Economic consequencess of environmental catastrophes

Titton, A.

Publication date
2025
Document Version
Final published version

This thesis contributes to understanding the economic consequences of climate change and environmental disasters. It does so from three perspectives. First, from the point of view of industrial organisations, it studies firms' diversification incentives and the formation of supply chains in response to possibly correlated environmental shocks. Then, it takes a step back and examines the macroeconomic consequences of rapid and potentially irreversible climate shifts. Finally, it focuses on how different world regions can react to such rapid changes.

Andrea Titton is a PhD student at the University of Amsterdam. His work focuses on environmental economics, specifically the economic consequences and drivers of climate change.

Economic Consequences of Environmental Catastrophes Andrea Titton

874

tinbergen institute
GRADUATE PROGRAM

Economic Consequences of Environmental Catastrophes

Andrea Titton

Universiteit van Amsterdam

Economic Consequences of Environmental Catastrophes

Andrea Titton

ISBN 978 90 361 0814 0

Cover image: *Valley of Aosta: Snowstorm, Avalanche, and Thunderstorm*, J. M. W. Turner, 1836–37, oil on canvas, 92.2 × 123 cm. Frederick T. Haskell Collection, Art Institute of Chicago (Acc. No. 1947.513).

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. 874 of the Tinbergen Institute Research Series, established through cooperation between Rozenberg Publishers and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

Economic Consequences of Environmental Catastrophes

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op donderdag 2 oktober 2025, te 16.00 uur

door Andrea Titton
geboren te Feltre

Promotiecommissie

Promotores:	prof. dr. C.G.H. Diks	Universiteit van Amsterdam
	dr. ir. F.O.O. Wagener	Universiteit van Amsterdam
Overige leden:	prof. dr. F. van der Ploeg	Universiteit van Amsterdam
	prof. dr. J. Tuinstra	Universiteit van Amsterdam
	dr. T.R.G.R. Douenne	Universiteit van Amsterdam
	prof. dr. K. Schubert	Université Paris 1 Panthéon-Sorbonne
	prof. dr. S. Dietz	London School of Economics and Political Science

Faculteit Economie en Bedrijfskunde

Contents

Introduction

The surge in environmental catastrophes, driven by climate change, is a pressing concern for the economy. A warming climate alters the conditions in which human economic activity occurs, making previously sporadic events such as large environmental catastrophes more likely and spatially correlated (Stott, 2016). Sudden and irreversible regime shifts are now a concrete possibility (Seaver Wang et al., 2023; Smolders, van Westen and Dijkstra, 2024). These factors must be considered in the decision-making process of both firms and policymakers. Yet, the economic toolbox used by decision-makers was developed in a time of relative environmental stability, when the most significant consequences of climate change were not yet felt. Economic research has been working towards addressing the assumptions that have made traditional economic models unfit to discuss climate change and its consequences[1]. This has brought theory and modelling work to the forefront as we attempt to predict how economic agents and the economy will react to a changing climate. Integrating climate science data and models into economic analysis to provide realistic and quantitatively meaningful assessments has become increasingly necessary.

This path presents significant challenges. Realistic modelling of a highly interconnected economy or a complex climate requires economic agents to keep track of many moving parts when making decisions they perceive to be optimal. This often makes economic models intractable. To face this challenge, one can either forgo agents' optimisation entirely and concentrate on scenario-based analysis or attempt to distil the essential features of reality into a tractable, simplified model. These two approaches entail a trade-off. On the one hand, sticking to optimisation allows one to predict how agents would respond in a new environment and to make assertions regarding optimal policies for welfare. However, the accuracy of one's predictions hinges on the way one distils reality in the model[2]. On the other hand, a scenario-based analysis in which agents do not optimise but follow simple decision rules

[1] I refer the reader to Carvalho and Tahbaz-Salehi (2019) for a review about production networks and Li, Crépin and Lindahl (2024) for a review on climate tipping points.

[2] Two recent examples of such dangerous simplifications are the Hulten theorem (Baqaee, 2018; Hulten, 1978) for production networks and the attempt to integrate climate dynamics into climate economics (Dietz et al., 2021*a*).

or heuristics places the researcher on firmer ground, as it allows them to integrate data and complex models into the economy. The cost is the inability to incorporate agents' behaviour into counterfactual analysis. Economists have faced this trade-off since the inception of the discipline. Nevertheless, as the object of study becomes more complex, as with an interconnected supply chain or a warming climate, the merits and dangers of optimisations become more salient.

In this thesis, I adopt the stance that the benefits of optimisation outweigh the costs. I develop three models, each attempting to capture the key dynamics of the underlying problem while remaining sufficiently simple to make predictions based on agents' optimisation decisions. In the first chapter, I focus on how supply chains are shaped to anticipate significant and correlated disruptions, such as hurricanes or droughts. Building on recent developments in the economic study of production networks, I develop an analytical framework to study the costs and benefits firms face when diversifying their input provision in an environment where the supply chain is deep and upstream producers are susceptible to correlated disruptions. I then aggregate individual incentives to study how the supply chain evolves as a whole when disruptions among upstream producers become more correlated. The findings suggest that supply chains become endogenously less diversified and more fragile as shocks become more correlated. This effect is exacerbated if firms acquire more information on the supply chain structure. The model raises concerns about the ability of firms to endogenously diversify in anticipation of more frequent and spatially correlated supply chain disruptions. It also suggests that attempting to fix this by incentivising firms to map their supply chains can backfire and lead to sparser and more fragile networks.

In the second chapter, I take a macroeconomic perspective and study how costly it is not to know when a new climate regime could arise. I develop an integrated economy-climate model in which the climate is subject to an endogenous regime shift if a tipping point is crossed. The new regime is irreversible and characterised by larger climate damages. I calibrate the climate model to capture the key dynamics predicted by larger ensembles commonly used in climate science. Using this model, I study how costly it is that the tipping point is not known to a planner scheduling emission abatements. These costs amount to a significant fraction of GDP growth. This highlights how sensitive our current policies are to mistakes or misspecification in climate modelling.

The third chapter extends the second one by considering the potential strategic implications of climate tipping points on different regions' emission abatement policies. I study a model in which OECD countries do not face the repercussions from a tipping point, while the rest of the world does. I demonstrate that OECD coun-

tries can enforce ambitious abatement policies on the rest of the world by scaling back their own. This policy is inefficient and fragile. It is inefficient because every dollar spent on emission abatement in the rest of the world forgoes larger future output growth than a dollar spent in the OECD country, given the higher productivity growth rate in the rest of the world. It is fragile because, if the rest of the world fails to prevent the tipping point, their incentive to abate emissions fall, requiring the OECD countries to suddenly and rapidly curb their own emissions. This result underscores the potential danger of assuming that large environmental catastrophes can coordinate the world towards more ambitious emission abatement paths. When accounting for strategic interactions, tipping points might have the opposite effect, making policymakers succumb to the temptation of forcing more affected regions to bear the brunt of emission abatement. These findings have significant implications for global climate policy, highlighting the need for a coordinated and equitable approach to emission abatement.

Chapter 1

Endogenous Fragility of Supply Chains and Correlated Disruption Risk

In August 2020, hurricane Laura hit one of the world's largest petrochemical districts, in the U.S. states of Louisiana and Texas. As polymer producers in the area were forced to halt production, up to 15% of the country's polypropene producers were unable to source polymer inputs, which in turn caused shortages across the economy (Vakil, 2021). This is just one example of agglomeration of economic activity increasing the correlation between disruptions among suppliers of crucial goods for the economy. In face of such correlated risk, how do downstream producers make sourcing decisions? And, do these decisions yield supply chains resilient to such correlated disruptions?

In this chapter, I study the feedback between the risk of a disruption in sourcing inputs and the endogenous formation of supply chains. A widespread approach to mitigate risk is to diversify it by multisourcing. This practice consists of procuring the same inputs from multiple suppliers, sometimes redundantly (Zhao and Freeman, 2019). Yet, when deciding how many suppliers to source from, a firm faces decreasing marginal benefits in risk reduction, because each additional supplier's failure to deliver is increasingly likely to be correlated with that of the firm's current suppliers. In the presence of marginal costs of sourcing, such as contractual costs or higher prices, the uncertainty behind the correlation of a firm's potential suppliers might induce it to diversify risk less than socially optimal. The wedge between endogenous firm decisions and social optimality arises because downstream firms would be willing to compensate their suppliers for increased diversification of inputs. This under-diversification can generate aggregate fragility in production networks. To understand the relationship between the firm's diversification

decisions and supply chain fragility, I study the properties of a stylised production model. In equilibrium, correlation in the risk of disruption among suppliers generates fragility via two channels. First, it directly introduces endogenous correlation in downstream firms' risk, which amplifies through the production network. This increases the probability of cascading failures, in which the entire production network is unable to produce. Second, it indirectly affects firms' decisions by reducing the expected marginal gain from adding a source of input goods. The latter channel leads to firms diversifying increasingly less, such that small increases in the expected disruption probability can yield fragile production networks.

The role that production networks play in determining economic outcomes has been long recognised. As far back as Leontief (1936), economists have studied how networks in production can act as aggregators of firm-level activity. Following a foundational paper by Hulten (1978), which proved, under strict assumptions, that the first order impact of a productivity shock to an industry is independent of the production network structure, macroeconomics has since de-emphasised this role. Baqaee and Farhi (2019), relaxing these assumptions, illustrated how the structure of the production network can aggregate micro shocks via second order effects.[1] Furthermore, the degree of competition in an industry also interacts with the production network to aggregate shocks, which can lead to cascading failures (Baqaee, 2018). Once established that production networks play a central role in aggregating shocks, two natural questions arise. First, which networks can we expect to observe, given that firms endogenously and strategically choose suppliers? Second, are these endogenous network formations responsible for the growth or fragility that large economies display? These questions fuelled a number of recent papers studying endogenous production network formation. Focusing on growth, Acemoglu and Azar (2020) show that endogenous production networks can be a channel through which firms' increased productivity lowers costs throughout the supply chain and allows for sustained economic growth. In parallel, a vast literature dealt with studying the role of endogenous production networks and firm incentives in determining fragile or resilient economies. Erol and Vohra (2022) showed that in networks with strategic link formation, systemic endogenous fragility arises if the shocks experienced by firms are correlated. Later work, by Amelkin and Vohra (2023), shows that uncertainty in the time of production is crucial in determining whether production networks in equilibrium are sparse and, hence fragile. Finally, Elliott, Golub and Leduc (2022) illustrate how complexity in the production process can also be a key driver of endogenous fragility in production networks. [2]

[1] These results build on a vast literature, for example Acemoglu and Tahbaz-Salehi (2020); Baqaee and Farhi (2019); Carvalho and Tahbaz-Salehi (2019); Carvalho et al. (2020); Gabaix (2011)

[2] The literature on production networks is vast and it is unfortunately impossible to give a fair

A less understood link is that between the correlation of risk within the supply chain, how firms deal with it, and the consequences this has on the economy. Kriegler et al. (2017) studied the effect of uncertainty in endogenous production network formation on firms' productivity and business cycles. They find that higher uncertainty can lead to lower economic growth. In contrast, this chapter focuses on the role of uncertainty in generating endogenous fragility to cascading failures using a more stylised production network model. In line with the existing literature (Elliott, Golub and Leduc, 2022), in the model, small idiosyncratic shocks can be massively amplified. The degree of amplification depends on the equilibrium behaviour of firms. This phenomenon holds true in vertical economies producing simple goods. This chapter extends the analysis of production network formation to an environment in which firms aim to minimise risk while accounting for correlation between suppliers. To do so, I develop a tractable analytical framework that describes the propagation of idiosyncratic shocks through the supply chain when firms make sourcing decisions endogenously in an imperfect information environment. The model describes the evolution of risk through the supply chain as a dynamical system over its depth. The social planner solution shows that endogenous fragility can impose large welfare losses. Importantly, these losses might be discontinuous: an arbitrarily small increase in the correlation of risk among basal firms can generate large welfare losses. Finally, I study a benchmark case where firms have perfect information over idiosyncratic risk. In this case, despite each individual firm being able to achieve a smaller disruption risk, the supply chain is maximally fragile and there is a high probability of large disruptions.

