sum_{i=1}^n V_i + \sum_{i < j} V_{ij} + \dots + V_{12\dots n}, \tag{12.3}$$

where

$$\begin{aligned} V_i &:= \text{Var}_{X_i}(\mathbb{E}[Y|X_i]), \\ V_{ij} &:= \text{Var}_{X_{ij}}(\mathbb{E}[Y|X_i, X_j]) - V_i - V_j, \\ &\vdots \end{aligned}$$

where we use the notation $\text{Var}_{X_{ij}}$ denoting the variance obtained when integrating over the input parameters X_i and X_j . This implies that the total variance $\text{Var}(Y)$ can be decomposed into a sum of partial variances

accounting for the variance purely induced by a single parameter or by the interaction between a set of parameters.

The original intent of this method was to determine how much of the output variance can be attributed to each input parameter. This leads to the definition of the *relative* partial variance:

$$\begin{aligned} S_i &= \frac{V_i}{\text{Var}(Y)}, \\ S_{ij} &= \frac{V_{ij}}{\text{Var}(Y)}, \end{aligned} \quad (12.4)$$

and so on. Since all variance terms are non-negative, these quantities are always between 0 and 1.

Most of the time, it is only necessary to look at the *first order* variance terms, giving the relative variance purely from one input parameter, without taking into account interactions, or the *total variance*, which is the relative variance obtained when taking into account all higher order interactions. The former is equal to S_i , while the latter would require $2^n - 1$ calculations. For problems with a high number of parameters, this becomes infeasible. Fortunately, since the total relative variances sum up to 1, it is enough to calculate:

$$S_{T,i} = 1 - \frac{V_{\mathbf{x}_{\sim i}}}{\text{Var}(Y)}, \quad (12.5)$$

where $V_{\mathbf{x}_{\sim i}}$ denotes the interaction variance between all parameters except X_i .

12.2. MONTE-CARLO SIMULATION METHOD

The expression for the relative partial variance, Eq. (12.4), requires the calculation of the variance of an expectation value. For instance, for the first-order variance, such calculation would require the conditional expectation $\mathbb{E}[Y|X_i]$, and the corresponding variance over the variable X_i :

$$S_i = \frac{\text{Var}_{X_i}(\mathbb{E}[Y|X_i])}{\text{Var}(Y)}$$

A naive implementation of this calculation would be to approximate the expectation value through sampling for a particular value of X_i , then repeating this process for many X_i in order to estimate the variance. If to achieve a certain accuracy, one such estimation would require 10 000 samples, then to estimate the variance, we would have to obtain 10^8 samples, which quickly becomes infeasible.

However, a much faster Monte-Carlo method has been proposed by [Saltelli \(2002\)](#). This method relies on the fact that the above expression for S_i can be rewritten as:

$$S_i = \frac{U_i - (\mathbb{E}[Y])^2}{\text{Var}(Y)},$$

where the term U_i is equal to:

$$U_i = \int_0^1 \mathbb{E}^2[Y|X_i = x_i] dx_i,$$

again using the property that X_i is uniformly distributed. To estimate this quantity efficiently, we first need to create two matrices A and B , both containing N independent samples from every distribution of X_1, \dots, X_n :

$$A = \begin{pmatrix} x_{1,1} & x_{2,1} & \cdots & x_{n,1} \\ \vdots & \vdots & & \vdots \\ x_{1,N} & x_{2,N} & \cdots & x_{n,N} \end{pmatrix} \quad B = \begin{pmatrix} x'_{1,1} & x'_{2,1} & \cdots & x'_{n,1} \\ \vdots & \vdots & & \vdots \\ x'_{1,N} & x'_{2,N} & \cdots & x'_{n,N} \end{pmatrix}$$

While both matrices are drawn from the same distributions, the second matrix can be seen as a “re-sampling” of the distribution. Since calculating a conditional expectation is effectively calculating the expectation while keeping the conditioned variable fixed, we create a new matrix C_i for every parameter, which is equal to B , except for column i which is taken from A :

$$C_i = \begin{pmatrix} x'_{1,1} & x'_{2,1} & \cdots & x_{i,1} & \cdots & x'_{n,1} \\ \vdots & \vdots & & \vdots & & \vdots \\ x'_{1,N} & x'_{2,N} & \cdots & x_{i,N} & \cdots & x'_{n,N} \end{pmatrix}$$

This way, by comparing the model output using row j from matrix A with the model output using the corresponding row j from matrix C_i , we are varying all parameters, except parameter X_i , which is the same in A and C_i . As a consequence, U_i can be approximated as (see Appendix A.2.2 for details):

$$U_i = \int_0^1 \mathbb{E}[Y|X_i]^2 dX_i \approx \frac{1}{N} \sum_{r=1}^N f(x_{1,r}, \dots, x_{n,r}) \cdot f(x'_{1,r}, \dots, x_{i,r}, \dots, x'_{n,r}).$$

For simplicity, we can write this in matrix notation (Saltelli et al., 2008):

$$U_i \approx \frac{1}{N} f(A) \cdot f(C_i),$$

where $f(A)$ is the vector of model outputs obtained when using every row of A as input parameter values and where $f(A) \cdot f(C_i)$ is the standard dot product.

The quantity $(\mathbb{E}[Y])^2$ can be approximated by:

$$(\mathbb{E}[Y])^2 \approx \left(\frac{1}{N} \sum_{j=1}^N f(A)_j \right)^2 =: (\hat{f}_0)^2,$$

or, if fast convergence is required, Homma and Saltelli (1996) have shown that the expectation value can be approximated more accurately by:

$$(\mathbb{E}[Y])^2 \approx \frac{1}{N} \sum_{j=1}^N f(A)_j f(B)_j = \frac{1}{N} f(A) \cdot f(B),$$

i.e. by using both sample points A and B . Finally, the total variance $\text{Var}(Y)$ is approximated by:

$$\text{Var}(Y) \approx \frac{1}{N} f(A) \cdot f(A) - (\hat{f}_0)^2$$

Combining everything, the first order partial variance can be calculated by:

$$S_i = \frac{\text{Var}_{X_i}(\mathbb{E}[Y|X_i])}{\text{Var}(Y)} \approx \frac{\frac{1}{N} f(A) \cdot f(C_i) - (\hat{f}_0)^2}{\frac{1}{N} f(A) \cdot f(A) - (\hat{f}_0)^2}.$$

Similarly, it can be shown that the total relative variance $S_{T,i}$ is calculated by:

$$S_{T,i} = 1 - \frac{V_{\mathbf{X}_{\sim i}}}{\text{Var}(Y)} \approx 1 - \frac{\frac{1}{N} f(B) \cdot f(C_i) - (\hat{f}_0)^2}{\frac{1}{N} f(A) \cdot f(A) - (\hat{f}_0)^2}.$$

12.3. STATISTICAL CONSIDERATIONS

The method to calculate partial variances using a Monte Carlo procedure, as described in the previous section, requires sampling values from the distributions of the random variables X_i . While for some distributions, numerical implementations that sample accordingly are readily available, this is not the case for arbitrary, continuous distributions. One such distribution that we will use extensively is the β -PERT distribution.




Figure 12.1: Probability Density Function (left) and inverse Cumulative Density Function (right) of a β -PERT distributed random variable with minimum $a = 1$, maximum $c = 2$ and three most likely values b .

12.3.1. THE β -PERT DISTRIBUTION

An often used distribution for assessing risk and uncertainties in the field of operational research and risk management, is the β -PERT distribution (Benny, 1989). Using the PERT (Program Evaluation and Review Technique) method, estimates of the minimal required time to finish a project can be made using three parameters: the optimistic time estimation, the most likely time and the pessimistic estimation.

In this spirit, the β -PERT distribution was created: based on the Beta distribution, it requires three parameters: the minimum a , most likely b and maximum value c that a random variable can take. Its Probability Density Function (PDF) is defined, for a random variable $X \sim \beta\text{-PERT}(a, b, c)$, as:

$$f_X(x) = \frac{(x - a)^{\alpha-1}(c - x)^{\beta-1}}{B(\alpha, \beta)(c - a)^{\alpha+\beta-1}},$$

where $\alpha = \frac{4b+c-5a}{c-a}$ and $\beta = \frac{5c-a-4b}{c-a}$, and $B(\alpha, \beta)$ is the Beta function:

$$B(p, q) = \int_0^1 t^{p-1}(1 - t)^{q-1} dt.$$

The PDF is shown in Fig. 12.1 for several values of b . Clearly, for values of b farther away from the midpoint of a and c , the distribution becomes more skewed.

Finally, the Cumulative Distribution Function (CDF) of the β -PERT distribution can be calculated by integrating the PDF:

$$F_X(x) = B\left(\frac{x - a}{c - a}; \alpha, \beta\right),$$

where $B(z; p, q)$ is the incomplete Beta function:

$$B(z; p, q) = \int_0^z t^{p-1}(1 - t)^{q-1} dt.$$

12.3.2. SAMPLING FROM AN ARBITRARY DISTRIBUTION

As mentioned above, sampling from the β -PERT distribution, like many other distributions, is often not directly implemented in numerical software packages. However, by using the Inverse CDF sampling method, a

(pseudo)-random number generator producing uniformly distributed numbers between 0 and 1 is enough to obtain statistically the same result.

This becomes clearer when looking at the inverse CDF of the β -PERT distribution in Fig. 12.1: for the very skewed distribution with $b = 1.95$, a sample value smaller than 0.1 would have to be obtained from the uniform sampler between 0 and 1 to obtain a β -PERT distributed value smaller than 1.6. In other words, there is a much larger probability (90%) that the resulting value is between 1.6 and 2, and only a 10% probability that it is between 1 and 1.6. In general, when the inverse CDF can be calculated, any distribution can be sampled using a random number generator with uniform output between 0 and 1.

12.3.3. PERCENTILES FROM DISTRIBUTIONS

The last mathematical consideration we want to discuss before introducing the model, is the calculation of *percentiles*. In statistics, when analysing a series of N observations x_i , we define the $p\%$ percentile to be the value below which $p\%$ of the observations are located. To calculate this, we first sort the values x_i from low to high, then take the value at position $\lceil p\% \cdot N \rceil$, where $\lceil x \rceil$ is the nearest integer larger than x .

Chapter 13

The model

As mentioned earlier, the goal of this part of the thesis is to combine the geophysical uncertainties in the relationship between CO₂ emissions and global temperature, with the economic and societal uncertainties of the global costs associated with keeping CO₂ concentrations below a given level. In this chapter, we first create a simple model for the geophysical quantities, then include variance from non-CO₂ emissions and finally create a probabilistic model of abatement costs.

13.1. CLIMATE SENSITIVITY

The first part we focus on, is the *climate sensitivity*: a measure of how sensitive our climate, and in particular global mean temperature, is to increases of CO₂ in the atmosphere. This sensitivity has been calculated using a variety of large, global climate models (IPCC, 2014). Different models with different parameters have produced a range of output values for the global mean temperature as function of the amount of CO₂ in the atmosphere, shown as the pink plume in Fig. 11.1. The most likely values are highlighted as the solid lines.

Temperature change does not actually directly result from higher CO₂ concentrations: the CO₂ contributes to an increase in radiative forcing, which in turn induces temperature increases.

An increased forcing can also be caused by non CO₂ gases: the total temperature increase is therefore larger when considering all greenhouse gases then when looking at only the CO₂. This difference is shown in Fig. 11.1 as the difference between the gray plume and the pink plume, and is discussed in more detail in the next section.