The remainder of the chapter is structured as follows. Section 1.1 discusses the assumptions on the supply chain disruptions and the problem of the firm, and establishes the results that allow to analyse firm sourcing decisions. Section 1.2 derives the law of propagation of the disruption events through the supply chain. Section 1.3 establishes the firm's optimal sourcing strategy and how this endogenously determines the fragility of the supply chain. These results are then compared, in Section 1.4, to the social planner solution to determine the welfare losses induced by the firm's endogenous decisions. Finally, in Section 1.5, the role of imperfect information is isolated by solving the model under perfect information.

overview in this introduction. For a more comprehensive review of the literature, I refer the reader to Carvalho and Tahbaz-Salehi (2019) and Amelkin and Vohra (2023)

1.1 Model

1.1.1 Production Technology and the Firm Objective

Consider an economy producing $K+1$ goods, indexed by $k \in \{0, 1, \dots K\}$. Each firm produces a single good and each good is produced by infinitely many firms indexed by $i \in \mathbb{N}$. Production of the *basal good* $k = 0$ does not require any input and is at risk of random exogenous disruptions in the production process. A *disrupted* basal firm is unable to deliver its good as input to downstream producers. The economy is vertical as each downstream good $k > 0$ requires only good $k-1$ as input. If a firm producing good k is unable to source its input good $k-1$, the firm is itself *disrupted* and hence unable to deliver downstream. In other words, the i-th firm producing good k, indexed by (k, i), is able to produce if at least one of its suppliers is able to deliver, that is, if not all of its suppliers are disrupted. To avoid being disrupted, the firm chooses which firms to source from, among the producers of its input good. In other words, letting $\mathcal{D}_k$ be the random set of disrupted firms in layer k and $\mathcal{S}_{k,i}$ the set of suppliers of firm (k, i), we can say that $(k, i) \in \mathcal{D}_k$ if and only if all of its suppliers $(k-1, j) \in \mathcal{S}_{k,i}$ are in $\mathcal{D}_{k-1}$. I refer to the set of the firm's suppliers $\mathcal{S}_{k,i}$ as its *sourcing strategy*. The disruption events are random and the probability that a firm is disrupted can be written as

$$p_{k,i} := \mathbb{P}\big((k,i) \in \mathcal{D}_k\big) = \mathbb{P}\big(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1}\big). \tag{1.1}$$

Figure 1.1 illustrates this mechanism.

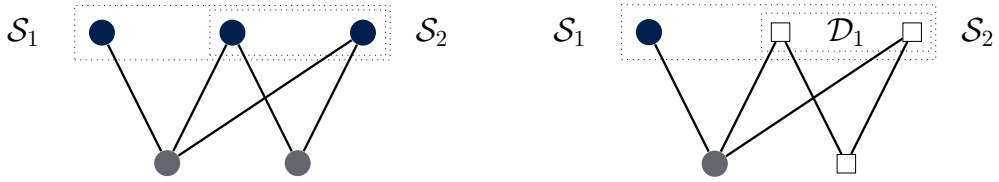

Figure 1.1: The supply chain is depicted in the left panel. The left firm is sourcing its input good from all three suppliers, $\mathcal{S}_1$, while the right firm only from the latter two, $\mathcal{S}_2$. As a disruption occurs, some upstream firms are unable to supply the input good (white box). Unlike the left firm, the right firm is unable to source its inputs and is hence disrupted.

If a firm is not disrupted, it obtains an exogenous profit π. Implementing a given sourcing strategy costs the firm $C\big(|\mathcal{S}_{k,i}|\big)$. The cost C is assumed to be non-negative, increasing in the number of suppliers $|\mathcal{S}_{k,i}|$, and with $C(0) = 0$ and $C(s) \to \infty$ as $s \to \infty$. The problem of firm (k, i) is then to maximise the expected profit[3]

$$\Pi_{k,i}(\mathcal{S}_{k,i}) = \Big(1 - \mathbb{P}\big(\mathcal{S}_{k,i} \subseteq \mathcal{D}_{k-1}\big)\Big)\, \pi - C\big(|\mathcal{S}_{k,i}|\big) \tag{1.2}$$

[3]The expectation is taken over the random set $\mathcal{D}_{k-1}$.

by picking a sourcing strategy $\mathcal{S}_{k,i}$. Before moving to the solution of the model, it is useful to discuss the assumptions presented in this section. The production game is highly stylised: first, firms do not adjust prices but only quantities, such that failure to produce only arises in case no input is sourced; second, they are able to obtain profits by simply producing; third, contracting with new suppliers has a cost; and lastly, there are infinitely many firms in each layer. There are both theoretical and empirical reasons behind these choices. Theoretically, a simpler model allows us to isolate the interplay between the variables of interest: correlation in the risk of suppliers, supply chain opacity, and endogenous production network fragility. Empirically, these assumptions capture well the rationale behind firms' multisourcing. There is strong evidence that firms, first, when faced with supply chain shocks, adjust quantities rather than prices in the short run (di Giovanni and Levchenko, 2010; Jiang, Rigobon and Rigobon, 2022; Lafrogne-Joussier, Martin and Mejean, 2022; Macchiavello and Morjaria, 2015), second, that production shutdowns can have significant costs (Barrot and Sauvagnat, 2016; Hameed and Khan, 2014), and third, that fostering relationships with suppliers is costly, but important in guaranteeing operational performance (Cousins and Menguc, 2006). Finally, modelling infinitely many firms yields tractable analytical results, which can then be mapped back to the finite case by considering a finite subset of firms. Furthermore, this is a good approximation of real supply chains, which are increasingly complex and deep (Elliott, Golub and Leduc, 2022). The model establishes a link between these issues faced by firms when choosing a sourcing strategy and the fragility of the production network.

1.1.2 Imperfect Information and Ex-Ante Symmetry

The supply chain is opaque: firms cannot observe the sourcing decisions of their potential suppliers before making their own. Furthermore, firms do neither know how risky individual basal producers are, nor how their risk is correlated. Yet, firms know the distribution from which the probabilities of disruption in the basal layer are drawn. To motivate this assumption, recall the introductory example of Hurricane Laura. A downstream firm producing polypropene, might not be able to trace back the production steps from its input to individual polymer producers in Louisiana or Texas, and, hence, the exact exposure of its production process to hurricanes. Yet, the firm can estimate the aggregate risk the polymer industry faces in the region. Given this information about the basal layer and their own depth k in the production network, firms can derive the distribution of risk among their suppliers and make sourcing decisions based on it. The risk of two firms downstream sourcing from the same number of suppliers is ex-ante identical, albeit possibly

correlated.

I now formalise this idea in the context of the model. Let

$$X_{k,j} \coloneqq \begin{cases} 1 & \text{if firm } (k,j) \text{ is disrupted and} \\ 0 & \text{otherwise.} \end{cases} \tag{1.3}$$

To encode the uncertainty in the probability $p_{0,j}$ that a basal firm experiences a disruption event $X_{0,j} = 1$, I assume that these probabilities $\{p_{0,j}\}_j$ are independent samples of a random variable P_0, which has a known distribution. I assume this distribution to be the Beta(μ, ρ). The Beta distribution allows to flexibly model shocks that might happen due to the spacial or technological proximity of basal producers, which cannot be diversified. Consider, for example, how oil extraction plants must be located near oil reserves and are hence all subject to correlated weather shocks that might force them to shut down. In this case, despite the small expected probability that an individual firm is disrupted, as a hurricane is a rare occurrence, disruptions are highly correlated, as all firms are vulnerable to the same hurricane.

Definition 1. *A Beta(μ, ρ) distribution, with mean $\mu \in [0,1]$ and correlation $\rho \in [0,1]$, is a continuous probability distribution over the unit interval with cumulative density function, if $\rho \in (0,1)$,*

$$F(p;\mu,\rho) = \begin{cases} 0 & \text{if } p < 0, \\ \frac{1}{C(\mu,\rho)} \int_0^p u^{(1-\mu)\left(\frac{1-\rho}{\rho}\right)-1}(1-u)^{\mu\left(\frac{1-\rho}{\rho}\right)-1}\,\mathrm{d}u & \text{if } 0 \le p \le 1 \\ 1 & \text{if } p > 1, \end{cases} \tag{1.4}$$

where $C(\mu,\rho) \coloneqq \int_0^1 u^{(1-\mu)\left(\frac{1-\rho}{\rho}\right)-1}(1-u)^{\mu\left(\frac{1-\rho}{\rho}\right)-1}\,\mathrm{d}u$, if $\rho = 0$,

$$F(p;\mu,0) = \begin{cases} 0 & \text{if } p < \mu, \\ 1 & \text{if } p \ge \mu, \end{cases} \tag{1.5}$$

and, if $\rho = 1$,

$$F(p;\mu,1) = \begin{cases} 0 & \text{if } p < 0, \\ 1-\mu & \text{if } 0 \le p < 1, \\ 1 & \text{if } p \ge 1. \end{cases} \tag{1.6}$$

Assumption 1. *The probability $p_{0,j}$ that basal firm j experiences a disruption $X_{0,j} = 1$ is sampled independently, for all j, from a Beta-distributed random variable with mean μ_0 and*

correlation ρ_0*. That is,*

$$P_0 \sim \textit{Beta}(\mu_0, \rho_0) \textit{ and} \tag{1.7}$$

$$(X_{0,j} \mid P_0 = p_{0,j}) \overset{\text{i.i.d.}}{\sim} \textit{Bernoulli}(p_{0,j}). \tag{1.8}$$

One can think of this process in terms of a simple analogy. Each basal firm samples independently a coin, with bias probability of head $p_{0,j}$, from the same collection P_0 of coins. If the outcome of the coin is head, the firm is disrupted. Under this assumption, the disruption events $\{X_{0,j}\}_j$ experienced by basal firms are identically distributed but not independent, as their probabilities $p_{0,j}$ are sampled from the same collection P_0 of coins. The parameter μ_0 is the average fraction of disrupted firms in the basal layer, that is,

$$\mathbb{E}[X_{0,j}] = \mu_0. \tag{1.9}$$

For this reason, I hereafter refer to it as *basal risk*. The parameter ρ_0 is the degree of correlation in disruption events among basal producers, that is,

$$\text{Corr}[X_{0,j}, X_{0,i}] = \rho_0 \;\; \text{for all } i, j. \tag{1.10}$$

If $\rho_0 = 0$, the disruption events are independent. If $\rho_0 = 1$, the disruption events are perfectly correlated. Hence, hereafter I refer to it as *basal correlation*. Figure 1.2 illustrates the effect of the basal correlation ρ_0 on the basal disruption events. Consider an economy with a basal risk $\mu_0 = 0.3$. In case of a low basal correlation $\rho_0 = 0.02$ (lighter line), the disruption probability of basal firms is concentrated around μ_0. If ρ_0 increases to 0.1 (darker line), the distribution of disruption probabilities becomes more spread. While basal firms are on average as risky as before, the degree of correlation between their shocks is larger.

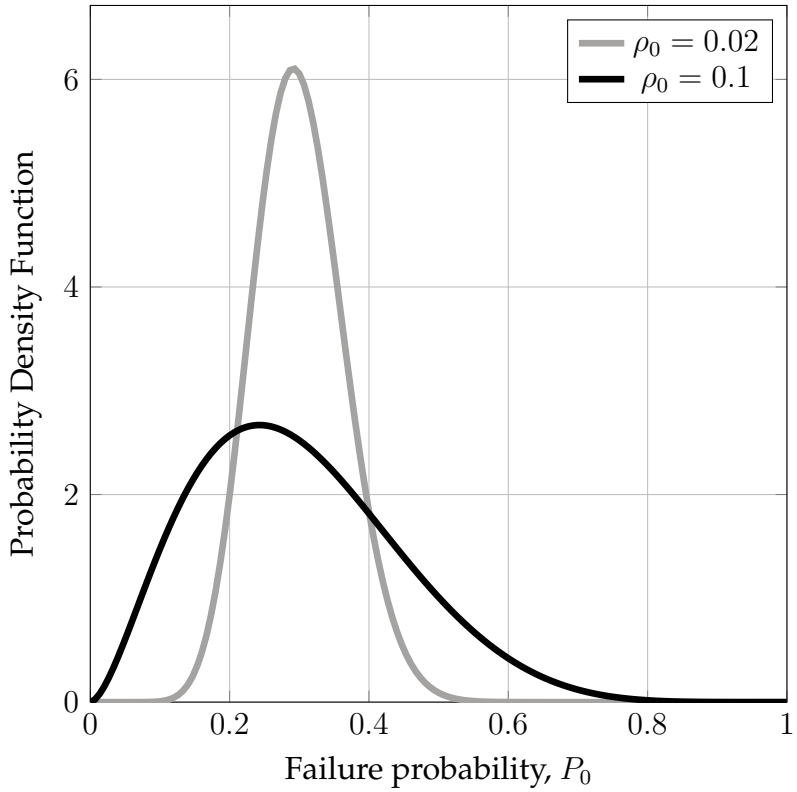

Figure 1.2: Distribution of disruption probabilities P_0 of basal firms for different levels of basal correlation ρ_0. In both cases $\mu_0 = 0.3$.

In the limit as $\rho_0 = 1$, all firms will either be disrupted, with probability μ_0, or not disrupted, with probability $1 - \mu_0$, that is, $X_{0,j} = X_{0,i}$ for all firms i and j.

Using this framework, the following assumption formalises the supply chain opacity.

Assumption 2. *Downstream firms (k, j) for $k \geq 1$ observe the distribution of the failure probability P_0, but do not observe the distribution of disruptions $\{X_{k,j}\}_j$.*

Under this assumption, all basal firms are ex-ante symmetric from the point of view of downstream producers. Going back to the example of polymer producers, downstream firms using plastic understand how hurricane risk can impact polymer production, which is an input to their production, as they observe P_0. Yet, they cannot estimate the risk that individual polymer producers $(0, j)$ face, since they do not observe $p_{0,j}$. The fact that basal producers are ex-ante symmetric is a crucial feature of the model: downstream producers must make sourcing decisions in a regime in which they cannot evaluate the risk of a disruption to specific suppliers. This notion of ex-ante symmetry, can be formalised using the concept of exchangeability.

Definition 2. *A sequence $\{X_i\}_i$ of random variables is exchangeable if its distribution is independent of the order of the sequence. That is, given any permutation σ,*

$$X_1, X_2, X_3 \ldots \sim X_{\sigma(1)}, X_{\sigma(2)}, X_{\sigma(3)} \ldots \tag{1.11}$$

Proposition 1. *The basal disruption events* $\{X_{0,j}\}_j$ *are exchangeable.*

The proof is provided in Appendix A.2.1. Due to the opacity of supply chains, a downstream producer i in layer $k = 1$ must choose its set of suppliers $\mathcal{S}_{1,i}$ without knowing the disruption risk of individual basal producers. However, by the exchangeability result, the firm's risk depends only on the number of suppliers $|\mathcal{S}_{1,i}|$ it selects, not on their specific identities. Therefore, I now impose an assumption on how firms select upstream producers when they are indifferent among suppliers and only the number of sources matters. Since downstream firms are indifferent between upstream suppliers, the selection procedure should be as "unbiased" as possible, meaning the upstream index (k, j) should have minimal influence on whether a firm is chosen as a supplier. A fully unbiased selection is not possible, as there is no uniform distribution over countably infinite upstream producers. Thus, I assume the selection procedure depends on an *uniformness* parameter $\theta > 0$. As $\theta \to 0$, firms are more likely to select suppliers with lower indices (k, j). As $\theta \to \infty$, the bias in supplier selection becomes arbitrarily small. The following formalises this idea.

Assumption 3. *If firm* $(k+1, i)$ *chooses to supply its input good* k *from* $s_{k+1,i} < \infty$ *sources and is indifferent among them, it picks suppliers* $\mathcal{S}_{k+1,i}$ *as follows.*

For an arbitrary uniformness $\theta > 0$*, a sequence of random variables*

$$u_j \overset{\text{i.i.d.}}{\sim} \textit{Beta}\left(\frac{1}{1+1/\theta}, \frac{1}{2+\theta}\right), \tag{1.12}$$

is sampled. Then, each upstream firm (k, i) *is assigned a weight*

$$\alpha_{k,1} = u_1 \textit{ and } \alpha_{k,j} = u_j \prod_{i<j}(1 - u_i). \tag{1.13}$$

These weights define a random probability mass function as

$$\lim_{n\to\infty} \mathbb{P}\left(1 - \sum_{j=1}^{n} \alpha_{k,1} < \delta\right) = 0, \tag{1.14}$$

for any $\delta > 0$*.*

Each downstream firm $(k+1, i)$ *constructs its set of sources* $\mathcal{S}_{k+1,i}$ *by sampling without replacement from* $\mathbb{N}$ *according to the random probability mass function* $\alpha_{k,i}$*.*

This construction, illustrated in Figure 1.3, is based on the stick-breaking process introduced in Ferguson (1973) and commonly used in constructing Dirichlet processes (Kallenberg, 2005). This assumption allows firm to "break ties" whenever multiple possible sourcing strategies yield the same payoff. To do so, it imposes a

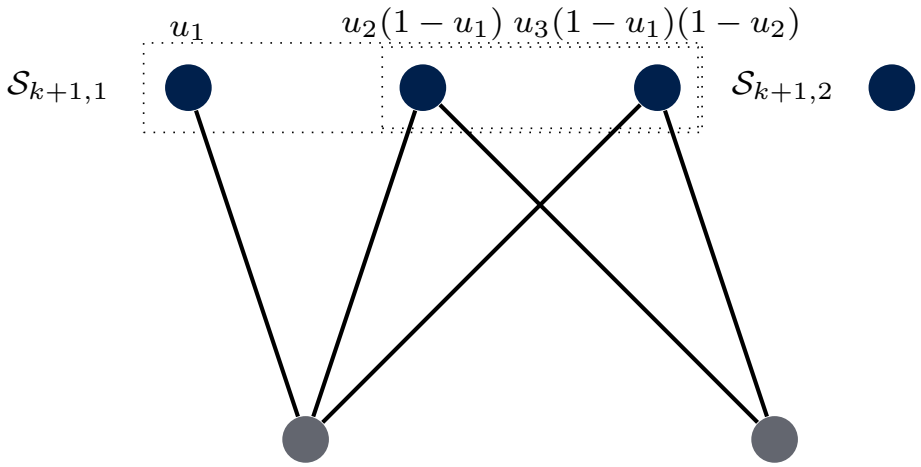

Figure 1.3: Illustration of the procedure to compute indifferent matching, Assumption 3

random preferential attachment weight u_j to each upstream producer (k, j). Crucially, this weight is random and does not depend on the downstream producer. That is, conditioning on the weights $\{u_j\}_j$, changing the index of a downstream firm $(k+1, i)$ does not change its sampling process. The parameter θ controls the degree to which the weight u_j is biased towards lower indices j. As $\theta \to 0$, more weight is placed on lower indices, hence all downstream firms are more likely to source from lower indices. As $\theta \to \infty$, more weight is placed on higher indices and firms are decreasingly likely to share suppliers.

Under Assumption 3, the exchangeability can be extended to all downstream layers of the supply chain.

Proposition 2. *In each downstream layer $k \geq 1$, the disruption events*

$$X_{k,1}, X_{k,2}, X_{k,3} \ldots$$

are exchangeable.

The proof is provided in Appendix A.2.2. Intuitively, in each layer, the sampling procedure is common for all downstream firms $(k+1, i)$, it is hence independent of the firm index i. This independence of the index i immediately implies that a permutation of the indices does not affect the downstream distribution of disruptions. Proposition 2 extends the exchangeability from the basal to all downstream layers. Given the exchangeability in each layer, from the point of view of a downstream firm, all suppliers are ex-ante symmetric. This implies that the profit of the firm depends exclusively on *how many* suppliers it chooses, rather than *which* suppliers it chooses. Then, a firm producing good k can first infer the distribution of the number $D_{k-1} := |\mathcal{D}_{k-1}|$ of disrupted firms among its potential suppliers and then choose the optimal number $s_{k,i} := |\mathcal{S}_{k-1,i}|$ of firms from which to source its input

good. Furthermore, by symmetry, all firms in layer k choose the same number s_k of sources, that is,

$$s_{k,i} = s_k \text{ for all } i. \tag{1.15}$$

As a result, the sourcing strategies $\mathcal{S}_{k,i}$ and $\mathcal{S}_{k,j}$ of any two firms i and j are such that their disruption probabilities $p_{k,i}$ and $p_{k,j}$ are identically distributed[4].

Before proceeding, it is worth discussing to what extent the result of Proposition 2 extends to a finite case, and, hence, this model represents a good approximation of a finite supply chain. Diaconis and Freedman (1980) have shown that, if the finite length n sequence of disruptions $X_{k,1}, X_{k,2}, \ldots, X_{k,n}$ is exchangeable, the subsequence $X_{k,1}, X_{k,2}, \ldots, X_{k,s_k}$ has variation distance at most $4s_k/n$ to the closest mixture of i.i.d. random variables. In each layer k of the production network this error accumulates, such that the total accumulated variation is at most

$$4\sum_{k=1}^{K} s_k/n. \tag{1.16}$$

Hence, the model represents a good approximation whenever the depth of the supply chain K and the number of suppliers chosen by firms s_k is small compared to the total number of firms n. Hence, as long as the supply chain is flat and sparse, relative to the total number of firms, the model yields a good finite approximation. In the real world there is ample evidence that production networks are sparse (Acemoglu et al., 2012). For example, Dhyne et al. (2020) found the fraction of sourcing connections to firms to be around $1/151$.

Using Proposition 2 allows us to characterise the firms' sourcing decisions, the formation of the production network, and the propagation of the disruptions in terms of a representative firm for each layer k.

1.2 Disruption Propagation

Using the framework introduced in the previous section, I now turn to characterising the propagation of the disruptions across layers. I consider the limit case in which $\theta \to \infty$, that is, vanishingly less weight is given towards lower indexed suppliers when selected by downstream firms using the procedure of Assumption 3. First, I will show that, in each layer k, the disruption events $\{X_{k,i}\}_i$ are i.i.d. Bernoulli trials conditional on a random variable $P_k = \mathbb{P}(X_{i,k} = 1)$. Recall, that the $X_{i,k}$ are constructed via a random probability mass function, hence the probability of a disruption P_k is itself a random variable. This can be thought of as the

[4]This approach is widely used in the study of random graphs, see, for example, Diaconis and Janson (2007); Kallenberg (2005).

random variable from which firms in layer k sample their disruption probabilities. It suffices to track how P_k evolves across the layers $k \in \{0, 1, 2 \dots K\}$ to study the propagation of the disruption. Furthermore, if $P_0 \sim \text{Beta}(\mu_0, \rho_0)$, as in Assumption 1, then $P_k \sim \text{BetaPower}(\mu_0, \rho_0, S_k)$ for all k, that is, the family of these distributions is closed as disruptions propagate through the layers. In the following, I introduce the BetaPower distribution and prove this result.

Definition 3. *A random variable P_k follows a BetaPower(μ, ρ, S_k) distribution, with risk $\mu \in [0, 1]$, correlation $\rho \in [0, 1]$, and power $S_k \in \mathbb{N}$, if it can be written as*

$$P_k = P_0^{S_k}, \tag{1.17}$$

where P_0 follows a Beta(μ, ρ). The probability density function is given by

$$f(p; \mu, \rho, S_k) \propto p^{\frac{(1-\mu_0)(1-\rho_0)}{S_k \rho_0} - 1} (1 - p^{1/S_k})^{\frac{\mu_0(1-\rho_0)}{\rho_0} - 1}. \tag{1.18}$$

Proposition 3. *If the upstream disruption probability $P_k \sim$ BetaPower(μ_0, ρ_0, S_k) and downstream firms select upstream suppliers using the procedure of Assumption 3 with $\theta \to \infty$, the downstream disruption probability is given by $P_{k+1} = P_k^{s_{k+1}}$, such that,*

$$P_{k+1} \sim \text{BetaPower}(\mu_0, \rho_0, s_{k+1} S_k). \tag{1.19}$$

The proof is provided in Appendix A.2.3. Proposition 3 guarantees that the distribution of disrupted firms remains in the same distribution family. This result allows us to describe disruptions propagation in the supply chain by mapping the evolution of the distribution P_k through the layers. Furthermore, the distribution P_k is fully determined by the risk μ_0 and correlation ρ_0 in the basal layer and the sourcing decisions $\{s_k\}_k$ the representative firms makes through the supply chain. Consider again the analogy of basal producers sampling biased coin from a bag with distribution P_0. From the point of view of a downstream firm in layer $k = 1$, before the coin is sampled from the bag by the basal producers, all basal producers are ex-ante symmetric. The downstream firm experiences a disruption if the outcome of the sampled coin is head for all of its suppliers. Hence, in choosing s_1 suppliers, the downstream firm is tying its disruption to s_1 coin tosses, drawn from the basal bag with distribution P_0. To illustrate how the sourcing decision of the representative downstream firm $k = 1$, influences the probability of a downstream disruption, consider

$$P_1 = P_0^{s_1}. \tag{1.20}$$

Figure 1.4 shows the probability density function of P_1 in case the firm has a single supplier $s_1 = 1$ (dotted lines) or two suppliers $s_1 = 2$ (solid line). The basal

disruption probability P_0 follows the same distribution as in Figure 1.2. If downstream producers supply from a single source (dashed line), the disruption probability downstream follows the same distribution as that of upstream producers, $P_1 = P_0$. To diversify risk, firms can choose to contract an additional supplier (solid line). Regardless of the basal correlation ρ_0, constructing an additional supplier reduces the downstream risk $\mu_1 = \mathbb{E}[P_1]$. If the upstream correlation is low, $\rho_0 = 0.02$ the distribution of P_1 remains concentrated around the mean. If the upstream correlation is large, $\rho_0 = 0.1$, the suppliers' disruption events become more correlated and the downstream disruption probabilities become fat-tailed, that is, a significant fraction of firms is likely to be disrupted and, as a consequence, diversification is less effective. Hence, a downstream firm can mitigate risk by contracting an additional supplier, but a large disruption probability remains.