To keep our model as simple as possible, we disregard the intermediate forcing step, and approximate the temperature change as function of cumulative CO₂ to a linear relation:

$$T(\text{CO}_2) = T_{1870} + \text{TCRE} \cdot \text{CO}_2, \quad (13.1)$$

where T_{1870} is the initial temperature (average temperature around 1870), CO_2 denotes the cumulative CO₂ difference starting from the moment T_{1870} was measured and TCRE is the Transient Climate Response to cumulative carbon Emissions. The TCRE is often used in climate science to describe the amount of temperature increase given a unit increase of CO₂ in the atmosphere.

13.1.1. PARAMETER DISTRIBUTION

The two quantities T_{1870} and TCRE are estimated from Fig. 11.1. Because of their statistical nature, it is not sufficient to simply find fixed values for them, we have to find their probability distribution.

The initial temperature relative to 1861-1880 has, by definition, a zero mean, since it gives the temperature increase given zero extra CO₂ in the atmosphere compared to 1870. However, because of measurement uncertainties and the difficulty of estimating the global temperature over a century ago: we have to associate a distribution to it. We will only assume that T_{1870} has a normal distribution, with zero mean and standard deviation σ , which we will calculate in the next subsection:

$$T_{1870} \sim \mathcal{N}(0, \sigma).$$

Next, we examine the TCRE. The slope of the solid line in Fig. 11.1 equals 0.62, and also has an associated distribution. However, we cannot use the normal distribution again here, for two reasons:

- the distribution has to be asymmetric, as seen in the IPCC figure;
- it is generally assumed that the TCRE is strictly positive, and the normal distribution can yield negative values.

Because of the boundedness and skewness requirements, the β -PERT distribution is chosen to model the TCRE distribution. As discussed in section 12.3.1, this distribution requires a minimum, a mode and a maximum value:

$$\text{TCRE} \sim \beta\text{-PERT}(\text{min}, 0.62, \text{max})$$

While the mode is known and equal to 0.62 (the average slope of the solid line in Fig. 11.1), we will estimate the minimum and maximum value using the range of the pink plume.

13.1.2. FITTING THE DISTRIBUTIONS TO THE PLUME

Since the pink plume in Fig. 11.1 represents the 90% confidence range, we can assume that the bottom of the plume represents the 5% percentile of the climate model runs, while the top corresponds to the 95% percentile. We are now posed with an optimisation problem: find values for σ , min and max such that the resulting joined distribution best matches the bottom and top percentiles.

While one might first estimate the value for σ by examining the initial range of the plume around $\text{CO}_2 = 0$, and the range of the slope of the TCRE by looking at the slope of the bottom and top part of the pink plume, this method would yield a different joined distribution than required, since there exists a correlation between the two parameters (the distributions are independent, but the resulting temperature does depend on both parameters). We now present a more rigorous way of fitting the distributions.

First of all, since the temperature is linearly dependent on the CO_2 , the probability density function of the transformed distribution of the temperature can be calculated analytically through a convolution integral. The PDF of T , given a CO_2 level, is given by:

$$f_T(t; \text{CO}_2) = \int_{-\infty}^{\infty} f_{T_{1870}}(t') \cdot f_{\text{TCRE} * \text{CO}_2}(t - t') dt',$$

where $f_{T_{1870}}$ and $f_{\text{TCRE} * \text{CO}_2}$ are the respective PDFs of T_{1870} and $\text{TCRE} * \text{CO}_2$. The first is the standard normal distribution PDF, while the second is the transformed β -PERT distribution, given by:

$$\text{TCRE} * \text{CO}_2 \sim \beta\text{-PERT}(\text{min} \cdot \text{CO}_2, 0.62 \text{ CO}_2, \text{max} \cdot \text{CO}_2)$$

The above integral has to be evaluated numerically, since no analytical solution exists for this. However, since we have to find values for σ , min and max, we are actually dealing with an inverse problem. In fact, we have to compare the 5% and 95% percentiles of this transformed distribution to the upper and lower values of the pink plume. These percentiles can be calculated using the inverse CDF:

$$F_T^{-1}(0.05; \text{CO}_2) \quad \text{and} \quad F_T^{-1}(0.95; \text{CO}_2),$$

where the cumulative distribution function $F_T(\cdot)$ is calculated as:

$$F_T(t; \text{CO}_2) = \int_{-\infty}^t f_T(s; \text{CO}_2) ds.$$

Calculating this integral, however, poses some numerical problems, given the complicated nature of the PDF $f_T(\cdot)$. One property helps us: the fact that we are dealing with linear functions. To compare the distribution to a line, we only need to compare it at two points of CO_2 . We use $\text{CO}_2 = 1000 \text{ GtCO}_2$ and $\text{CO}_2 = 4500 \text{ GtCO}_2$. The temperature change at these points are, as read from Fig. 11.1, equal to:

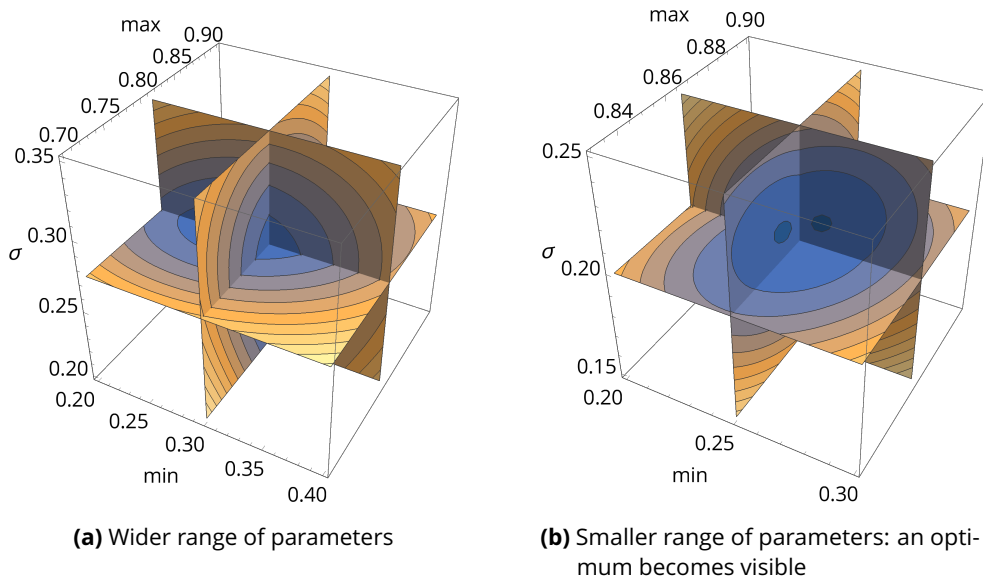


Figure 13.1: Minimisation of the difference between the calculated inverse CDF and the values of the pink plume, for two ranges of the distribution parameter values σ , min and max. The bluer, the smaller the difference with the pink plume.

CO ₂	Plume	
	bottom (5% percentile)	top (95% percentile)
1 TtCO ₂	0 °C	0.8 °C
4.5 TtCO ₂	1.8 °C	3.6 °C

The goal now is to find the distribution parameter values such that $F_T^{-1}(0.05; \text{CO}_2)$ and $F_T^{-1}(0.95; \text{CO}_2)$ are as close as possible to the values in this table, for the two values of CO₂ chosen.

In order to calculate the inverse CDF values, we first approximate the CDF $F_T(t; \text{CO}_2)$. The integral cannot be calculated easily using traditional quadrature schemes, since such schemes would require too much calculation power. Instead, we calculate the PDF $f_T(t; \text{CO}_2)$ at 20 points, and interpolate it using cubic interpolation. This method gives a good approximation because of the smoothness of the PDF. Using this interpolated function, we calculate the inverse CDF using a standard Newton-Raphson scheme, giving us $F_T^{-1}(0.05; \text{CO}_2)$ and $F_T^{-1}(0.95; \text{CO}_2)$ for a chosen set of parameters. These values are then compared to the values in the table above.

Most optimisation schemes require some form of gradient in the parameters σ , min and max, which is unfeasible here. For this reason, we use a rather naive brute-force method. Let $\Delta(\sigma, \text{min}, \text{max})$ be the difference between calculated model inverse CDF and plume values, defined as:

$$\Delta(\sigma, \text{min}, \text{max}) = \left(F_T^{-1}(0.05; \text{CO}_2 = 1) - 0 \text{ °C} \right)^2 + \left(F_T^{-1}(0.95; \text{CO}_2 = 1) - 0.8 \text{ °C} \right)^2 \\ + \left(F_T^{-1}(0.05; \text{CO}_2 = 4.5) - 1.8 \text{ °C} \right)^2 + \left(F_T^{-1}(0.95; \text{CO}_2 = 4.5) - 3.6 \text{ °C} \right)^2,$$

where the quantities $F_T^{-1}(\cdot; \cdot)$ depend on the parameters σ , min and max. This difference is then evaluated for many combinations of the parameter values. The best values for the parameters are then those that minimise Δ :

$$\sigma^*, \text{min}^*, \text{max}^* = \arg \min_{\sigma, \text{min}, \text{max}} \Delta(\sigma, \text{min}, \text{max})$$

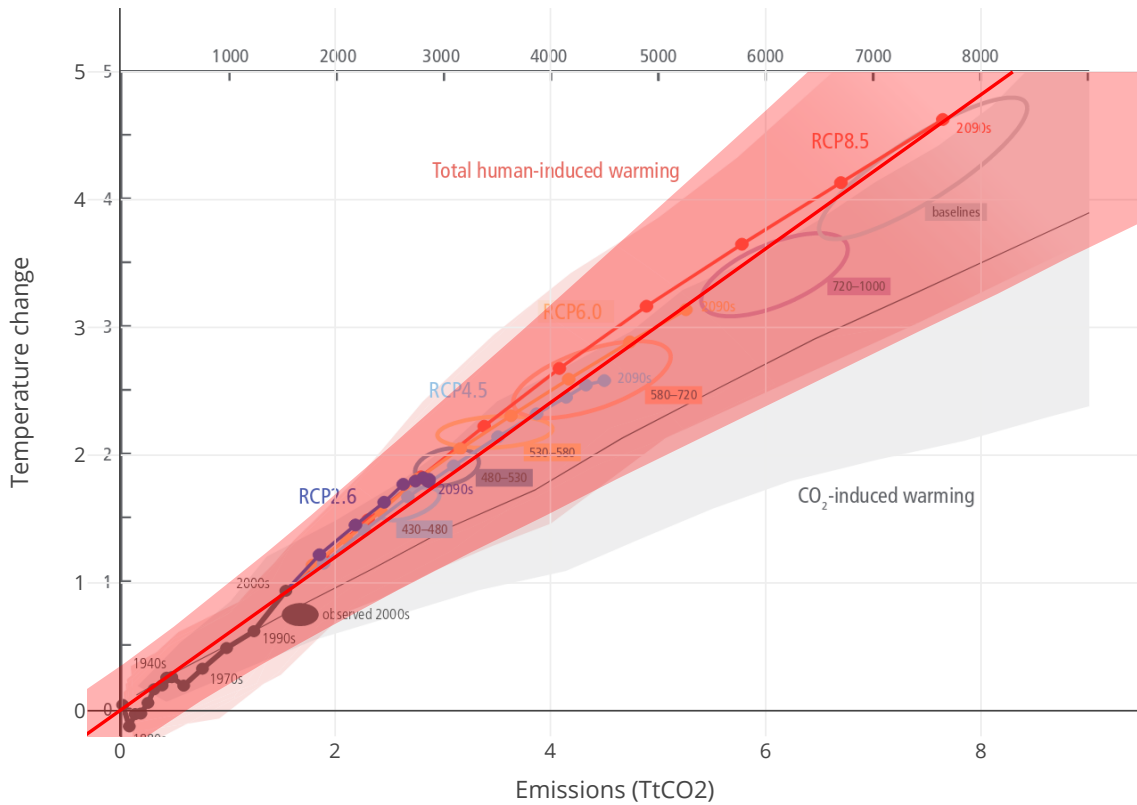


Figure 13.2: Results from sampling our linear model compared to the IPCC figure. The bottom and top lines of the plume correspond to the empirical 5% and 95% percentiles of the sampling results.

By using this method, we make no assumption on the optimisation landscape. We start with a rather large range of parameter values (Fig. 13.1a), and rerun the algorithm focussing on the parameter range giving the smallest difference (Fig. 13.1b).

We find that when using the parameters $\sigma^* = 0.21$, $\min^* = 0.255$ and $\max^* = 0.855$, our model best replicates the pink plume. To summarise, we have found the distributions of T_{1870} and TCRE to be:

$$T_{1870} \sim \mathcal{N}(0, 0.21)$$

$$\text{TCRE} \sim \beta\text{-PERT}(0.255, 0.62, 0.855)$$

We realise that our brute force can be replaced by much more elegant methods, like Markov Chain Monte-Carlo, or even Neural Networks, but since this very simple method already yields good enough results, we have chosen to keep it.

We can test the validity of these distributions by sampling T_{1870} and TCRE for a number of values of the CO_2 , between 0 and 9000 GtCO_2 , and calculating the empirical mode, 5% percentile and 95% percentile, and overlaying these values on top of Fig. 11.1. This empirical distribution is shown in Fig. 13.2. Clearly, our model yields samples in very good accordance with the pink plume.

13.2. NON CO₂ CONTRIBUTIONS

While CO₂ emissions are the largest contributor to global temperature increases, other gases also play a significant role. Methane (CH₄), for example, primarily resulting from agriculture and biomass burning, contributes to 16% of all greenhouse gas emissions, but also nitrous oxide (N₂O) from fertilisers and fluorinated gases (HFCs, PFCs...) contribute to the greenhouse effect (IPCC, 2014).

Like with any measurement, there is a physical uncertainty to the extent of the contribution of these gases to the total forcing, which in turn induces global temperature change. However, the largest uncertainty comes from societal uncertainties. In fact, the amount of non-CO₂ gases strongly depends on which economic activities are developed in the future, which activities are reduced and how current and future climate policy addresses the abatement of non-CO₂ gases.

The figure from the IPCC report (Fig. 11.1) relating temperature increase to increased atmospheric CO₂ levels does include the contribution of non-CO₂ gases to temperature change, giving the difference between the pink and the gray plume. However, this graph does not account for variability in the amount of non-CO₂ gases given an amount of CO₂ emissions (see explanation box on RCPs below).

To quantify the total non-CO₂ related uncertainty, we analyse a wide variety of scenarios from the Shared Socio-economic Pathways (SSP) database (van Vuuren et al., 2014; O'Neill et al., 2014; Kriegler et al., 2014), created using six different Integrated Assessment Modelling teams from all over the world. The database consists of a wide variety of model runs, exploring various socio-economic and climatic scenarios. For each of these runs, a large number of output variables is produced, among which:


- total CO₂ emitted during the period 2010-2100,
- forcing due to CO₂ emissions,
- forcing due to non-CO₂ emissions,
- temperature change compared to pre-industrial times.

We use these variables to estimate the uncertainty related to the non-CO₂ emissions. First of all, since our model does not consider forcing, but rather temperature increases, we look at the relationship in these SSP outputs between temperature and forcing, where forcing is the sum of CO₂ and non-CO₂ forcing. As shown in Fig. 13.3, the temperature increase follows an almost linear relation with the total forcing, with best fit equation:

$$T = 0.60713F. \quad (13.2)$$

SSP'S: CLASSIFYING OUR FUTURE SOCIETY

In order to calibrate scenarios across different IAM teams, five distinct Shared Socioeconomic Pathways have been created: these look purely at the geo-political and social structure of the world. The SSPs are divided along two axes: difficulty of overcoming adaptation challenges and difficulty of overcoming mitigation challenges, giving a measure of economic development, rivalry between countries and population growth:



RCP'S AND THE CO₂/NON-CO₂ CORRELATION

In order to create the pink plume of Fig. 11.1, different modelling teams set out to devise scenarios which reach different levels of climate change, by imposing a total forcing level for the end of the scenario: 2.6, 4.5, 6.0 and 8.5 W/m² by 2100. These scenarios are called the RCP's, and shown as solid lines in Fig. 11.1. These lines overlap rather well. Assuming that there is a clear relation between CO₂ and temperature, the only way these lines would overlap, is if the contribution of the non-CO₂ gases would be the same in the scenarios, relatively. In other words, these lines can only overlap if there is a very strong correlation between CO₂ emissions and non-CO₂ emissions. Since the lines do almost overlap, we can assume that the different RCP's have been created using very similar assumptions on non-CO₂ emissions. For this reason, we can state that the whole socio-economic variability of non-CO₂ emissions has not been accounted for in the pink plume.

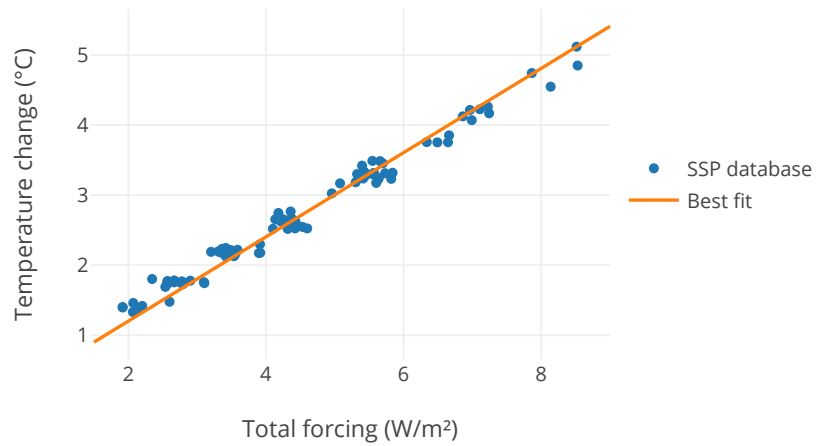


Figure 13.3: Relationship between temperature change and total forcing, calculated from the SSP database, combined with the best fit relation $T = 0.60713F$. The total forcing is the sum of CO₂ and non-CO₂ induced forcing.

The linear relation is a consequence from the fact that all IAMs used in this database use the very simplified MAGICC climate model (Model for the Assessment of Greenhouse Gas Induced Climate Change). From now on, we assume that we can transform any forcing value to a temperature change using Eq. (13.2).

Next, we can examine the individual CO₂ and non-CO₂ contribution to the forcing, for every model run in the SSP database. We directly convert these forcings to temperatures: $T_{\text{CO}_2} = 0.60713F_{\text{CO}_2}$ and $T_{\text{non-CO}_2} = 0.60713F_{\text{non-CO}_2}$. In Fig. 13.4, we show T_{CO_2} and $T_{\text{non-CO}_2}$ as function of the cumulative CO₂ of the database entries. From this figure, we first observe that the temperature increase purely from CO₂ emissions follows a clear, almost linear, relation, without significant variance. This also comes from the fact that these models all use the simple MAGICC climate model. One of the simplifications is that the effect of CO₂ on the global temperature is fixed and calculated without variance (Meinshausen et al., 2011), while the pink plume of the IPCC report uses a wide variety of very detailed, large-scale climate models, which take much longer to run.

Next, we observe in Fig. 13.4 a large spread in the values of $T_{\text{non-CO}_2}$. We assume a linear relation, with as best fit the orange line in Fig. 13.4. The reason that there are less data points for higher CO₂ values, is simply that there have been much less scenario runs for these parameters. We also include the 5% and 95% lines, where we shift the best fit line such that 5% and 95% of the data points are below it. These lines are respectively 0.19°C below and 0.25°C above the best fit line.

While there is an asymmetry in the uncertainty, we consider this to be small enough to neglect, and take the average of the upper and lower deviation: 0.222°C.

This approximation helps us to model the contribution of non-CO₂ emissions to the temperature increase as a linear function (the best fit), plus a variance term $\sigma_{\text{non CO}_2}$, with zero mean and standard deviation σ :

$$T_{\text{non-CO}_2}(\text{CO}_2) = T_{\text{best fit}}(\text{CO}_2) + \sigma_{\text{non CO}_2}, \quad \text{where } \sigma_{\text{non CO}_2} \sim \mathcal{N}(0, \sigma)$$

To calculate σ , we use the property of the normal distribution that the inverse CDF of a normally distributed random variable X evaluated at 5% is given by:

$$F_X^{-1}(0.05) = \mu - 1.64485\sigma,$$

where F_X represents the cumulative distribution function. In our case, the previous quantity should equal -0.222°C. By solving for σ and using the zero mean assumption, we obtain a standard deviation of $\sigma = -0.222/(-1.64485) = 0.135$.

Finally, we assume that since the mean contribution of non-CO₂ emissions was already accounted for in the pink plume of Fig. 11.1, but not the variance, the temperature increase given an increase of CO₂ is calculated by:

$$T(\text{CO}_2) = T_0 + \text{TCRE} \cdot \text{CO}_2 + \sigma_{\text{non CO}_2}, \quad (13.3)$$




Figure 13.4: Temperature change contribution from CO₂ (left) and non-CO₂ (right) gases, as function of total amount of CO₂ from 2010 to 2100 for every model run in the SSP database. Note that the y-axes have a different range. The non-CO₂ contribution graph also includes the best linear fit, along with the 90% confidence range.

where $\sigma_{\text{non-CO}_2} \sim \mathcal{N}(0, 0.135)$, the distribution of TCRE was determined in the previous section, and T_0 corresponds to the initial temperature that is being considered. In the IPCC figure, T_0 was the temperature in 1870, since the CO₂ emissions corresponded to cumulative emissions *from the year 1870*. However, we are more interested in the effect of extra future emissions, and will therefore only consider emissions *from the year 2010*. The initial temperature associated with this, T_{2010} , is discussed in the next section.

13.3. CURRENT GLOBAL TEMPERATURE

The main goal of this part of the thesis is to determine the total costs required to reach a certain temperature goal by the end of the century. These temperature goals are measured relative to pre-industrial times. To calculate how much more CO₂ can be emitted for a given goal, it is important to know the current global temperature, also relative to pre-industrial times.

This quantity is determined to be $T_{2010} = 0.909 \pm 0.15^\circ\text{C}$ (Visser et al., 2018), where the uncertainty range corresponds to twice the standard deviation. We assume that the underlying statistical distribution of T_{2010} is normal, with a mean of 0.909°C . The standard deviation σ is therefore equal to $\sigma = 0.075$, giving the total distribution:

$$T_{2010} \sim \mathcal{N}(0.909, 0.075)$$

13.4. COSTS

The last part of our model consists of the correspondence between cumulative CO₂ emissions and costs. More specifically, to keep the cumulative emissions from 2010 to 2100 below a certain carbon budget, large scale investments and mitigation policies will have to be put in place. In the same SSP database as we used to calculate the non-CO₂ contribution, several economic and societal scenarios have been investigated to reach these carbon budgets. While in the first part of this thesis, we defined the total mitigation costs as being the area under the MAC (the integral of the Marginal Abatement Cost curve, summed and discounted over the years), the underlying models report costs using three different measures:

- Area under MAC: the same measure as we have used throughout this thesis.
- Consumption loss: from an economists view point, imposing a carbon tax means that extra money has to be invested in clean technologies, money that otherwise could have been used for general consumption.