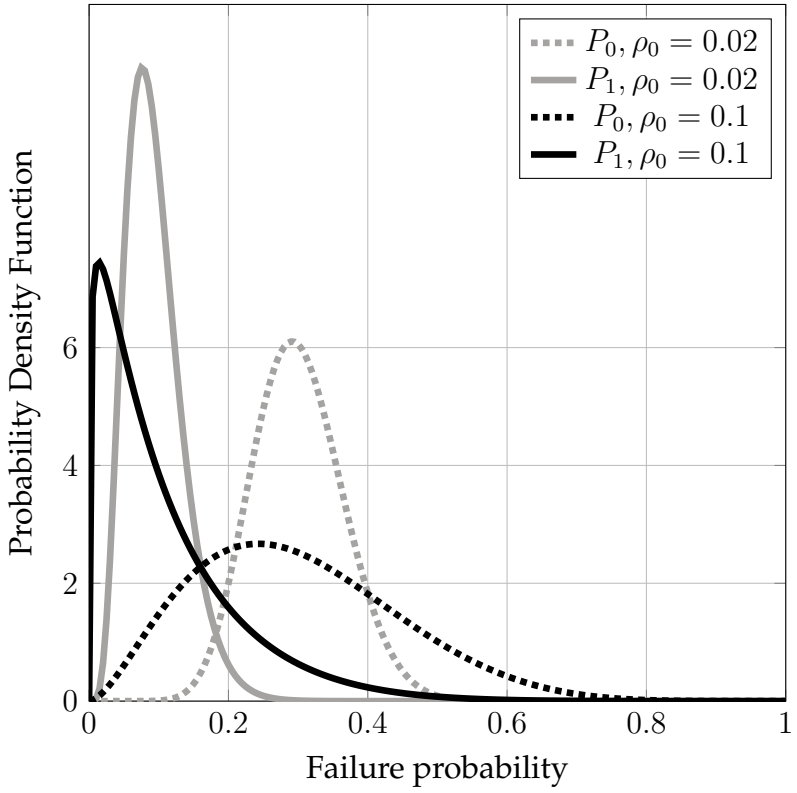

Figure 1.4: Distribution of disruption probabilities of downstream firms for different levels of upstream correlation ρ_0, in the cases of single sourcing (dotted) and multisourcing (solid). In both cases $\mu_k = 0.3$.

Having established the link between upstream and downstream probabilities of disruptions

$$P_{k+1} = P_k^s, \tag{1.21}$$

I now turn to the analysis of how risk $\mu_k = \mathbb{E}[P_k]$ propagates through the supply chain, before studying how firms make decisions endogenously.

Proposition 4. *The average disruption probability between one layer k and the next $k+1$ depends on the sourcing strategy s_{k+1} via the map*

$$\mu_{k+1} = \begin{cases} \eta(s_{k+1}, S_k)\ \mu_k & \text{if } s_{k+1} > 0, \\ 1 & \text{if } s_k = 0, \end{cases} \tag{1.22}$$

where $S_k := \prod_{j=1}^{k} s_j$ is the diversification level up to layer k and η is the "risk reduction factor", which is given by

$$\eta(s_{k+1}, S_k) := \left(\frac{\mu_0 \frac{\rho_0}{1-\rho_0} + S_k}{\frac{\rho_0}{1-\rho_0} + S_k} \right) \left(\frac{\mu_0 \frac{\rho_0}{1-\rho_0} + S_k + 1}{\frac{\rho_0}{1-\rho_0} + S_k + 1} \right) \times \cdots \times \left(\frac{\mu_0 \frac{\rho_0}{1-\rho_0} + S_k s_{k+1} - 1}{\frac{\rho_0}{1-\rho_0} + S_k s_{k+1} - 1} \right). \tag{1.23}$$

The risk reduction factor $\eta(s_{k+1}, S_k)$, derived in Appendix A.2.4, governs how the firm's choice s_{k+1}, the choices along the firm's production chain S_k, and the basal conditions μ_0, ρ_0 affect the expected number of disruptions downstream. This interplay is illustrated in the following figures. Figure 1.5 shows how the risk reduction factor varies with basal correlation ρ_0 for different sourcing strategies s_k, fixing the upstream diversification of $S_k = 2$. If correlation ρ_0 in the basal layer grows, to obtain a given level of risk reduction η, firms producing good k need to source more suppliers. If $\rho_0 = 1$, diversification becomes impossible, as $\eta \to 1$ and $\mu_{k+1} \to \mu_k$ for any sourcing strategy s_{k+1}.

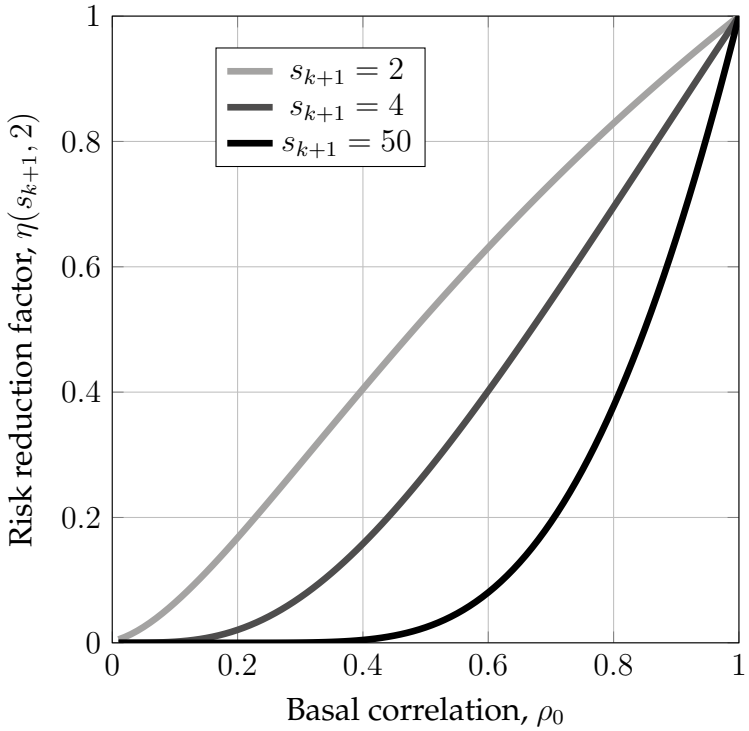

Figure 1.5: Risk reduction factor μ_{k+1}/μ_k at different basal correlation levels ρ_0 and for different sourcing strategies s_{k+1}. $S_k = 2$.

As the above, Figure 1.6 shows the response of the risk reduction factors to different levels of basal correlation, but instead of varying the strategy s_k of the firm, it varies the level of upstream diversification S_k. For low levels of basal correlation ρ_0, more upstream diversification S_k allows downstream producers to achieve lower risk with fewer suppliers. Yet, there is a level of basal correlation after which more diversification is detrimental for the downstream firm, as this high upstream diversification simply exacerbates tail risk. This represents a crucial externality the upstream suppliers impose on downstream producers. For a low level of correlation, sourcing downstream represents a positive externality downstream. This externality shrinks as correlation increases until it becomes a negative externality. Section 1.4 explores the welfare consequences of this mechanism.

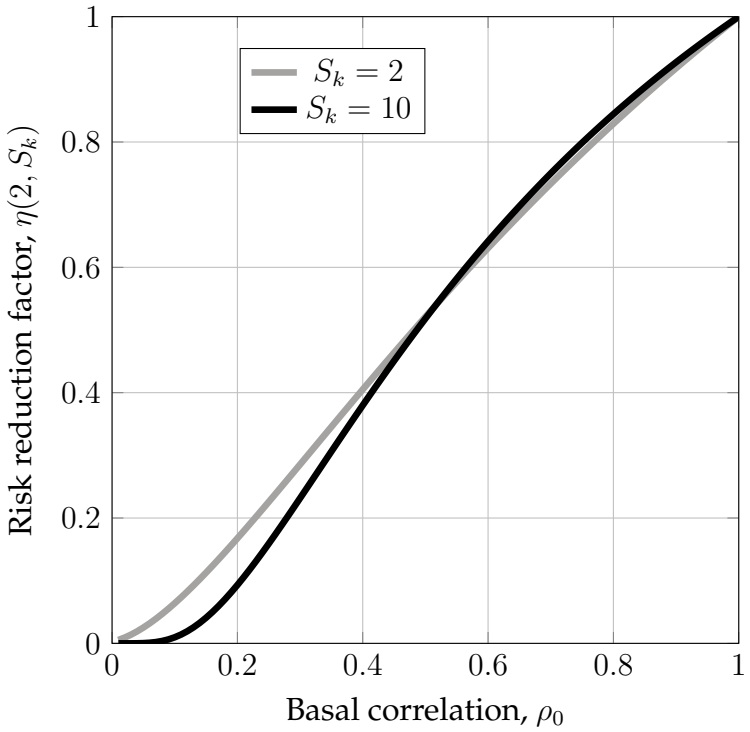

Figure 1.6: Risk reduction factor μ_{k+1}/μ_k at different basal correlation levels ρ_0 and for different upstream diversification levels strategies S_k. $s_{k+1} = 2$.

1.3 Firm Optimal Diversification and Equilibrium

The mechanics of disruption propagation, derived in the previous section, determine the firm's optimal sourcing strategy. This section derives such optimal strategies. Importantly, due to Proposition 2, all firms in a given layer are identical before the shock and so is their optimisation problem. We can hence focus on the problem of the representative firm in layer $k+1$: to choose how many suppliers in layer k

to source from, based on the inferred distribution of their probability of experiencing a disruption event. This, in turn, is fully determined by the average disruption probability μ_0, the correlation ρ_0 of the disruptions, and the sourcing strategies $\{s_1, s_2, \dots s_k\}$ of the representative firms upstream. Henceforth, I assume firms face quadratic costs of sourcing, with cost parameter c, such that expected profits (1.2) can be written as

$$\Pi_k(s) = \left(1 - \mathbb{E}_s\left[P_k\right]\right)\pi - \frac{c}{2}s^2, \tag{1.24}$$

The optimisation problem of the firm is to then choose the optimal sourcing strategy

$$s_k = \arg \max_{s \in \{0,1,2,\dots\}} \Pi_k(s). \tag{1.25}$$

1.3.1 Limit Case: Uncorrelated Disruptions

Before turning towards the general framework, I first analyse a limit case in which suppliers' risk is not correlated, that is, $\rho_0 = 0$. This limit case gives a useful interpretation of the incentives behind multisourcing and allows us to establish a benchmark against which to study the introduction of correlated shocks.

Proposition 5. *If risk among basal firms is uncorrelated, that is, $\rho_0 = 0$, disruption events in layer k are independent and happen with probability*

$$\mu_{k+1} = \mu_k^{s_{k+1}}. \tag{1.26}$$

Proof. By the definition of P_k, $P_k = \mu_k$ if $\rho_0 = 0$. □

In this limit case, the representative firms' profits (1.24) are given by

$$\Pi_{k+1}(s) = \left(1 - \mu_k^s\right)\pi - \frac{c}{2}s^2. \tag{1.27}$$

The function Π_{k+1} is illustrated in Figure 1.7. Using this, we can derive the optimal sourcing strategy s_k of a firm producing good k. A firm with s suppliers contracts an extra one only if doing so yields a positive marginal profit

$$\begin{aligned} \Delta\Pi_{k+1}(s) &\coloneqq \Pi_{k+1}(s+1) - \Pi_{k+1}(s) \\ &= \mu_k^s(1-\mu_k)\pi - c\left(s + \frac{1}{2}\right). \end{aligned} \tag{1.28}$$

The marginal benefit $\Delta\Pi_{k+1} : \mathbb{R} \to \mathbb{R}$ is strictly decreasing and that it has a unique root. Hence, the optimal number of suppliers s_{k+1} is the smallest integer s for which the expected marginal profit is negative, that is, $\Delta\Pi_{k+1}(s) < 0$.