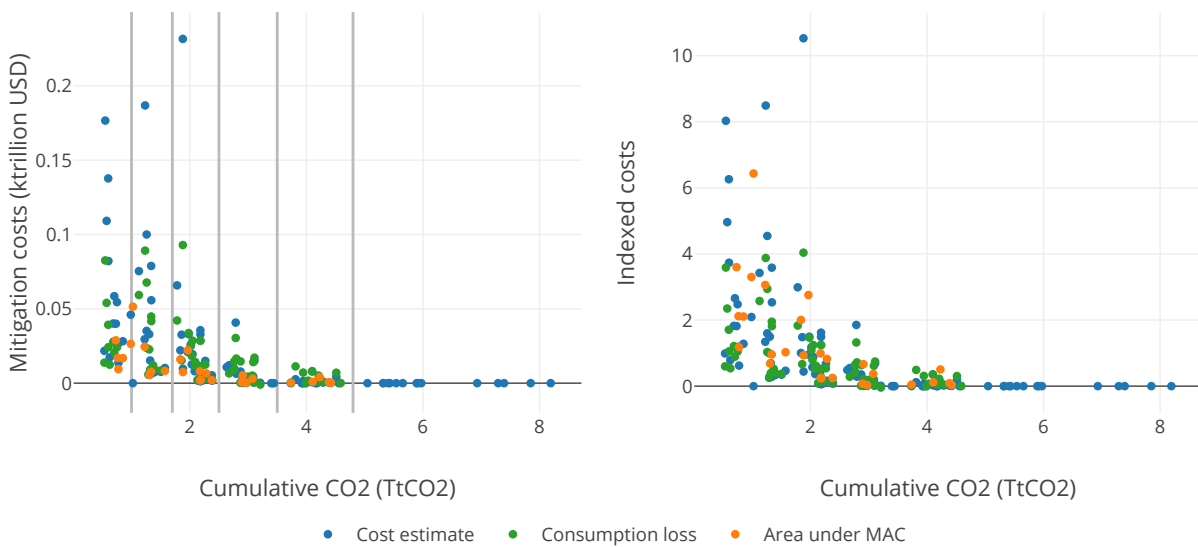


Figure 13.5: Mitigation costs of the tree cost quantities produced in the SSP database (left). The indexed costs (right) are normalised such that the medians of the three measures are identical.

Therefore, a way of measuring total costs is by calculating the loss of consumption.

- **Cost estimate:** when total costs have not been reported directly, the costs can be estimated by multiplying the total abatement by the carbon tax, and sum this over the years. This method is similar to the Area under MAC, except for the fact that we do not know the shape of the MAC, which makes it difficult to integrate it, hence this estimate.

These quantities, as function of the cumulative CO₂ from the corresponding model run, are shown in Fig. 13.5 (left). Not all model runs provide all three cost measures, since some of them are not calculated by the respective IAM.

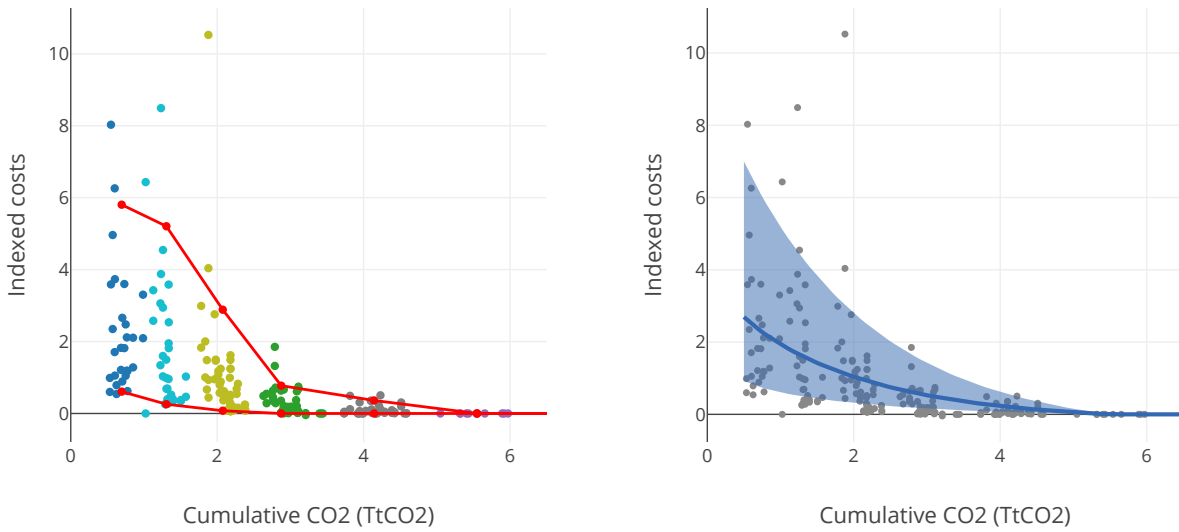
13.4.1. COMPARING THE THREE COST MEASURES

While all three cost measures are expressed in US dollars, they represent different economic quantities. In fact, in Fig. 13.5, the Area under the MAC is consistently lower than the two other measures. While one could debate on the economic validity and origin of this, we are merely interested in the change in costs for different carbon budgets.

To properly compare the three measures, we need to normalise them. We use the method proposed in the master thesis of Marsman (2018): we attempt to equalise the median values of the three units. To do this, we split the cost results in Fig. 13.5 in six groups of CO₂ values, since the results are clustered around these groups, as shown in the same figure. For each group, we calculate the median value of each of the three cost measures. These will be our index values, with which we normalise the three datasets. For better consistency, we choose one single group for which we calculate the median values, namely the second group, with CO₂ values between 1 and 1.7 Tt CO₂. Then, the indexed costs are calculated by dividing all costs by their respective index. The resulting indexed costs are shown in Fig. 13.5 (right).

13.4.2. TOWARDS A CONTINUOUS DESCRIPTION

In the probabilistic framework that we are using, it is necessary to use a continuous function (with associated distribution) describing the costs as function of CO₂ emissions. We approach this problem by creating three functions, describing the mode, 5% percentile and 95% percentile of the data points. These functions have the same functional form: we use an exponentially decreasing function, with the extra assumption that the function



(a) Empirical 5% and 95% percentiles per coloured group of CO₂ values.

(b) Comparison to the statistical model, where the range corresponds to the 90% confidence range.

Figure 13.6: 90% confidence range calculated empirically per group of CO₂ (left) and using the statistical model (right).

is zero for CO₂ budgets higher than 5.5 Tt CO₂. The functions have the following form:

$$\text{costs}(\text{CO}_2; a, b) = \begin{cases} a e^{-b \text{CO}_2} - a e^{-5.5b} & \text{CO}_2 < 5.5 \\ 0 & \text{CO}_2 \geq 5.5 \end{cases}, \quad (13.4)$$

where the parameters a and b have to be determined to match the mode, 5% and 95% percentiles. First, the mode is calculated using a best fit, linear regression, routine:

$$\text{costs}_{\text{best fit}}(\text{CO}_2) = \text{costs}(\text{CO}_2; 4.31, 0.724).$$

Next, the parameters a and b for the 5% and 95% percentiles are calculated such that respectively 5% and 95% of the data points reside below the calculated function. However, there are many ways to calculate this, all leading to different parameter values. In order to stay as statistically accurate as possible, we proceed by again grouping all indexed cost points using the same groups as Fig. 13.5 (right). For every group, we calculate the empirical 5% and 95% percentile values (red lines in Fig. 13.6a). Then, we fit the exponential function of Eq. (13.4) to these lines, to obtain:

$$\begin{aligned} \text{costs}_{5\%}(\text{CO}_2) &= \text{costs}(\text{CO}_2; 1.705, 1.477), \\ \text{costs}_{95\%}(\text{CO}_2) &= \text{costs}(\text{CO}_2; 9.844, 0.540), \end{aligned}$$

13.4.3. CALCULATING THE STATISTICAL DISTRIBUTION

In the previous section, we found three functions describing respectively the mode, 5% percentile and 95% percentile of the costs of reaching a given carbon budget. To describe the variance between these functions, it is not possible anymore to simply add a constant variance term to the best fit line, since the variance is much higher for low values of CO₂ than for high carbon budgets.

To still account for the variability using a single parameter, we define p as being a parameter between 0 and 1, such that the costs are a linear combination of some minimum and some maximum cost function:

$$\text{costs}(\text{CO}_2, p) = \text{costs}_{\text{min}}(\text{CO}_2) + p \left(\text{costs}_{\text{max}}(\text{CO}_2) - \text{costs}_{\text{min}}(\text{CO}_2) \right), \quad (13.5)$$

where

$$\begin{aligned} \text{costs}_{\min}(\text{CO}_2) &= \text{costs}(\text{CO}_2; a_{\min}, b_{\min}), \quad \text{and} \\ \text{costs}_{\max}(\text{CO}_2) &= \text{costs}(\text{CO}_2; a_{\max}, b_{\max}). \end{aligned}$$

We again use a β -PERT distribution for p , since it has a clearly defined minimum, 0, and maximum, 1, and an asymmetric distribution is needed to account for the higher spread for higher costs than for lower costs:

$$p \sim \beta\text{-PERT}(0, p^*, 1)$$

To represent the data faithfully, we need to find the set of parameters p^* , a_{\min} , b_{\min} , a_{\max} and b_{\max} such that when evaluating $\text{costs}(\text{CO}_2, p)$ with $p = p^*$, we obtain the best fit cost function, and similarly with the 5% and 95% percentile values of p , which should correspond to $\text{costs}_{5\%}(\text{CO}_2)$ and $\text{costs}_{95\%}(\text{CO}_2)$.

This brings us to an optimisation problem, similar to the one used to calculate the distribution of TCRE and T_{1870} . We define a penalty function which measures the distance between $\text{costs}_{\text{best fit}}(\text{CO}_2)$, $\text{costs}_{5\%}(\text{CO}_2)$ and $\text{costs}_{95\%}(\text{CO}_2)$ and the respective predictions from the statistical distribution, for a number of values of CO_2 :

$$\begin{aligned} \text{penalty}(p^*, a_{\min}, b_{\min}, a_{\max}, b_{\max}) &= \sum_{\{\text{CO}_2\}} \left(\text{costs}_{\text{best fit}}(\text{CO}_2) - \text{costs}(\text{CO}_2, p^*) \right)^2 \\ &\quad + \left(\text{costs}_{5\%}(\text{CO}_2) - \text{costs}(\text{CO}_2, p_{5\%}) \right)^2 \\ &\quad + \left(\text{costs}_{95\%}(\text{CO}_2) - \text{costs}(\text{CO}_2, p_{95\%}) \right)^2, \end{aligned}$$

where $p_{x\%}$ is the $x\%$ percentile of p using the distribution defined above, and $\{\text{CO}_2\}$ is a discretisation of the range $[0, 5.5]$ Tt CO_2 in 250 steps. It should be noted that the functions $\text{costs}(\text{CO}_2, p)$ depend on the parameters a_{\min} , b_{\min} , a_{\max} and b_{\max} .

The penalty function is then minimised using the numerical Newton-Raphson scheme. This yields the following values:

$$\begin{aligned} \text{costs}_{\min}(\text{CO}_2) &= \text{costs}(\text{CO}_2; 0.8905, 3.4815), \\ \text{costs}_{\max}(\text{CO}_2) &= \text{costs}(\text{CO}_2; 15.012, 0.5476), \end{aligned}$$

and, for the distribution of p :

$$p \sim \beta\text{-PERT}(0, 0.242, 1).$$

This concludes the statistical description of the costs. To test our model, we show the cost points from the database combined with the theoretical prediction of the statistical model in Fig. 13.6b.

13.5. PUTTING IT ALL TOGETHER

In this chapter, we have extensively discussed how we estimate the statistical distributions of the different parameters of our model. We combine this all here. First of all, the temperature and CO_2 relation is given by:

$$T(\text{CO}_2) = T_{2010} + \text{TCRE} \cdot \text{CO}_2 + \sigma_{\text{non-CO}_2}. \quad (13.6)$$

Second, we describe the costs to reach a given carbon budget CO_2 by the function:

$$\text{costs}(\text{CO}_2, p),$$

which was a linear combination of $\text{costs}_{\min}(\text{CO}_2)$ and $\text{costs}_{\max}(\text{CO}_2)$ previously defined.