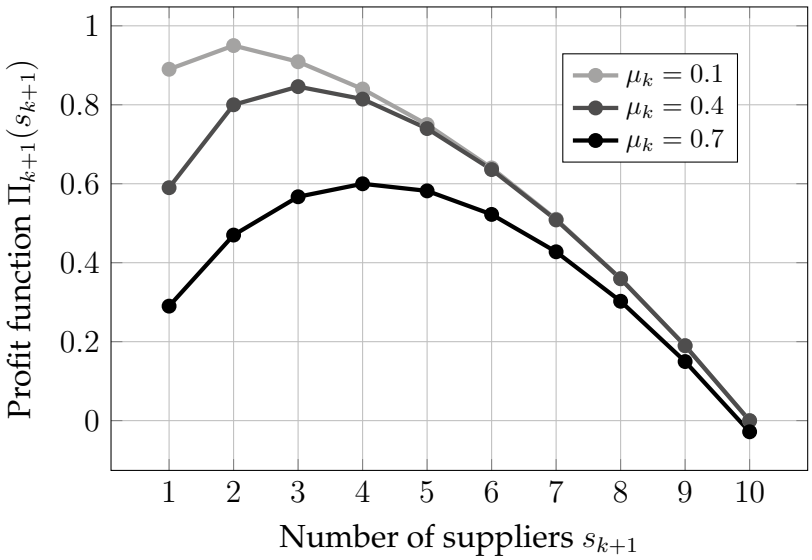

Figure 1.7: Firms profits Π_{k+1} as a function of the sourcing strategy s_{k+1} for different levels of upstream risk μ_k.

Definition 4. *Let $\tilde{s}_{k+1} \in \mathbb{R}$ be the unique root of $\Delta\Pi_{k+1}$. I refer to this quantity as the "desired sourcing strategy" of the firm.*

The optimal sourcing strategy s_{k+1} is then the smallest integer larger than the firm's desired sourcing strategy $\tilde{s}_{k+1}$, that is,

$$s_{k+1}(\mu_k) = \begin{cases} \lceil \tilde{s}_{k+1} \rceil & \text{if } \tilde{s}_{k+1} > 0 \text{ and} \\ 0 & \text{otherwise,} \end{cases} \tag{1.29}$$

where $\lceil \cdot \rceil$ is the ceiling function. The optimal sourcing strategy s_{k+1}, immediately implies an upper bound on the suppliers' risk μ_k firms are willing to tolerate.

Proposition 6. *Introduce the threshold*

$$\mu^{\text{frag}} \coloneqq 1 - rc, \tag{1.30}$$

where $rc \coloneqq \frac{c/2}{\pi}$ is the real marginal costs of an additional supplier. If the average disruption probability μ_k is larger than μ^{frag}, the downstream firm does not source any inputs, that is, $s_{k+1} = 0$.

Proof. Suppose a firm optimally does not source any inputs. This implies that the marginal benefit of adding the first supplier is negative, that is, $\Delta\Pi_{k+1}(0) < 0$, which yields the desired inequality. □

For any value of initial basal risk $\mu_0 > \mu^{\text{frag}}$, firms do not source and the production network shuts down.

As expected, the desired $\tilde{s}_{k+1}$ and the optimal sourcing strategy s_{k+1} are determined by the upstream average disruption probability μ_k and the real marginal costs of contracting a new supplier rc. Figure 1.8 illustrates the effect these two conditions have on the optimal sourcing strategy. First, higher real marginal costs rc reduce the firm's number of sources. Second, as the upstream average disruption probability μ_k increases, initially, the firm seeks higher diversification, until a level above which the desired sourcing strategy starts decreasing steeply.

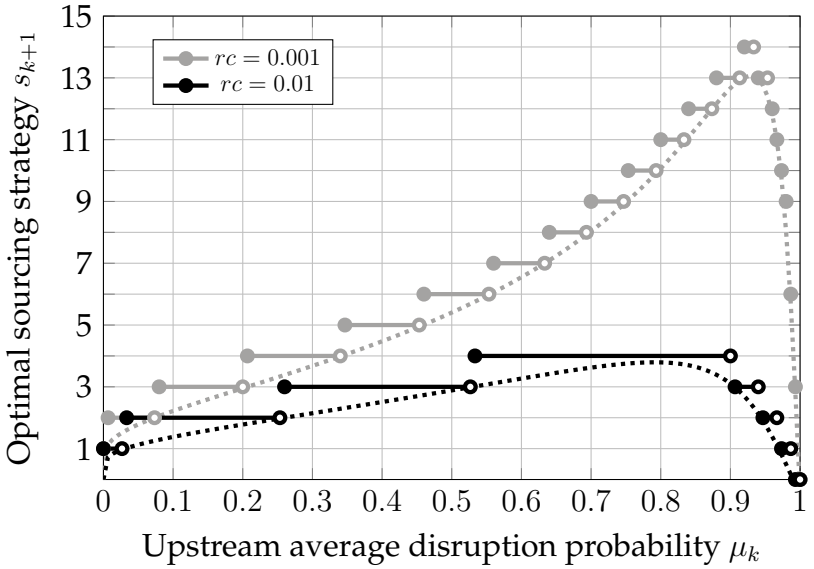

Figure 1.8: The desired $\tilde{s}_{k+1}$ (dotted) and optimal s_{k+1} sourcing strategy (solid) as a function of the upstream average disruption probability μ_k

Having studied how risk affects the firm's optimal sourcing, I now look at the opposite channel, that is, how the firm sourcing strategy affects risk propagation. To do so, I think of the average disruption probability

$$\mu_{k+1} = \phi(\mu_k) \coloneqq \mu_k^{s_{k+1}(\mu_k)} \tag{1.31}$$

from suppliers to downstream producers as a dynamical system, not in time but in layers $k \in \{0, 1, 2 \ldots\}$ of the supply chain. Given a *basal condition* μ_0, a fixed point $\bar{\mu}$ of the map (1.31) is then a level of disruption probability $\bar{\mu}$ such that all firms downstream of a layer $\overline{k}$ single-source, namely $s_l = 1$ for all $l \geq \overline{k}$, and hence all share the same disruption probability $\mu_l \equiv \bar{\mu}$. When looking at the production network through this lens, a natural question arises: which basal levels of disruption probabilities μ_0 are not endogenously diversified by the production network, that is, $\bar{\mu} \geq \mu_0$? To answer this, first I characterise the downstream disruption probability $\bar{\mu}$.

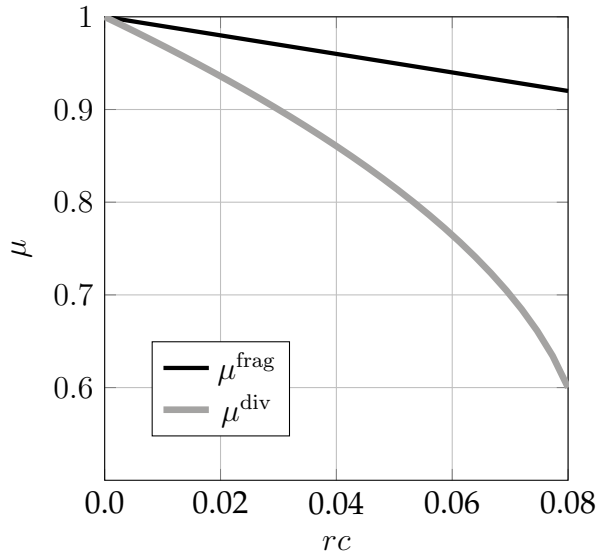

Figure 1.9: The critical thresholds μ^{div}, above which are unable to diversify, and μ^{frag} above which firms stop producing, as a function of the relative marginal costs rc.

Proposition 7. *The downstream disruption probability $\bar{\mu}$ satisfies*

$$\bar{\mu}\,(1-\bar{\mu}) \leq 3rc. \tag{1.32}$$

Proof. A steady state is attained at level $\bar{k}$ if $s_{\bar{k}} = 1$. This implies that the marginal benefit of multisourcing is negative $\Delta\Pi_{\bar{k}}(1) \leq 0$. This yields the desired inequality. □

Corollary 7.1. *Introduce the critical threshold*

$$\mu^{\text{div}} := \frac{1}{2} + \sqrt{\frac{1}{4} - 3rc}. \tag{1.33}$$

If $\mu_0 > \mu^{\text{div}}$, the endogenous supply chain is unable to diversify risk, that is, $\bar{\mu} \geq \mu_0$.

This result, proven in Appendix A.2.5, links the firm real marginal costs of sourcing rc and the production network risk. As relative marginal costs increase, the capacity of the production network to endogenously diversify decreases and firms' under-diversification yields endogenous fragility. Notice that, comparing the threshold μ^{div} of endogenous diversification with the threshold μ^{frag} of firm shutdown, illustrated in Figure 1.9, for some levels of basal probability of disruption μ_0, despite no firm shutting down production $\mu_0 < \mu^{\text{frag}}$, the production network as a whole is still unable to endogenously diversify risk $\mu_0 > \mu^{\text{div}}$. This is true even in this special case, where the firms' risk is uncorrelated. In the next section, I introduce correlation risk $\rho > 0$ and investigate how doing so changes the dynamics illustrated here.

1.3.2 Optimal Sourcing with Correlated Distributions

If disruption events are not independent, that is, $\rho_0 > 0$, the risk among suppliers throughout the production network is correlated, which affects the firm's optimisation incentives. In this case, the problem of a firm in layer $k+1$ is still to choose the number of suppliers $s_{k+1} \in \{0, 1, 2 \ldots\}$ that maximises the profits Π, but, by Proposition 4, the firm's disruption probability is given by the average disruption probability of its suppliers μ_k multiplied by a factor $\eta(s_{k+1}, S_k)$ which depends on the upstream diversification $S_k = \prod_{l<k} s_l$. As in the limit case analysed in the previous section, the firm will increase diversification as long as the expected increase in profits obtained by adding a supplier outweighs the costs of contracting that additional supplier. These expected marginal profits are given by

$$\Delta\Pi_{k+1}(s_{k+1}) = \Big(\eta(s_{k+1}, S_k) - \eta(s_{k+1}+1, S_k)\Big)\, \mu_k \pi - c\left(s_{k+1} + \frac{1}{2}\right). \tag{1.34}$$

The characterisation of the optimal sourcing strategy is analogous to the limit case without the correlation discussed above. $\tilde{s}_{k+1} \in \mathbb{R}$ is the desired sourcing strategy for which the marginal benefits and marginal costs of diversification are equal, such that $\Delta\Pi_{k+1}(\tilde{s}_{k+1}) = 0$. As the marginal profits are strictly decreasing in the number of suppliers (see Appendix A.2.6), the firm will, as in the limit case, choose its optimal sourcing strategy as $s_{k+1} = \lceil \tilde{s}_{k+1} \rceil$ if $\tilde{s}_{k+1} > 0$ and chooses not to source any inputs otherwise. Figure 1.10 illustrates how the optimal sourcing strategy s_{k+1} changes with upstream correlation ρ_k for different levels of relative costs rc. As upstream correlation increases, the firm increases its sources to diversify risk. Yet, for large levels of correlation, the disruption of an additional source of the input good is likely correlated to a disruption among the firm's existing suppliers, which reduces the firm's incentive to multisource. As disruptions among suppliers tends towards perfect correlation, $\rho_k \to 1$, the firm sources from a single supplier.