However, we are interested in calculating the costs as function of temperature, and not of CO_2 . For this reason, we use Eq. (13.6) to express the carbon budget in terms of temperature change:

$$\text{CO}_2(T) = \frac{T - T_{2010} - \sigma_{\text{non-CO}_2}}{\text{TCRE}} \quad (13.7)$$

By plugging the previous equation into the expression for the costs, we obtain the desired form:

$$\text{costs}(T) = \text{costs}\left(\frac{T - T_{2010} - \sigma_{\text{non-CO}_2}}{\text{TCRE}}, p\right), \quad (13.8)$$

where we use the following distributions:

$$\begin{aligned} \text{TCRE} &\sim \beta\text{-PERT}(0.255, 0.62, 0.855) \\ T_{2010} &\sim \mathcal{N}(0.909, 0.075) \\ \sigma_{\text{non CO}_2} &\sim \mathcal{N}(0, 0.135) \\ p &\sim \beta\text{-PERT}(0, 0.242, 1). \end{aligned}$$

Chapter 14

Results — Uncertainty in mitigation costs

Using the statistical model defined in the previous chapter, we proceed to calculate two specific results. First, we estimate the model output by sampling the underlying distributions, and calculating the resulting costs as function of the temperature goal. Second, we analyse the sensitivity of the model output to the underlying parameters, to determine their contribution to the total variance.

14.1. MODEL OUTPUT

The model described in the previous chapter, summarised in Eq. (13.8), calculates the abatement costs required to reach a given temperature goal. Since it is not possible to calculate the resulting distribution of the costs analytically, given the four input parameters and their distributions, we use a sampling method. Specifically, for a range of values of CO_2 , we sample a total of N values from the distributions of TCRE , T_{2010} , $\sigma_{\text{non CO}_2}$ and p and calculate separately the resulting temperature using Eq. (13.6) and the abatement costs using Eq. (13.5). The resulting point has than as abscissa (x -coordinate) the temperature, and as ordinate (y -coordinate) the costs. This is shown in Fig. 14.1 for 6500 values of CO_2 between 0.5 and 6 Tt CO_2 .

It should be noted that we could also have sampled without the separation step using the CO_2 . In fact, by choosing N values of the temperature T , and sampling the parameter distributions for every point, we directly obtain the empirical distribution of the costs as function of temperature. While mathematically, these methods converge to the same distribution, we would disregard the information about the *data-availability*: in fact, while the temperature distribution is well known for a wide range of CO_2 values, except very high emission levels, the costs are only calculated using scenarios with a carbon budget of at least 0.5 Tt CO_2 . Therefore, the reason that in Fig. 14.1 (bottom), there is a smaller point density towards lower temperature goals, is not just because of the higher variance in costs, but also because of the fact that we don't use samples with CO_2 values lower than 0.5 Tt CO_2 .

14.1.1. CARBON BUDGETS

By using the first part of the model, namely temperature as function of CO_2 , we can calculate the distribution of an important quantity: carbon budgets. As already introduced in section 4.2, these budgets represent the maximum cumulative amount of CO_2 emissions allowed to reach a given temperature goal. To account for the uncertainties in the calculation of these budgets, scientists and policy makers often refer to the *likely* carbon budget. This is the amount of CO_2 emissions such that there is a 66% chance of reaching the temperature goal.

For five temperature goals ranging from 1.5°C to 3.5°C, we calculated the distribution of the carbon budget by using Eq. (13.7) and sampling the distributions of the underlying parameters. The resulting histograms are shown in Fig. 14.2. A first approximation to the likely carbon budget is by taking the 34% percentile of these data points, since this means that for 66% of the runs, a higher, less stringent, carbon budget still reaches the same temperature goal.

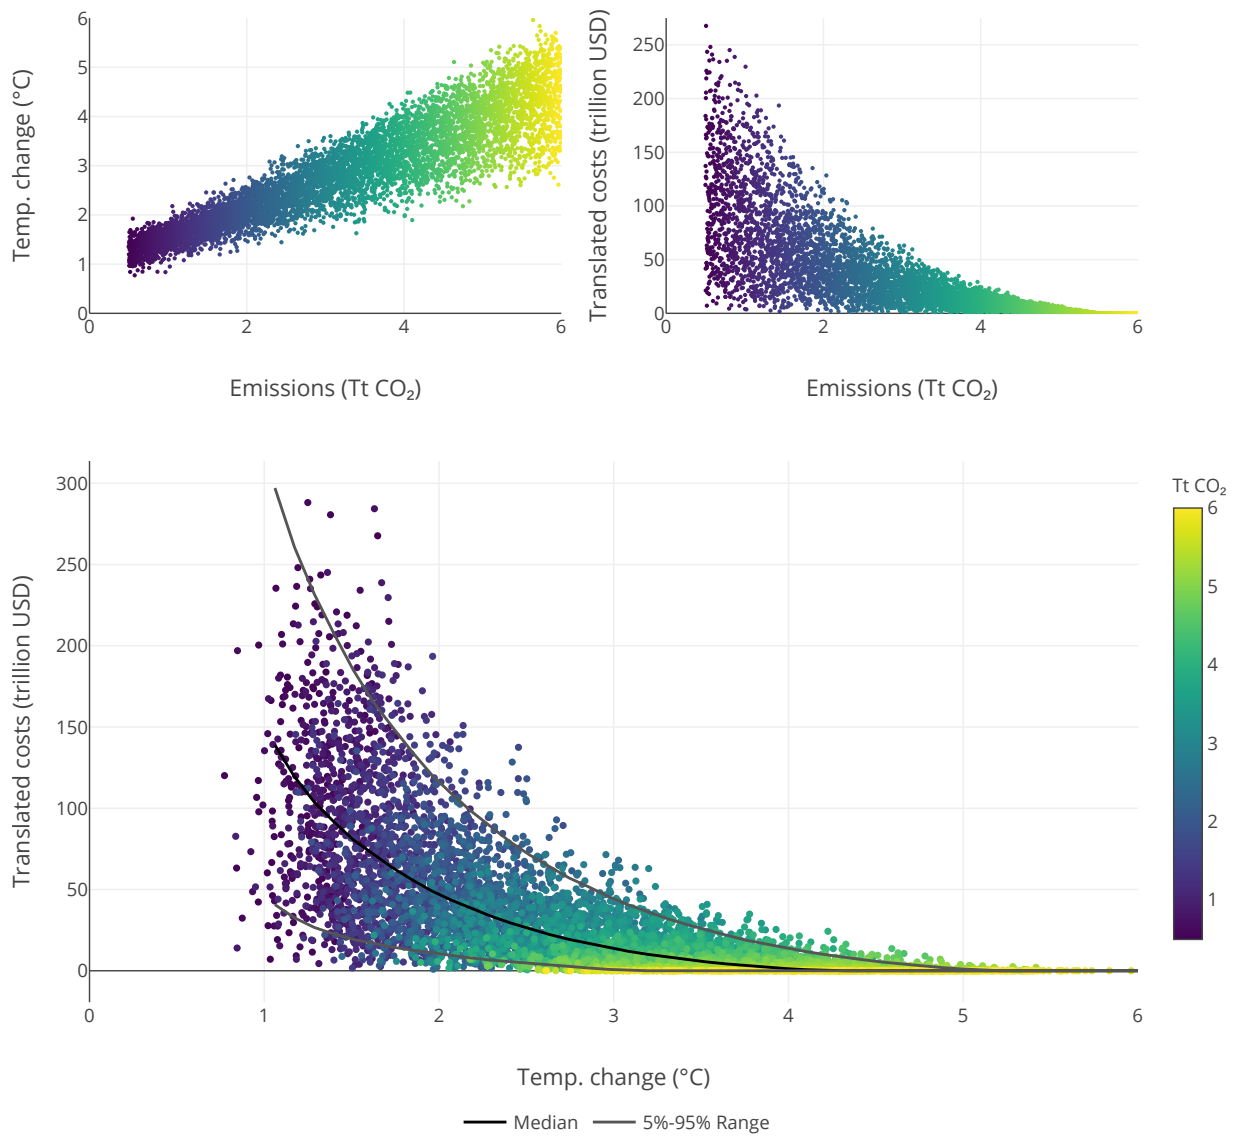


Figure 14.1: Samples of the temperature as function of CO₂ (top left) and costs as function of CO₂ (top right) are merged to obtain the statistical distribution of the costs required to reach a temperature target (bottom).




Figure 14.2: Probability distribution of the carbon budgets given 5 temperature goals. The lines indicate the *likely* carbon budget: with this budget, there is a 66% chance of reaching the temperature goal.

However, this approach is not quite accurate. In fact, since the definition of the likely budget is the amount of CO₂ such that there is a 66% chance of reaching the temperature goal, we are actually dealing with an inverse problem. Let $f(\text{CO}_2; \text{perc})$ be the function which returns the perc percentile of the resulting temperature, given a level of CO₂. This function is calculated by sampling N times the function $T(\text{CO}_2)$ defined in Eq. (13.6), and calculating its perc empirical percentile (the value such that 66% of the sample points are smaller than that temperature value). The *likely* carbon budget LCB is then the inverse of this function:

$$\text{LCB}(T) = f^{-1}(T; 66\%)$$

This yields the following likely carbon budgets, which we compare with the original, first approximation of the carbon budget:

T	Carbon budget (Tt CO ₂)	
	$\text{LCB}(T)$	first approximation
1.5 °C	0.852	0.861
2.0 °C	1.634	1.644
2.5 °C	2.410	2.420
3.0 °C	3.171	3.187
3.5 °C	3.933	3.954

These carbon budgets are significantly higher than the carbon budgets given in the 2014 IPCC report, namely 0.4 TtCO₂ for 1.5°C, 1.0 TtCO₂ for 2.0°C and 2.4 TtCO₂ for 3.0°C warming. These are carbon budgets from 2011, whereas the values in the previous table start in 2010, which means that about 40 GtCO₂ should be added to the IPCC values for the extra year. However, even then, the differences remain large.

Interestingly, as of writing this paragraph on 1st of October, a new IPCC report is being discussed at the 48th session of the IPCC in South Korea, in which new, updated carbon budgets will be proposed which are much closer to the carbon budgets found in our model.

14.2. SENSITIVITY ANALYSIS

As described above, each parameter used in the model has an underlying distribution: the climate sensitivity TCRE was determined using the IPCC pink plume, the current global mean temperature T_{2010} from a recent publication (Visser et al., 2018), the variance from the contribution to climate forcing from non CO₂ gases, $\sigma_{\text{non CO}_2}$, where the uncertainty mainly comes from different societal scenarios in the SSP database, and finally p , representing the uncertainty in the costs required to reach a certain carbon budget.

Using the model described in section 13, we analyse the contribution of these individual parameter uncertainties to the total uncertainty of the costs as function of the temperature goal. This is performed using the variance decomposition method introduced in section 12.2: by sampling the distributions and keeping certain sets of parameters fixed, we can split the total variance into partial variances. Specifically, we can estimate the *first order* variances S_i , giving the influence of purely the parameter X_i individually, without taking into account the interaction between terms. These interactions are calculated using higher order terms. Besides the first order terms, we also focus on the *second order* terms S_{ij} , giving the variance occurring specifically from the interaction between parameter X_i and X_j (for example, between TCRE and $\sigma_{\text{non CO}_2}$). We calculate these partial variances for a range of temperature goals between 1°C and 5°C. For lower temperatures, the mitigation costs are too little know, and for higher temperatures, the costs are chosen to be zero anyway.

The results from this calculation are shown in Fig. 14.3. First, by looking only at the first order terms (the four bottom most areas in the figure), the main observation is that for low temperature goals (1.5°C- 2°C), the uncertainty in the abatement cost dominate the output variance, while the higher the temperature goal, the more important the uncertainty of TCRE becomes. The impact of the first order term for $\sigma_{\text{non CO}_2}$ is rather constant over the temperature range, accounting for, on average, between 5% and 10% of the total variance.

For temperatures higher than 4°C, second order interaction terms become increasingly important. In fact, for a 5.0°C temperature goal, the interaction between TCRE and $\sigma_{\text{non CO}_2}$ and the interaction between TCRE and p alone are both responsible for 15% of the variance.

The results of the variance decomposition for the very low temperature range should be considered very cautiously. In fact, while mathematically correct, they have been obtained by extrapolating the cost function to carbon budgets way below the 0.5 TtCO₂ for which we still have data on the costs. We are thereby using information that is actually unknown, since there are no cost estimates of such low carbon budgets. In fact, such scenarios would require an almost immediate stop of global CO₂ emissions, along with huge amounts of negative emission technologies, which is considered unfeasible. For this reason, we suggest to only consider the results from Fig. 14.3 for temperature values starting from 1.4 or 1.5°C.

Lastly, we observe that the uncertainty in T_{2010} is too small to cause any significant variance in the model output.




Figure 14.3: Variance decomposition for different temperatures. We included the first order partial variances (bottom four: ‘1st: TCRE’, ‘1st: T₂₀₁₀’, ‘1st: p’ and ‘1st: σ_{nonCO₂}’ for the partial variances of the parameters TCRE, T₂₀₁₀, p and σ_{nonCO₂} respectively), the second order partial variances (where the interaction term between a parameter A and a parameter B is denoted by ‘2nd: param A ↔ param B’) and one of the third order terms. The rest is combined in the *others* term.

Chapter 15

Wrap-up of part II

15.1. CONCLUSION

In the second part of this thesis, we developed a simple IAM expressing the mitigation costs as function of a temperature goal. This model uses four parameters to capture the high level geophysical and socio-economic relations provided in current literature:

- First, the climate sensitivity was derived from the iconic IPCC pink plume of Fig. 11.1. While the value of TCRE obtained already accounts for the influence of non-CO₂ gases, the uncertainty due to these non-CO₂ gases was not included.
- The second parameter, $\sigma_{\text{non-CO}_2}$ captures this uncertainty, and was derived from a wide range of scenarios from the Shared Socio-economic Pathways database.
- The third parameter gives the global temperature in 2010, relative to pre-industrial times, as found in current literature.
- The fourth parameter, p , comprises the uncertainty in the costs of reaching a carbon budget. The costs were modelled as a linear combination of an exponential function fitted through the minimum of the SSP cost points and a similar exponential function fitted through the maximum of those points.

Using our IAM, we calculated carbon budgets, defined as the amount of CO₂ emissions that could still be emitted such that a temperature goal is reached with a 66% probability. These budgets were significantly higher than the budgets proposed in the last IPCC Assessment Report. Especially for stringent budgets, our carbon budget was almost twice as large as the AR5 budget. However, the IPCC budgets have been increased in the recent 1.5°C Special Report, coming closer to our budgets.

The original reason we developed our simple IAM was to investigate the partial variance resulting from each parameter. Sobol's Monte Carlo method provided this variance decomposition. As main result, we observe that for stringent temperature goals, the socio-economic uncertainty prevails over the geophysical uncertainty. The inverse holds for looser goals, where the geophysical uncertainty dominates the total variance. Because of the importance of stringent temperature goals, like the 1.5°C or 2°C goals in the Paris Agreement, socio-economic modelling uncertainties should first be reduced to improve the total mitigation cost estimate for the most policy-relevant goals.

This result could require a possible shift in funding priorities, from more geophysical research towards more socio-economic modelling. However, such a conclusion is not straightforward, and requires further investigation, as detailed in the Discussion section.

The main value of our uncertainty analysis resides in the fact that we provide a clear dissection of the sources of uncertainties of the future mitigation costs for different temperature goals.

15.2. DISCUSSION

MODEL PARAMETERS

While the estimation of every model parameter used in our IAM represents an approximation of current literature, we acknowledge that some assumptions would benefit from further research.

- In our model, we have assumed that the contribution of non-CO₂ gases can be simply included as a variance term with zero mean. However, this constant variance assumption might not be correct. Moreover, fitting a line through the non-CO₂ SSP points is not the only option when fitting a function through them. For instance, a convex function could be investigated, which would be flatter through scenarios with low CO₂ emissions, and steeper for high CO₂ emissions.
- The assumption that all three cost measures can be compared by indexing them could also be researched further. Mainly, the large price difference between Area under the MAC and the two other units should be investigated before translating indexed costs to more relevant units, like percentage of baseline GDP loss.
- Some argue that mitigation costs can also become negative for stringent carbon budgets, because of radical economic transformations and climate benefits. While there is no direct data supporting this hypothesis in the SSP database, possibly negative mitigation costs can be modelled by choosing a probability distribution for p that is still asymmetric, but allows for negative values as well with a small probability.
- When communicating the mitigation costs as function of the temperature goal, one should also always compare them with the benefits from reduced climate change, since this is the reason that scientists, policy makers and politicians are pushing for higher emission reductions.

CARBON BUDGETS

As mentioned in the conclusion, our calculated carbon budgets deviate from the previous IPCC estimations. In the 1.5°C IPCC Special Report, published in October 2018, new estimates of carbon budgets are provided, higher than those in the AR5 report. Further comparison is required to test the legitimacy and accuracy of the carbon budgets calculated using our IAM.

UNCERTAINTY CONTRIBUTION

A surprising result from our uncertainty dissection is the small contribution of the T_{2010} temperature uncertainty. Especially for small carbon budgets, one might expect a 0.1°C uncertainty to have a higher influence on the carbon budget than suggested by our analysis. This effect might be cancelled by the exponential increase in mitigation cost uncertainty for small carbon budgets. One way to test this is by performing the full analysis without the cost factor.

Finally, given our research, it might be too soon to suggest shifts in research priorities. One might argue that economical models are so uncertain that it might not be possible to reduce the mitigation cost uncertainty, because of the very nature of our society. Moreover, it might be possible that more funding to socio-economic IAM modelling would actually increase the uncertainty. However, if this last consequence would be true, it would suggest that we are currently underestimating the socio-economic uncertainty, which would also be an important finding.

Bibliography

- M. R. Allen, D. J. Frame, C. Huntingford, C. D. Jones, J. A. Lowe, M. Meinshausen, and N. Meinshausen. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242):1163–1166, 2009. ISSN 00280836. doi: 10.1038/nature08019. URL <http://dx.doi.org/10.1038/nature08019>.
- K. J. Arrow and E. L. Lehmann. *Biographical Memoirs*, volume 87. The National Academies Press, 1979. ISBN 978-0-309-02549-2. doi: 10.17226/573. URL <http://www.nap.edu/catalog/573>.
- N. Bauer, K. Calvin, J. Emmerling, O. Fricko, S. Fujimori, J. Hilaire, J. Eom, V. Krey, E. Kriegler, I. Mouratiadou, H. Sytze de Boer, M. van den Berg, S. Carrara, V. Daioglou, L. Drouet, J. E. Edmonds, D. Gernaat, P. Havlik, N. Johnson, D. Klein, P. Kyle, G. Marangoni, T. Masui, R. C. Pietzcker, M. Strubegger, M. Wise, K. Riahi, and D. P. van Vuuren. Shared Socio-Economic Pathways of the Energy Sector – Quantifying the Narratives. *Global Environmental Change*, 42:316–330, 2017. ISSN 09593780. doi: 10.1016/j.gloenvcha.2016.07.006. URL <http://dx.doi.org/10.1016/j.gloenvcha.2016.07.006>.
- J. Berny. A New Distribution Function for Risk Analysis. *Journal of the Operational Research Society*, 40(12): 1121–1127, 1989. URL <https://www.jstor.org/stable/2582921>.
- P. Billingsley. Section 34. Conditional Expectation. In *Probability and Measure*, page 445. John Wiley & Sons, 1995.
- V. Bosetti, M. Tavoni, E. De Cian, and A. Sgobbi. The 2008 Witch Model: New Model Features and Baseline. *Ssrn*, 2009. doi: 10.2139/ssrn.1512475.
- P. Chebyshev. Théorie des mécanismes connus sous le nom de parallélogrammes. *Mémoires des Savants étrangers présentés à l'Académie de Saint-Petersbourg*, 7:539–586, 1854.
- R. Courant and D. Hilbert. *Methods of Mathematical Physics*. 1:560, 1953. doi: 10.1016/B978-0-12-373750-2.50021-5.
- P. Dasgupta. Discounting climate change. *Journal of Risk and Uncertainty*, 37(2-3):141–169, 2008. ISSN 08955646. doi: 10.1007/s11166-008-9049-6.
- J. Emmerling, L. Drouet, L. A. Reis, M. Bevione, L. Berger, V. Bosetti, S. Carrara, E. D. Cian, M. Malpede, G. Marangoni, F. Sferra, M. Tavoni, and P. Havlik. No Title. *Working Paper, Fondazione Eni Enrico Mattei*, 2016.
- W. H. Fleming and R. W. Rishel. *Deterministic and stochastic optimal control*. 1975. ISBN 9781461263807.
- C. Gough, S. Mander, P. Thornley, A. Lea-Langton, and N. Vaughan. Understanding Negative Emissions From BECCS. In *Biomass Energy with Carbon Capture and Storage (BECCS): Unlocking Negative Emissions*. 2018.
- L. H. Goulder and K. Mathai. Optimal CO₂ abatement in the presence of induced technological change. *Journal of Environmental Economics and Management*, 39(1):1–38, 2000. ISSN 00950696. doi: 10.1006/jeem.1999.1089.
- A. Grimaud and L. Rougé. Polluting non-renewable resources, innovation and growth: Welfare and environmental policy. *Resource and Energy Economics*, 27(2):109–129, 2005. ISSN 09287655. doi: 10.1016/j.reseneeco.2004.06.004.