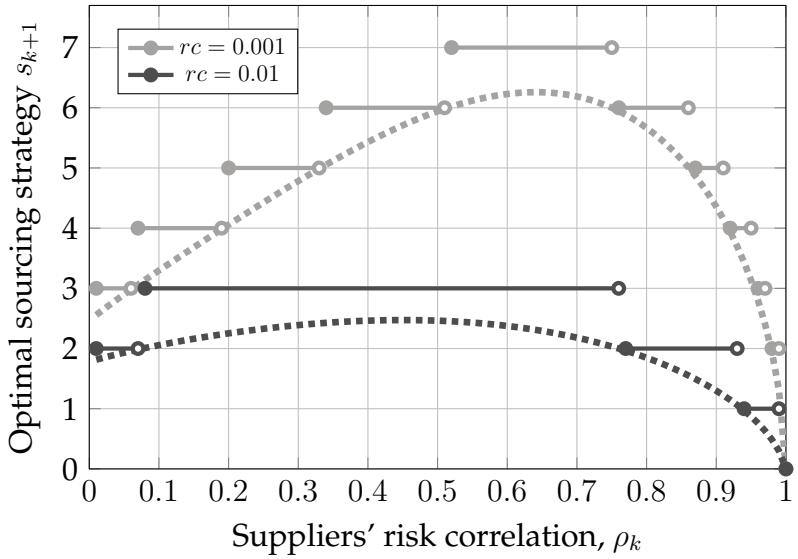

Figure 1.10: The desired $\tilde{s}_{k+1}$ (dotted) and optimal s_{k+1} sourcing strategy (solid) as a function of the upstream correlation ρ_k

To study the ramifications of this endogenous channel for the supply chain formation and its fragility, next I analyse the propagation of risk through the layers. As above, conditional on the choices of the upstream producers $S_k = \prod_{l<k} s_l(\mu_l)$, we can view the mapping of the probability of disruption between layers as a dynamical system through the layers of the production network, with map

$$\mu_{k+1} = \Phi(\mu_k; S_k) \coloneqq \eta\left(s_{k+1}(\mu_k), S_k\right)\mu_k. \tag{1.35}$$

A steady state of the dynamical system is then an average disruption probability $\bar{\mu} \coloneqq \mu_{\bar{k}}$ in some layer $\bar{k}$ such that all downstream layers $l \geq \bar{k}$ have the same average disruption probability $\mu_l \equiv \bar{\mu}$. This can occur in two cases. Either the firm in layer $\bar{k}$ does not source, that is, $s_{\bar{k}} = 0$, or it sources from one supplier, that is, $s_{\bar{k}} = 1$. The former case is trivial: the production network shuts down, and all downstream firms do not produce, such that $\bar{\mu} = 1$. In the latter case, by single sourcing, the average disruption probability in layer $\bar{k}$ is the average disruption probability among the suppliers $\bar{k} - 1$, as the risk reduction factor $\eta(s_k, S_{k-1}) = 1$ if $s_k = 1$. Because the layers are symmetric, the firms in the downstream layer $\bar{k} + 1$ face the same problem as those in layer $\bar{k}$, such that they endogenously single source, that is, $s_{\bar{k}+1} = 1$. Inductively, this holds true for all $l \geq \bar{k}$, hence $\mu_l \equiv \bar{\mu}$. Hereafter, I refer to the situation in which the downstream average disruption probability is greater than the basal one, that is, $\bar{\mu} \geq \mu_0$, as *endogenous fragility*. Figure 1.11 shows the downstream average disruption probability $\bar{\mu}$ as a function of basal average disruption probability μ_0, for cases in which basal correlation ρ_0 is low or high. In both cases for large possible initial levels of basal average disruption probability μ_0 the supply chain is

endogenously resilient, as $\bar{\mu} < \mu_0$. But, as in the uncorrelated cases studied above, there is a threshold of average basal disruption probability $\mu_0 > \mu^{\text{div}}$ for which the firm is endogenously fragile and $\bar{\mu} \geq \mu_0$. The threshold effect is discontinuous. At $\mu_0 \equiv \mu^{\text{div}}$ an arbitrarily small increase in μ_0 can lead to discontinuously large downstream failure probabilities $\bar{\mu}$.

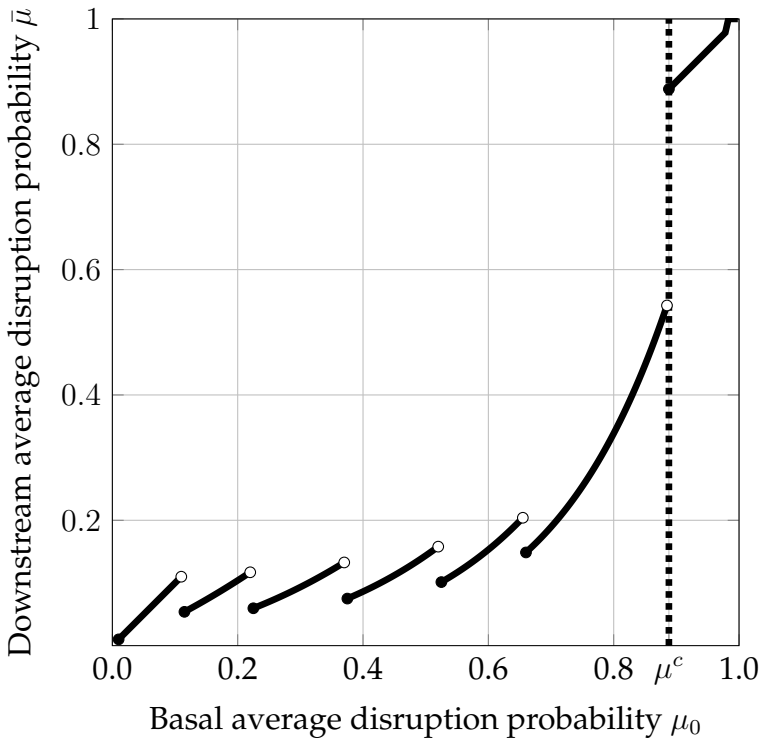

Figure 1.11: Downstream average disruption probability as a function of basal average disruption probability for $\rho_0 = 0.4$

The threshold μ^{div} is decreasing in the basal level of correlation ρ_0, as illustrated in Figure 1.12. This implies that a small increase in basal correlation leads to discontinuous increases in the downstream average disruption probability. This result highlights an additional channel to that studied by Elliott, Golub and Leduc (2022) by which supply chains can be endogenously fragile: even if the expected failure probability μ_0 of basal producers remains unchanged, an increase in the correlation of their disruptions ρ_0, can endogenously induce large fragilities. Hence, phenomena that can lead to increases in risk among upstream producers, such as offshoring, climate disruptions, and economic agglomeration, can generate underdiversification and endogenous fragility, even as they leave individual producer's risk unchanged.

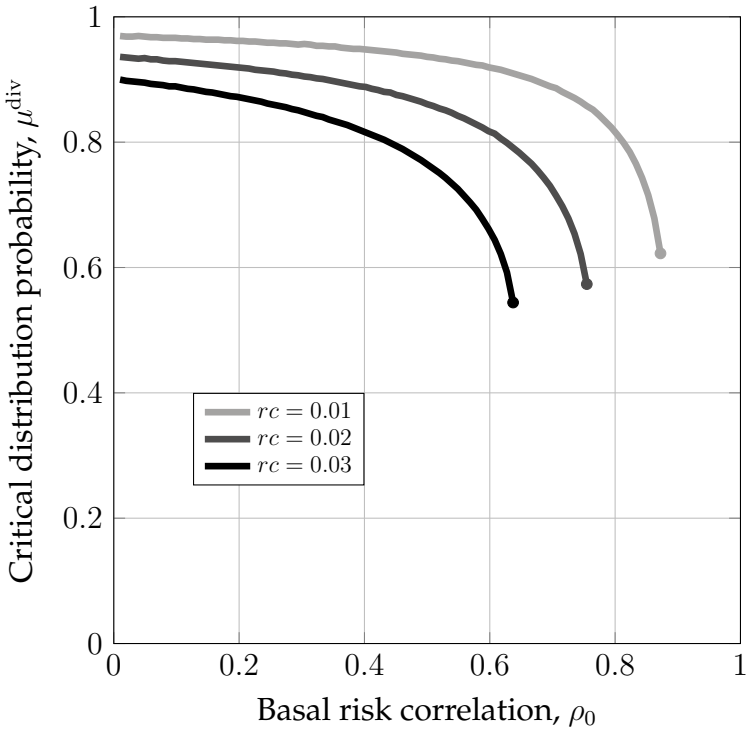

Figure 1.12: Critical level of basal average disruption probability μ^{div} as a function of the basal correlation.

1.4 Social Planner Problem

To establish a benchmark to which one can compare the competitive equilibrium analysed above, in this section I solve the model from the perspective of a social planner. The social planner attempts to, on the one hand, minimise the number of firms expected to fail, and, on the other, minimise the number of costly sourcing relations. To develop a useful benchmark, I define a social planner problem that can be meaningfully compared to the decentralised firm's problem by making the following two assumptions.

Assumption 4. *The social planner knows the distribution of failure in the basal layer $P_0 \sim Beta(\mu_0, \rho_0)$ and establishes the firms' sourcing decisions before P_0 is realised.*

Assumption 5. *The social planner makes sourcing decisions such that no two firms share suppliers, that is, $\mathcal{S}_{k,i} \cap \mathcal{S}_{k,j}$ is empty.*

By this assumption, the social planner can "diversify away" all the correlation that arises due to the network structure. Hence, the only source of risk in the model is represented by the shutdowns experienced by firms in the basal layer, which happen with non-idiosyncratic probabilities $\{p_{0,j}\}_j$ sampled from P_0 (Assumption 4). This assumption can be satisfied by the planner as long as the socially optimal

sourcing strategies $\mathcal{S}_{k,i}$ are finite, as they can match downstream firms in k and upstream firms in $k-1$ by picking a bijection $\gamma_k : \mathbb{N} \to \mathbb{N}^{\max_k |\mathcal{S}_{k,i}|}$.

Combining Assumptions 4 and 5, the social planner problem is then to maximise average expected payoff

$$W(\{\mathcal{S}_{k,i}\}) := \frac{1}{K}\sum_{k=0}^{K} \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \Big(1 - \mathbb{P}\big(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1}\big)\Big)\,\pi - \frac{c}{2}|\mathcal{S}_{k,i}|^2, \tag{1.36}$$

by choosing a sourcing strategy $S_{k,i} \subseteq \mathbb{N}$ for each firm in each layer such that $S_{k,i} \cap S_{k,j}$ is empty for all i, j. The social planner problem can be further simplified by noticing that, given that all firms in layer k, are identical if establishing an additional path from a firm in layer k to a basal firm has positive marginal benefits, then it has positive marginal benefits for all firms in layer k which share the same number of paths to basal firms. Hence, as in the decentralised firms' problem, the social planner can choose the optimal number of sources in each layer and then match downstream producers and upstream suppliers using γ_k. This allows to formulate the social planner problem recursively. Let V_k be the maximal average welfare from layer k to the last layer K. This can be recursively defined as

$$V_k(P_{k-1}) = \max_{s_k} \left\{ \left(1 - \mathbb{E}\big[P_{k-1}^{s_k}\big]\right)\pi - \frac{c}{2}s_k^2 + \mathbb{E}\big[V_{k+1}(P_k)\big] \right\}, \tag{1.37}$$

where the state $P_{k-1} \sim \text{BetaPower}(\mu_0, \rho_0, s_1 s_2 \dots s_{k-1})$ is the fraction of disrupted firms, which evolves as

$$P_k = P_{k-1}^{s_k}. \tag{1.38}$$

The average welfare in layer $K+1$ is given by $V_{K+1}(P_K) = 0$, since firms in the last layer are never sources to other firms. The state is initialised at $P_0 \sim \text{Beta}(\mu_0, \rho_0)$. This problem can be solved using standard backward induction techniques (see Appendix A.3). The optimum average social welfare (1.36) can then be written as

$$V_1(P_0) = \max_{s_1, s_2, \dots s_{K-1}} W(\{s_1, s_2, \dots s_{K-1}\}). \tag{1.39}$$

Let $\{s_k^p\}_{k=1}^K$ be the socially optimal sourcing strategies sequence and $\{\mu_k^p\}_{k=1}^K$ be the expected disruption in each layer given by such sourcing strategies. I assume that the number of layers K is sufficiently large such that for some downstream layer $\bar{k}$, the risk remains constant $\mu_k^p \equiv \overline{\mu}^p$ for all $k \geq \bar{k}$. Figure 1.13 shows this difference in downstream risk between the decentralised case and the social planner, that is, $\overline{\mu} - \overline{\mu}^p$ for the same two cost regimes. If pairing costs are low, the social planner achieves marginally lower risk levels of downstream risk for most initial conditions. If initial basal correlation ρ_0 is sufficiently large and the average basal disruption

probability μ_0 is sufficiently low, the firms over-diversify compared to the socially optimum $\overline{\mu} < \overline{\mu}^p$. If relative pairing costs are high, the social planner is able to diversify risk around the critical threshold μ_c, such that the decentralised equilibrium induces inefficiently high levels of average downstream disruption probability, that is, $\overline{\mu} > \overline{\mu}^p$. This result implies that the cascading failures that occur around the critical threshold are fully attributable to firms' endogenous under-diversification motives and are hence inefficient.