- M. Ha-Duong, M. J. Grubb, and J. C. Hourcade. Influence of socioeconomic inertia and uncertainty on optimal CO₂- emission abatement. *Nature*, 390(6657):270–273, 1997. ISSN 00280836. doi: 10.1038/36825.
- N. Hanley, J. F. Shogren, and B. White. *Environmental Economics in Theory and Practice*. 1997.
- T. Homma and A. Saltelli. Importance measures in global sensitivity analysis of nonlinear models. *Reliability Engineering & System Safety*, 52:1–17, 1996. ISSN 09518320. doi: 10.1016/0951-8320(96)00002-6. URL <http://www.sciencedirect.com/science/article/pii/0951832096000026>.
- H. Hotelling. The University of Chicago Press. *Journal of Political Economy*, 39(2):137–175, 1931. ISSN 10773711. doi: 10.1086/521238.
- IPCC. *Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. 2014. ISBN 9781107654815. doi: 10.1017/CBO9781107415416. URL <http://www.ipcc.ch/report/ar5/wg3/>.
- S. Kahouli-Brahmi. Technological learning in energy-environment-economy modelling: A survey. *Energy Policy*, 36(1):138–162, 2008. ISSN 03014215. doi: 10.1016/j.enpol.2007.09.001.
- H. J. Kappen. Optimal control theory and the linear Bellman Equation. In & S. C. D. Barber, A. Cemgil, editor, *Bayesian Time Series Models*, pages 363–387. Cambridge University Press, 2011. ISBN 9780511984679. doi: <http://dx.doi.org/10.1017/CBO9780511984679.018>.
- G. Knowles. *An introduction to Applied Optimal Control*. Mathematics in Science and Engineering, volume 159 edition, 1981. ISBN 0-12-416960-0. doi: 10.15713/ins.mmj.3.
- E. Kriegler, J. Edmonds, S. Hallegatte, K. L. Ebi, T. Kram, K. Riahi, H. Winkler, and D. P. van Vuuren. A new scenario framework for climate change research: The concept of shared climate policy assumptions. *Climatic Change*, 122(3):401–414, 2014. ISSN 01650009. doi: 10.1007/s10584-013-0971-5.
- F. G. Li and N. Strachan. Modelling energy transitions for climate targets under landscape and actor inertia. *Environmental Innovation and Societal Transitions*, 24(December 2017):106–129, 2017. ISSN 22104224. doi: 10.1016/j.eist.2016.08.002.
- J. Livernois. On the empirical significance of the hotelling rule. *Review of Environmental Economics and Policy*, 3(1):22–41, 2009. ISSN 17506816. doi: 10.1093/reep/ren017.
- A. Löschel. Technological change in economic models of environmental policy: A survey. *Ecological Economics*, 43(2-3):105–126, 2002. ISSN 09218009. doi: 10.1016/S0921-8009(02)00209-4.
- J. Macki and A. Strauss. *Introduction to Optimal Control Theory*, volume 51. 1982. ISBN 3540901639. doi: 10.1007/978-0-387-92712-1.
- S. Marsman. Master Thesis The uncertainty in meeting climate targets — an assessment of our current knowledge, 2018.
- H. D. Matthews, N. P. Gillett, P. A. Stott, and K. Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832, 2009. ISSN 00280836. doi: 10.1038/nature08047. URL <http://dx.doi.org/10.1038/nature08047>.
- M. Meinshausen, S. C. Raper, and T. M. Wigley. Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 - Part 1: Model description and calibration. *Atmospheric Chemistry and Physics*, 11(4):1417–1456, 2011. ISSN 16807316. doi: 10.5194/acp-11-1417-2011.
- W. D. Nordhaus. Rolling the ‘DICE’: An optimal transition path for controlling greenhouse gases. *Resource and Energy Economics*, 15:27–50, 1993.

- B. C. O'Neill, E. Kriegler, K. Riahi, K. L. Ebi, S. Hallegatte, T. R. Carter, R. Mathur, and D. P. van Vuuren. A new scenario framework for climate change research: The concept of shared socioeconomic pathways. *Climatic Change*, 122(3):387–400, 2014. ISSN 01650009. doi: 10.1007/s10584-013-0905-2.
- H. Pesch and M. Plail. The Cold War and the maximum principle of optimal control. *Optimization Stories. Documenta Mathematica*, I(2009):331–343, 2012. URL <http://www.emis.ams.org/journals/DMJDMV/vol-ismj/48{ }pesch-hans-josef-cold-war.pdf>.
- F. P. Ramsey. A Mathematical Theory of Saving. *The Economic Journal*, 38(152):543–559, 1928.
- A. Saltelli. Making best use of model evaluations to compute sensitivity indices. 145:280–297, 2002.
- A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global Sensitivity Analysis. The Primer*. 2008. ISBN 9780470059975 (ISBN). doi: 10.1002/9780470725184. URL <http://www.scopus.com/inward/record.url?eid=2-s2.0-84889461919{ }partnerID=40{ }&md5=367cd0af66837fc6e4ae661b69007cd9>.
- I. Sobol. On sensitivity estimation for nonlinear mathematical models. 2(1):112–118, 1990.
- The Editors of Encyclopaedia Britannica. Lev Semyonovich Pontryagin, 2018. URL <https://www.britannica.com/biography/Lev-Semyonovich-Pontryagin>.
- D. P. Van Vuuren. *Energy systems and climate policy - long-term scenarios for an uncertain future*. 2007.
- D. P. van Vuuren, E. Kriegler, B. C. O'Neill, K. L. Ebi, K. Riahi, T. R. Carter, J. Edmonds, S. Hallegatte, T. Kram, R. Mathur, and H. Winkler. A new scenario framework for Climate Change Research: Scenario matrix architecture. *Climatic Change*, 122(3):373–386, 2014. ISSN 01650009. doi: 10.1007/s10584-013-0906-1.
- H. Visser, S. Dangendorf, D. P. Van Vuuren, B. Bregman, and A. C. Petersen. Signal detection in global mean temperatures after "Paris": An uncertainty and sensitivity analysis. *Climate of the Past*, 14(2):139–155, 2018. ISSN 18149332. doi: 10.5194/cp-14-139-2018.

Appendix A

Derivations

A.1. ANALYTICAL RESULTS

A.1.1. LESS THAN EXPONENTIAL GROWTH OF CARBON TAX IN LBD CASE

We have shown in section 8.1.3 that, while the carbon tax is given by $p(t) = \lambda(t)e^{rt}$, the growth of $\lambda(t)$ is given, in the learning by doing case, by:

$$\begin{aligned}\dot{\lambda}(t) &= \frac{\partial \mathcal{H}}{\partial \text{CE}(t)} \\ &= -e^{-rt}B(t) \left[\frac{\partial}{\partial \text{CE}(t)} \left(\int_0^{\text{MAC}^{-1}(p(t); \text{CA}(t))} \text{MAC}(a; \text{CA}(t)) da \right) \right] + \lambda(t)B(t) \frac{\partial}{\partial \text{CE}(t)} \text{MAC}^{-1}(p(t); \text{CA}(t))\end{aligned}$$

By using the Fundamental Theorem of Calculus, the chain rule and the product rule, this expression becomes:

$$\begin{aligned}\dot{\lambda}(t) &= -e^{-rt}B(t) \left[\int_0^{\text{MAC}^{-1}(p(t); \text{CA}(t))} \left(\frac{\partial}{\partial \text{CE}(t)} \text{MAC}(a; \text{CA}(t)) \right) da \right. \\ &\quad \left. + \text{MAC}(\text{MAC}^{-1}(p(t); \text{CA}(t)); \text{CA}(t)) \frac{\partial}{\partial \text{CE}(t)} \text{MAC}^{-1}(p(t); \text{CA}(t)) \right] \\ &\quad + \lambda(t)B(t) \frac{\partial}{\partial \text{CE}(t)} \text{MAC}^{-1}(p(t); \text{CA}(t)) \\ &= -e^{-rt}B(t) \left[\int_0^{\text{MAC}^{-1}(p(t); \text{CA}(t))} \left(\frac{\partial}{\partial \text{CE}(t)} \text{MAC}(a; \text{CA}(t)) \right) da + p(t) \frac{\partial}{\partial \text{CE}(t)} \text{MAC}^{-1}(p(t); \text{CA}(t)) \right] \\ &\quad + \lambda(t)b(t) \frac{\partial}{\partial \text{CE}(t)} \text{MAC}^{-1}(p(t); \text{CA}(t))\end{aligned}$$

By noting that $\lambda(t) = e^{-rt}p(t)$, the last two terms cancel, giving the much shorter expression:

$$\dot{\lambda}(t) = -e^{-rt}B(t) \int_0^{\text{MAC}^{-1}(p(t); \text{CA}(t))} \left(\frac{\partial}{\partial \text{CE}(t)} \text{MAC}(a; \text{CA}(t)) \right) da$$

Finally, since we have defined the cumulative abatement to be equal to the cumulative baseline minus the cumulative emissions, $CA(t) = \left(\int_0^t B(s)ds\right) - CE(t)$, we can replace $\partial MAC(a; CA(t))/\partial CE(t)$ by:

$$\begin{aligned}\frac{\partial}{\partial CE(t)}MAC(a; CA(t)) &= \frac{\partial}{\partial CA(t)}MAC(a; CA(t))\frac{\partial CA(t)}{\partial CE(t)} \\ &= -\frac{\partial}{\partial CA(t)}MAC(a; CA(t))\end{aligned}$$

Putting this back in the expression for $\dot{\lambda}(t)$, we obtain:

$$\dot{\lambda}(t) = e^{-rt}B(t) \int_0^{MAC^{-1}(p(t); CA(t))} \left(\frac{\partial}{\partial CA(t)}MAC(a; CA(t)) \right) da,$$

which is equal to Eq. (8.6).

A.1.2. CLOSED FORM STATICS FOR CONVEX MAC AND CONCAVE BASELINE

Using the functional forms of the baseline and the MAC defined in Eq. 8.1 and 8.2, we can start by using Eq. 8.4 to calculate the initial carbon tax in this setting. This means that we have to calculate the following integral:

$$\begin{aligned}\int_0^T B(t)(1 - \alpha - MAC^{-1}(p_0 e^{rt}))dt &= 0 \\ \int_0^T (b_0 + b_1t + b_2t^2) \left(1 - \alpha - \left(\frac{p_0 e^{rt}}{\beta_0} \right)^{1/\beta_1} \right) dt &= 0\end{aligned}$$

By noting that $(e^{rt})^{1/\beta_1} = e^{rt/\beta_1}$ and using integration by parts, this integral can be evaluated, giving:

$$\begin{aligned}\Leftrightarrow (1 - \alpha) \left(b_0T + \frac{1}{2}b_1T^2 + \frac{1}{3}b_2T^3 \right) + \left(\frac{p_0}{\beta_0} \right)^{1/\beta_1} \left[\frac{b_0\beta_1}{r} - \frac{b_1\beta_1^2}{r^2} + \frac{2b_2\beta_1^3}{r^3} \right] \\ - \left(e^{rT} \frac{p_0}{\beta_0} \right)^{1/\beta_1} \left[\left(\frac{b_0\beta_1}{r} - \frac{b_1\beta_1^2}{r^2} + \frac{2b_2\beta_1^3}{r^3} \right) + b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1}{r} T^2 - \frac{2\beta_1^2}{r^2} T \right) \right] &= 0 \\ \Leftrightarrow (1 - \alpha) \left(b_0T + \frac{1}{2}b_1T^2 + \frac{1}{3}b_2T^3 \right) - \left(\frac{p_0}{\beta_0} \right)^{1/\beta_1} \left[\left(\frac{b_0\beta_1}{r} - \frac{b_1\beta_1^2}{r^2} + \frac{2b_2\beta_1^3}{r^3} \right) (e^{rT/\beta_1} - 1) \right. \\ \left. + e^{rT/\beta_1} \left(b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1}{r} T^2 - \frac{2\beta_1^2}{r^2} T \right) \right) \right] &= 0\end{aligned}$$

While this equation is still very large, we can find the expression for p_0 from it directly:

$$p_0 = \beta_0 \left(\frac{(1 - \alpha) \left(b_0T + \frac{1}{2}b_1T^2 + \frac{1}{3}b_2T^3 \right)}{\left(\frac{b_0\beta_1}{r} - \frac{b_1\beta_1^2}{r^2} + \frac{2b_2\beta_1^3}{r^3} \right) (e^{rT/\beta_1} - 1) + e^{rT/\beta_1} \left(b_1 \frac{\beta_1 T}{r} + b_2 \left(\frac{\beta_1}{r} T^2 - \frac{2\beta_1^2}{r^2} T \right) \right)} \right)^{\beta_1},$$

which is indeed the expression shown in Eq. 8.5.