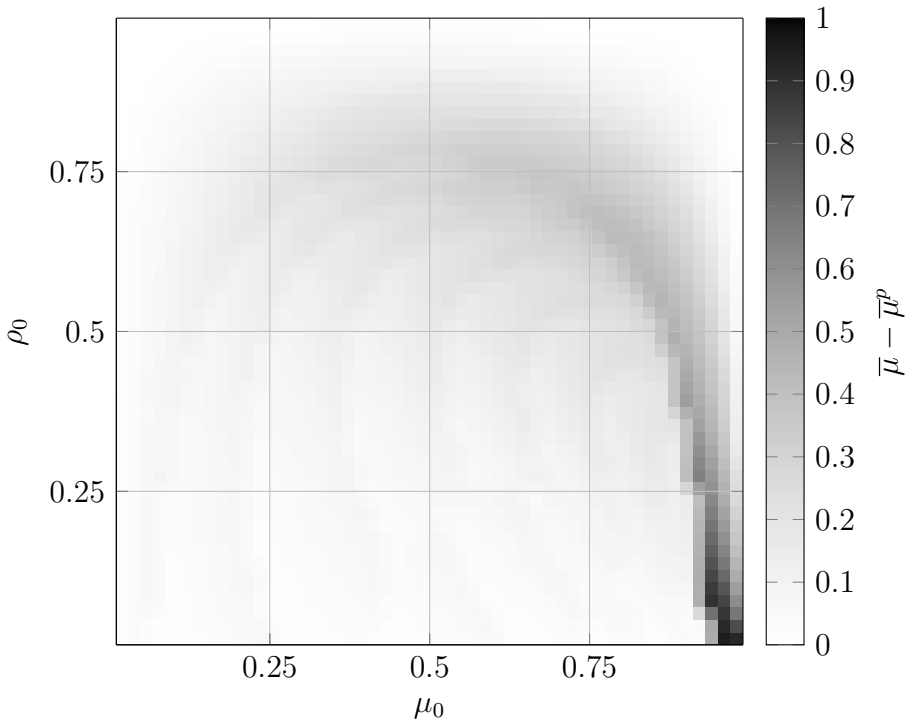

Figure 1.13: Change in downstream expected failure probability between the firms' $\overline{\mu}$ and the social planner $\overline{\mu}^p$ equilibrium, given different initial conditions μ_0 and ρ_0 in a low (left) and a high (right) relative pairing costs regime.

The differences between the firms' sourcing strategies and the social optimum generate welfare losses in the production network. Letting W be the average firm profit in the decentralised case and W^p be the average profit achieved by the social planner, Figure 1.14 illustrates the welfare loss due to the firms' diversification decisions $W - W^p$. The welfare loss is larger around the critical value μ^c, where the production network is endogenously fragile. At these levels of risk, firms' upstream firms' diversification incentives are weak, which creates large downstream resilience externalities. Crucially, both an increase in basal risk μ_0 and an increase in basal correlation ρ_0 can generate discontinuous welfare losses.

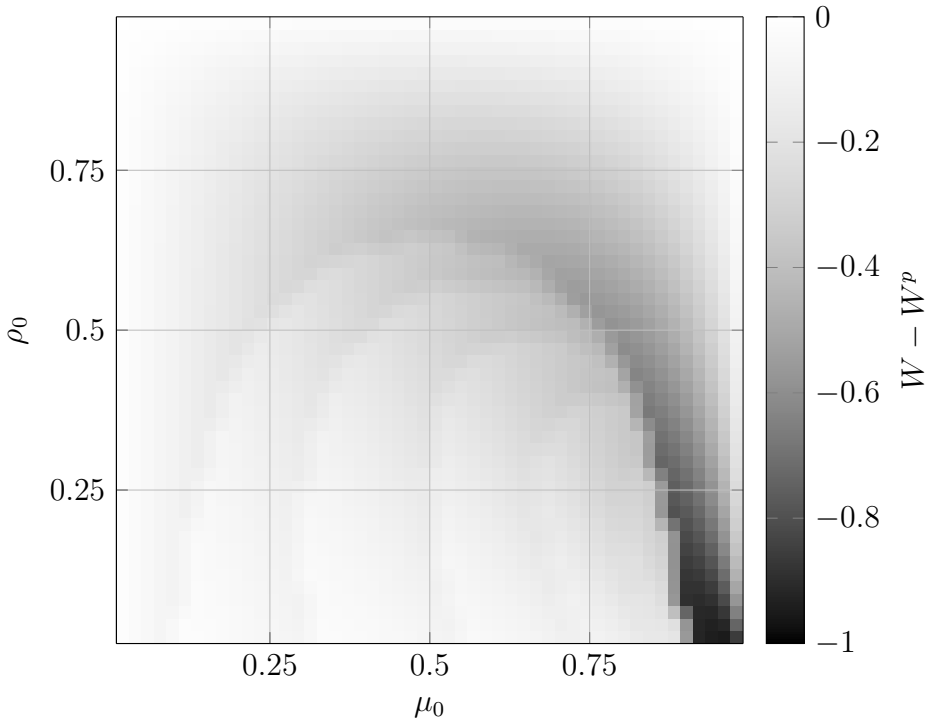

Figure 1.14: Welfare loss of the decentralised equilibrium $W - W^p$, given different initial conditions μ_0 and ρ_0.

1.5 The Role of Opacity

So far I assumed that firms cannot observe the realisation of the supply chain and the basal disruption probabilities P_0 when making sourcing decisions. To understand how this assumption affects optimal decisions and fragility within the supply chain, I now analyse the supply chain under perfect information. The following assumption clarifies what is meant by perfect information in the context of the model.

Assumption 6. *In a regime of perfect information, each firm i in level k is able to perfectly estimate the disruption probability of each potential supplier and the full correlation structure of the disruption events.*

Under this perfect information regime, the firm can assign correct probabilities to its own disruption risk

$$\mathbb{P}\Big(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1}\Big) \text{ for all possible sourcing strategies } \mathcal{S}_{k,i}.$$

The firm can hence rank suppliers by the marginal reduction in risk they provide and source from the "safest" s_k desired suppliers. As all firms downstream are ex-ante symmetric, the marginal benefits of diversification experienced by firm (i, j) are the same as those of all other firms in layer k, which implies that, in equilibrium,

all firms in layer k will employ the same sourcing strategy, given that they are ex-ante symmetric. This outcome is beneficial for any single firm, but detrimental to the stability of the production network. The following two propositions formalise this.

Proposition 8. *Compared to the opaque scenario, for the same number of sources, firm (k, i) is (weakly) less likely to be disrupted.*

Proof. Given the same number of sources, the firm with perfect information minimises its disruption risk with fewer constraints than in the opaque scenario. □

Proposition 9. *Under perfect information, the supply chain is maximally fragile: either all firms fail or none do.*

Proof. Without loss of generality, assume basal firms are sorted by their disruption risk $p_{0,1} > p_{0,2} > \ldots$. Suppose a firm in layer $k = 1$ chooses to contract $s_{1,i}$ suppliers. In choosing which suppliers to source from, the optimal choice is then to pick the first $s_{1,i}$-th basal producers, namely

$$\mathcal{S}_{1,i} = \{1, 2, \ldots s_{1,i}\}. \tag{1.40}$$

By symmetry, this is also the optimal choice of all other firms in layer $k = 1$, such that $\mathcal{S}_{1,i} = \mathcal{S}_{1,j}$ for all i, j. This further implies that disruption events in layer $k = 1$ are perfectly correlated, $X_{1,i} = X_{1,j}$ for all i, j, as each two firms share all suppliers. This in turn implies that, regardless of the diversification strategy, all firms downstream $k > 1$ experience perfectly correlated distribution events. Hence,

$$X_{k,i} = X_{k,j} \text{ for all firms } i, j \text{ and layers } k > 0. \tag{1.41}$$

□

Supply chain opacity, despite preventing firms from implementing an optimal diversification strategy, leads to more resilient supply chains. Hence, policy efforts to improve supply chain resilience via transparency might backfire if not paired with efforts to coordinate diversification of firms' sources.

1.6 Conclusion

Deeper and more globalised supply chains are more vulnerable to widespread, correlated disruptions. This chapter studies how firms diversify sourcing risk when disruptions are correlated, and how this affects the endogenous formation of supply chains. I show that, as disruption correlation rises, a firm sources its input good

from more suppliers to diversify the risk of a disruption. Yet, there is a level of correlation after which the expected risk reduction of adding a supplier is small, hence, the firm starts reducing its number of suppliers and lowers its diversification. This mechanism can create two forms of externalities imposed by upstream suppliers onto downstream producers. When firms upstream choose to under-diversify, this generates increased disruption risk downstream. When firms upstream choose to over-diversify, this generates a high correlation in their disruption risk, which has consequences for downstream diversification incentives. Due to this, I show that the supply chain is endogenously fragile to disruption correlation, that is, small increases in the correlation among upstream disruptions can trigger large under-diversification throughout the supply chain. This can leave the economy vulnerable to unlikely disruption events.

To construct a welfare benchmark for these results, I then solve the equivalent social planner problem. I show that a social planner can design a supply chain that is resilient to such correlated shocks. This result has two consequences. First, it illustrates that the fragility derives from the individual firm's diversification strategy and the risk externalities it induces downstream. It is hence entirely endogenous to the supply chain formation. Second, it implies the presence of large and discontinuous welfare losses. Finally, I study the role of imperfect information on supply chain formation. I show that, in the presence of perfect information on the structure and risk of the supply chain, firms are better able to individually diversify risk, yet, in doing so, they choose identical suppliers, which renders the supply chain vulnerable to small upstream disruptions. This suggests that recent efforts in increasing firms' visibility of the supply chain, despite reducing individual firms' supply chain risk, might have the unintended consequence of making supply chains more fragile.

Chapter 2

The Cost of Regret of Irreversible Climate Tipping Points

As the global average temperature rises due to greenhouse gas emissions from human activities, temperature feedback mechanisms in the climate system can drive the world past critical thresholds, known as tipping points, into a regime of high temperatures. These tipping points pose two challenges to determining optimal emissions abatement strategies that try to balance the current benefits of emissions and resulting future climate damages. First, tipping points are hard, if not impossible, to predict (Ben-Yami et al., 2024; Ditlevsen and Johnsen, 2010), giving policymakers no time to react once they are crossed. Second, as crossing a tipping point shifts the climate to a possibly irreversible new regime, the marginal climate damages of adding an additional tonne of CO_2 in the atmosphere suddenly increase. This discontinuity implies that designing optimal emissions abatement policies using marginalist analysis, such as the social cost of carbon, can lead policymakers astray.

In this chapter, I study the economic costs of unpredictable and irreversible climate tipping points. Using an integrated assessment model incorporating temperature feedback effects, I compute optimal emissions abatement strategies under two scenarios: one where a tipping point is imminent and one where it is remote. These scenarios are calibrated to reflect extreme cases from the climate literature (Armstrong McKay et al., 2022; Seaver Wang et al., 2023). I then compute an upper bound on the economic costs of tipping points through two experiments. The first measures the cost incurred by a *wishful thinker* social planner who mistakenly assumes the tipping point is remote, delaying optimal abatement until after it is crossed. This serves as an upper bound on the cost of under-abating. The second considers a *cautious* planner who incorrectly assumes the tipping point is imminent, leading to more aggressive abatement. This serves as an upper bound on the cost of over-

abating. Together, these scenarios provide an upper bound on the economic costs of tipping without requiring prior assumptions about the tipping point's exact location and avoiding marginalist calculations, which are ill-defined in the presence of tipping points. Finally, to compare the costs of over-abating and under-abating, I compute their certainty equivalent: the amount of output society would be willing to give up today to identify the tipping point and switch to the optimal strategy. This quantity measures the risk-adjusted net present cost of the uncertainty and irreversibility of the tipping point. I show that this can be as large as 10.2 trillion US $ per year or 13.5% of current world output. Reacting to this uncertainty with caution and minimising the risk of tipping requires quick abatement efforts, which incur large adjustment costs. Reacting to this uncertainty by wishful thinking and delaying abatement locks the world with high probability into a hot regime characterised by large climate damages. Both are costly approaches, yet wishful thinking can cost up to 1.79 trillion US $ per year more than caution. When faced with uncertain tipping points, it is better to pay the costs of emissions abatement, which are becoming cheaper over time, than to gamble with the risk of crossing a tipping point.