Next, we can also calculate the ratio between net positive E_+ and net negative E_- emissions. In other words, E_+ is the integral of the emissions from the beginning to t^* , and E_- the integral from t^* to T . Just like the expression for E_-/E_+ for the constant baseline / linear MAC case, the full expression for E_-/E_+ is a fraction with mostly the same terms in the numerator and denominator, except for a few extra terms in the nominator:

$$\frac{E_-}{E_+} = \frac{\text{constant terms} + \text{fraction term} + \text{log terms} - \alpha r^3 \left(b_0T + \frac{b_1T^2}{2} + \frac{b_2T^3}{3} \right)}{\text{constant terms} + \text{fraction term} + \text{log terms}},$$

where we have to define the following quantities:

$$\text{constant terms} = -(b_0 r^2 \beta_1 - b_1 r \beta_1^2 + 2b_2 \beta_1^3)$$

$$\text{fraction term} = \frac{r^3(1-\alpha) \left(b_0 T + \frac{b_1 T^2}{2} + \frac{b_2 T^3}{3} \right)}{-1 + e^{rT/\beta_1} \left(1 + \frac{b_1 r^2 T + b_2 r^2 T^2 - 2b_2 r T \beta_1}{b_0 r^2 - b_1 r \beta_1 + 2b_2 \beta_1^2} \right)}$$

$$\text{log terms} = \ln(\Omega) (b_0 r^2 \beta_1 - b_1 r \beta_1^2 + 2b_2 \beta_1^3) - \frac{1}{2} \ln^2(\Omega) (-b_1 r \beta_1^2 + 2b_2 \beta_1^3) + \frac{1}{6} \ln^3(\Omega) (2b_2 \beta_1^3),$$

and finally the quantity Ω , introduced purely for readability:

$$\Omega = \frac{b_0 r^2 \beta_1 - b_1 r \beta_1^2 + 2b_2 \beta_1^3}{\text{fraction term}}$$

A.2. SOBOL'S VARIANCE-BASED SENSITIVITY ANALYSIS

A.2.1. ZERO-MEAN CONDITION FOR DECOMPOSITION IN SUMMANDS OF DIFFERENT DIMENSIONS

While the definition of the decomposition terms using conditional expectations (Eq. 12.2) has an intuitive meaning, they still have to satisfy the zero-mean condition. We show this for f_i and f_{ij} , but the same reasoning can be generalised. Starting with f_i , we see:

$$\begin{aligned} \int_0^1 f_i(X_i) dX_i &= \int_0^1 \mathbb{E}[Y|X_i] - f_0 dX_i \\ &= \int_0^1 \mathbb{E}[Y|X_i] dX_i - f_0 \end{aligned}$$

Since the random variable X_i is uniformly distributed between 0 and 1, the integral of the conditional expectation can be written as:

$$\int_0^1 \mathbb{E}[Y|X_i] dX_i = \mathbb{E}[Y],$$

as shown in [Billingsley \(1995\)](#). Using the fact that $f_0 := \mathbb{E}[Y]$, we indeed obtain:

$$\int_0^1 f_i(X_i) dX_i = 0.$$

Using similar steps for f_{ij} , we obtain successively:

$$\begin{aligned} \int_0^1 f_{ij}(X_i, X_j) dX_i &= \int_0^1 \mathbb{E}[Y|X_i, X_j] - f_i(X_i) - f_j(X_j) - f_0 dX_i \\ &= \int_0^1 (\mathbb{E}[Y|X_i, X_j] - f_i(X_i)) dX_i - f_j(X_j) - f_0 \end{aligned}$$

We have just shown that the second term in this integral is 0, leaving us with only:

$$\int_0^1 f_{ij}(X_i, X_j) dX_i = \int_0^1 \mathbb{E}[Y|X_i, X_j] dX_i - f_j(X_j) - f_0.$$

Again, the integral of the conditional expectation can be written as $\int_0^1 \mathbb{E}[Y|X_i, X_j] dX_i = \mathbb{E}[Y|X_j]$:

$$\begin{aligned} \int_0^1 f_{ij}(X_i, X_j) dX_i &= \mathbb{E}[Y|X_j] - f_j(X_j) - f_0 \\ &= \mathbb{E}[Y|X_j] - (\mathbb{E}[Y|X_j] - f_0) - f_0 \\ &= 0, \end{aligned}$$

which is what we had to show.

A.2.2. FROM FUNCTION DECOMPOSITION TO VARIANCE DECOMPOSITION

When having expanded the integral of $f^2(\mathbf{X})$, the resulting sum does not directly look like a sum of variances. Since the full proof is rather cumbersome to read and write, we only show how to rewrite the first few terms to the desired variance form. The other terms require similar calculations.

As a reminder, the total variance was found to be equal to:

$$\text{Var}(Y) = \sum_{k=1}^n \sum_{i_1 < \dots < i_k}^d \int_{K^n} f_{i_1, \dots, i_k}^2(X_{i_1}, \dots, X_{i_k}) d\mathbf{X} - f_0^2$$

The sum can be rewritten in a slightly clearer way, to:

$$\text{Var}(Y) = \sum_{i=1}^k \int_0^1 f_i^2(X_i) dX_i + \sum_{i < j} \int_0^1 \int_0^1 f_{ij}^2(X_i, X_j) dX_i dX_j + \dots + \int_{K^n} f_{12\dots n}^2(\mathbf{X}) d\mathbf{X}$$

We now focus on the terms of the first sum. By using the definition of f_i , we obtain successively:

$$\begin{aligned} \int_0^1 f_i^2(X_i) dX_i &= \int_0^1 (\mathbb{E}[Y|X_i] - f_0)^2 dX_i \\ &= \int_0^1 \mathbb{E}[Y|X_i]^2 dX_i - 2f_0 \int_0^1 \mathbb{E}[Y|X_i] dX_i + f_0^2 \end{aligned}$$

As we've seen before, the second integral is equal to $\mathbb{E}[Y] = f_0$, giving:

$$\begin{aligned} \Leftrightarrow \int_0^1 f_i^2(X_i) dX_i &= \int_0^1 \mathbb{E}[Y|X_i]^2 dX_i - f_0^2 \\ &= \int_0^1 \mathbb{E}[Y|X_i]^2 dX_i - (\mathbb{E}[Y])^2 \end{aligned}$$

Since X_i is uniformly distributed, the integral $\int_0^1 g(X_i) dX_i$ is equal to the expectation value $\mathbb{E}[g(X_i)]$. Moreover, the second term, $(\mathbb{E}[Y])^2$ can be written as:

$$(\mathbb{E}[Y])^2 = \left(\mathbb{E}[\mathbb{E}[Y|X_i]] \right)^2$$

using the Law of Total Expectation. Finally, this leads us to:


$$\begin{aligned} \Leftrightarrow \int_0^1 f_i^2(X_i) dX_i &= \mathbb{E} \left[\mathbb{E}[Y|X_i]^2 \right] - \left(\mathbb{E}[\mathbb{E}[Y|X_i]] \right)^2 \\ &= \text{Var}_{X_i}(\mathbb{E}[Y|X_i]), \end{aligned}$$

which is indeed the definition of the variance. These steps can be repeated for the higher order terms.


Appendix B

Additional figures

B.1. OPTIMAL CARBON TAX IN LOG-PLOT



(a) Learning Over Time: linear MAC, cost reduction purely depending on time.



(b) Learning By Doing: linear MAC, cost reduction depending on cumulative abatement.

Figure B.1: For the carbon budget $\alpha = 0.3$, optimal carbon tax paths (right) are shown on a log-scale along with corresponding emission paths (left). In gray, the (proxy) Hotelling paths are shown.

B.2. OPTIMAL CARBON TAX UNDER EXTREME LEARNING BY DOING




Figure B.2: When using a progress ratio in the learning by doing scenario of $\rho = 0.4$, the resulting optimal carbon tax path starts to decline around 2070.

B.3. COST REDUCTION LEARNING BY DOING USING CONVEX MAC

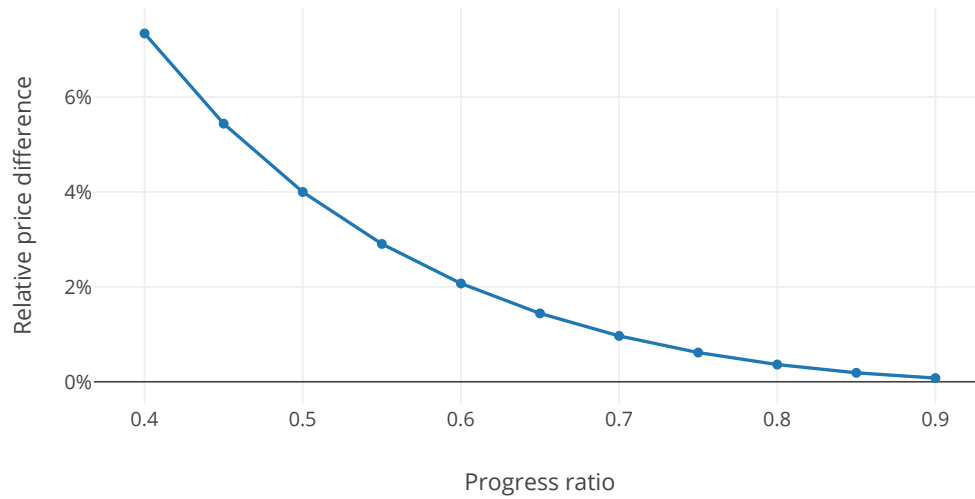


Figure B.3: Just like in Fig. 8.3a, we calculate the cost difference between the optimal tax path and the proxy Hotelling path, in the learning by doing scenario. Here, we used a convex (cubic) MAC. The cost reduction is lower than when using a linear MAC, but show similar behaviour.

B.4. SIGN OF DERIVATIVES OF NET-ZERO YEAR AND OVERSHOOT

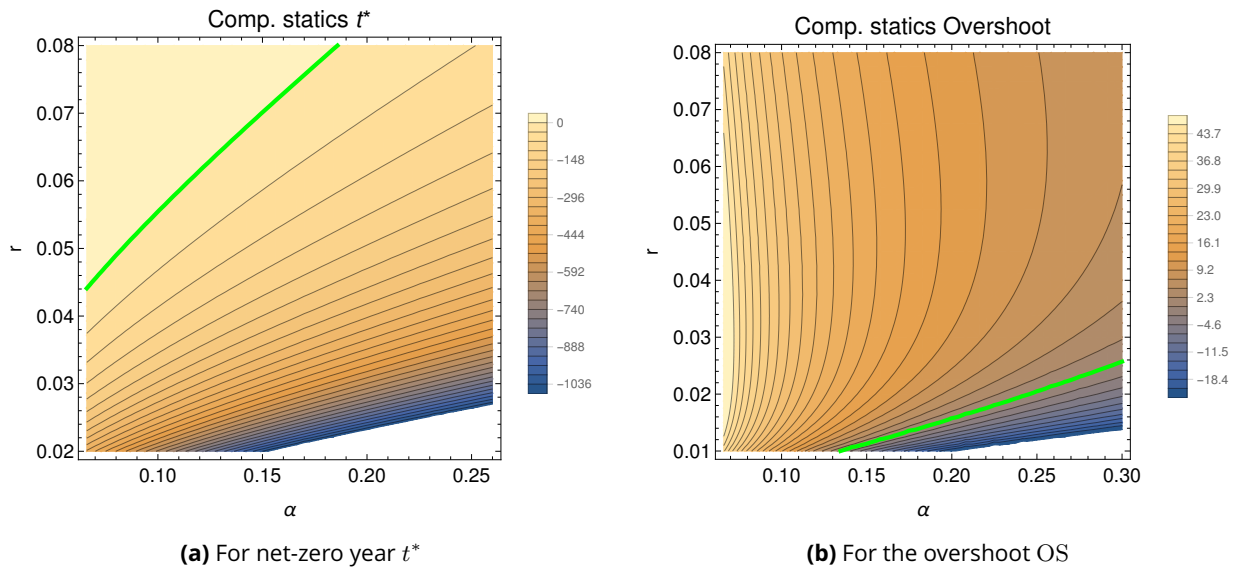


Figure B.4: Value of the derivative of t^* and OS for different values of the discount rate r and the relative carbon budget α . The green line represents the zero-derivative line, showing the boundary between positive and negative values.