The rest of the paper is structured as follows. Section 2.1 puts the current paper in the context of the climate and environmental economic literature and highlights its contributions. Section 2.2 presents a stylised model version with temperature feedback that generates an irreversible tipping point. This simplified model illustrates the main challenges faced in computing optimal abatement when faced with a tipping point, which carries over into the extended model used for calibration. The components of the extended model are then introduced in Section 2.3: the climate, the economy, and the social planner problem. Section 2.4 computes optimal emissions abatement policies when tipping points are known. Section 2.5 introduces the *wishful thinker* and *cautious* social planner problem and computes the economic damages from tipping points and the cost of uncertainty. Finally, Section 2.6 draws policy implications and concludes.

2.1 Related Literature

A large recent literature in economics has highlighted the importance of correctly incorporating climate dynamics when analysing the economic trade-off between current emissions and the reduction of future climate damages induced by raises in temperature. The key challenge is to develop integrated assessment models that, on the one hand, are sufficiently simple to be integrated into a dynamic optimisation problem, necessary to compute the economic costs of climate change and optimal

emission abatement policies, and, on the other hand, are able to reproduce the key dynamics of larger, more accurate climate models (Dietz et al., 2021*a*). One dynamic aspect of the climate that has drawn a lot of attention in economics are climate tipping points (Li, Crépin and Lindahl, 2024).

Early work dealing with uncertain tipping points in environmental economics has modelled tipping points via stochastic jump processes. Henceforth, I refer to this approach as *Barro disasters approach* as it builds on work on economic catastrophes by Barro (2009). At each moment in time, either temperature or atmospheric CO_2 concentration experiences a discontinuous jump with some probability known to the planner. This approach was introduced to climate economics by Pindyck and Wang (2013)[1]. It has been used in modelling climate tipping points by, for example, Lontzek et al. (2015), Cai, Lenton and Lontzek (2016), and Van der Ploeg and De Zeeuw (2018). Its wide adoption is due to its flexibility as it allows for greater analytical tractability (Li, Crépin and Lindahl, 2024; Lin and Van Wijnbergen, 2023; Van den Bremer and Van der Ploeg, 2021), the introduction of uncertainty, and calibration (Hambel, Kraft and Schwartz, 2021). Nevertheless, this modelling choice neglects two crucial features of climate tipping points that are the focus of this chapter. First, crossing a tipping point can bring the climate to a new persistent regime of high temperatures. Barro disasters are not suited to model these different regimes as they simply model transient catastrophic events. Second, it requires defining a probability distribution on the tipping event, be that real or the belief of the social planner, which is in reality hard to compute reliably.

A second approach is to assume that with some probability the climate system switches to a new, high temperature regime, such that the system dynamics follow a Markov chain (Lemoine and Traeger, 2016*a*,*b*, 2014). Henceforth, I refer to this approach as *Markov chain approach*. This model allows for irreversible tipping points, but displays two limitations. First, it requires an assumption on the distribution of the tipping point, or the social planner to have a prior over it. Lemoine and Traeger (2014) take this to be uniform. This choice is not consistent with the large uncertainty around such tipping points (Ben-Yami et al., 2024). Second, once a tipping point is triggered, its effect are immediately felt and the climate system jumps to a new regime. This is unrealistic as the climate system might converge to a new high temperature regime decades, or even centuries, after the tipping point is crossed.

A final approach is to model feedback in the climate dynamics explicitly. Henceforth, I refer to this as the *feedback approach*. This is closer to the real behaviour of the climate system (McGuffie and Henderson-Sellers, 2014), as it introduces mul-

[1]This approach has an earlier history in the optimal control of environmental systems (Kamien and Schwartz, 1971; Nævdal and Oppenheimer, 2007; Tsur and Zemel, 1996, among others). I refer to Li, Crépin and Lindahl (2024) for a more detailed review.

tiple climate regimes, and is the standard one adopted in more complex climate modelling (e.g. Smith et al. 2017). Furthermore, optimisation over such systems has a long history in economics (Skiba, 1978) and has been developed extensively in environmental economics (Mäler, Xepapadeas and de Zeeuw, 2003; Wagener, 2013) and, specifically, in climate economics (Greiner and Semmler, 2005; Nordhaus, 2019; Wagener, 2015). Modelling the feedback explicitly has one drawback: it requires assuming the size of the feedback and when it comes into effect. These two properties are objects of empirical contention (Armstrong McKay et al., 2022; Seaver Wang et al., 2023). Because of this, the feedback approach has not been widely employed in the climate economic literature, especially when trying to give quantitative estimates of costs via calibration.

In this chapter, by embedding and calibrating a temperature feedback in a state-of-the-art dynamic stochastic general equilibrium model with a climate component, I obtain empirical estimates of the costs of uncertainty of a realistic and irreversible tipping point. The focus on the two extreme scenarios, a *wishful thinker* and a *cautious* planner, allows me to obtain an upper bound on the cost of under-abating and suffering climate damages and over-abating and paying large adjustment costs. These then serve as an upper bound for the overall cost of tipping points. While only upper bounds, these estimates do not require an arbitrary prior on the tipping point and can be computed in a setting with multiple climate regimes. This approach allows to bridge the realistic climate dynamics of the *feedback approach*, the irreversibility of the *Markov chain approach*, and the uncertainty of the *Barro disasters approach*.

Finally, the numerical solver employed in the paper contributes to the literature on controlled stochastic processes by developing a parallelisation algorithm, based on Bierkens, Fearnhead and Roberts (2019), for the class of solver introduced in Kushner and Dupuis (2001), and extending their results to recursive utilities.

2.2 Stylised Example

Before turning to the extended model in the next section (2.3), in what follows I discuss a stylised integrated assessment model with an irreversible tipping point. The climate has a more stylised temperature feedback compared to the extended model. In addition, the planner objective is reduced to a simple optimal stopping problem. Despite its simplicity, this stylised model illustrates the two main challenges tipping points pose to the computation of optimal abatement. First, their unpredictability hinders evaluation of future damages. Second, their irreversibility makes analyses based on marginal emission benefits and marginal temperature

damages misleading.

Emissions of CO_2 from human economic activity E_t $[\mathrm{Gt\,CO_2 y^{-1}}]$ raise CO_2 atmospheric concentration M_t [p.p.m.]

$$\frac{\mathrm{d}M_t}{\mathrm{d}t} = \xi_m E_t, \tag{2.1}$$

where ξ_m converts $\mathrm{Gt\,CO_2}$ to p.p.m.. As CO_2 is a greenhouse gas, an increase in its concentration M_t compared to the pre-industrial levels M^p, in turn, increases average temperature T_t [°]. Hereafter, temperature is expressed in Celsius deviation from its pre-industrial level. The change in temperature is given by

$$\frac{\mathrm{d}T_t}{\mathrm{d}t} = G_1 \log\left(\frac{M_t}{M^p}\right) - \lambda_0 T_t, \tag{2.2}$$

where G_1 is the size of the greenhouse gas effect and λ_0 is a stabilisation rate of temperature. In this model, temperature tends towards the equilibrium

$$T_t \to \frac{G_1}{\lambda_0} \log\left(\frac{M_t}{M^p}\right) \text{ as } t \to \infty, \tag{2.3}$$

that is, the equilibrium temperature grows with the logarithm of atmospheric CO_2 concentration.

Consider now a social planner maximising societal utility. Society derives utility $u(E_t)$ from CO_2 emissions E_t and disutility $d(T_t) \coloneqq e^{\nu T_t}$ from temperature deviations T_t, and discounts both at a rate ρ. For simplicity assume the emissions are constant at a level $\bar{E}$ until a time τ when the social planner chooses to stop emitting

$$E_t = \begin{cases} \bar{E} & \text{if } t \leq \tau, \\ 0 & \text{if } t > \tau. \end{cases} \tag{2.4}$$

This implies that CO_2 concentration grows linearly as

$$M_t = M^p + \xi_m \bar{E} \min\{t, \tau\}. \tag{2.5}$$

For illustration purposes, assume that temperature reaches its equilibrium value immediately, such that

$$T_t = \frac{G_1}{\lambda_0} \log\left(\frac{M_t}{M^p}\right). \tag{2.6}$$

The objective function of the social planner is then to choose a stopping time $\tau > 0$

to maximise

$$\int_0^\infty e^{-\rho t} \left(u(E_t) - d(T_t)\right) \mathrm{d}t \,, \tag{2.7}$$

where T_t is given by (2.2). Finally, assume the planner approximate damages by their value at the stopping time τ, that is, $d(T_t) \approx d(T_\tau)$, such that objective function can be approximated by the function

$$J(\tau) \coloneqq \underbrace{\int_0^\tau e^{-\rho t} u(\bar{E}) \, \mathrm{d}t}_{\text{Emission benefits}} - \underbrace{\int_0^\infty e^{-\rho t} d(T_\tau) \, \mathrm{d}t}_{\text{Climate damages}} . \tag{2.8}$$

The function J is illustrated in Figure 2.1. This problem is readily solved by a

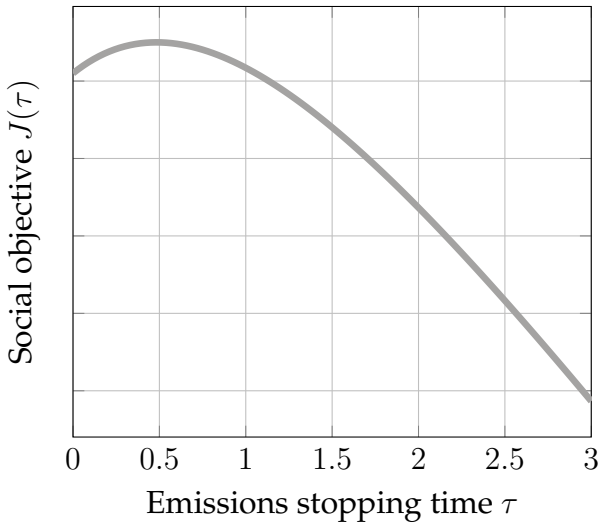

Figure 2.1: Social objective J (2.8) as a function of the emission stopping time τ, under a climate linear climate (2.2).

marginalist planner: at each moment in time $t \geq 0$, the planner can evaluate the marginal benefits of continuing emitting

$$J'(t) = e^{-\rho t} u(\bar{E}) - \frac{1}{\rho} d'(T_t) \partial_t T_t, \tag{2.9}$$

and choose $\tau = \inf\{t \geq 0 : J'(t) = 0\}$ if such optimal time exists or $\tau = \infty$ if this never occurs.

Unfortunately, unlike the stylised model in equation (2.2), the climate system is more complex: there are many temperature feedbacks that can generate tipping points (Seaver Wang et al., 2023). Consider, for example, the ice-albedo feedback (Riihelä, Bright and Anttila, 2021). As CO_2 concentration warms the planet, large bodies of ice melt, such as the west Antarctic or Greenland ice sheets. This reduces the surface albedo, that is, the fraction of solar radiation reflected by the Earth's

surface, which further contributes to warming. This feedback can generate a low temperature regime, in which the ice has not melted, and the albedo is high, and a high temperature regime, in which the low albedo prevents ice from reforming. If a tipping point is crossed, all ice melts and the Earth is a trapped in a warmer regime. In the new warmer regime, ice cannot reform. Hence, the amount of cooling needed for the ice to reform and hence restore the pre-tipping albedo is larger than the warming that melted the ice in the first place. Another possible feedback is that of the carbon release from thawing permafrost (Turetsky et al., 2019). Warming is causing frozen soils to thaw and microorganisms to start breaking down matter in the soil, which in turn releases greenhouse gasses into the atmosphere. Importantly, both of these examples have two features which are not captured in the economic literature: they are hard to predict and are de facto irreversible. The simplest way to introduce these into the stylised model is to assume that the stabilisation rate λ is a smooth function of temperature T_t with the property that

$$T_t \mapsto \mu(M_t, T_t), \text{ with} \tag{2.10}$$

$$\mu(M_t, T_t) \coloneqq G_1 \log\left(\frac{M_t}{M^p}\right) - \overbrace{\lambda(T_t)}^{\text{Feedback}} T_t, \tag{2.11}$$

has two non-degenerate critical points for some values of CO_2 concentration M_t. Temperature dynamics are then given by

$$\frac{\mathrm{d}T_t}{\mathrm{d}t} = \mu(M_t, T_t). \tag{2.12}$$

The simplest feedback function λ that models a climate tipping point and which I will use for illustration purposes in this section, is the quadratic

$$\lambda(T_t) = \lambda_0 + \lambda_1 T_t + \lambda_2 T_t^2, \tag{2.13}$$

where λ_1 and λ_2 are chosen such that current temperature is in equilibrium $T_0 \approx 1.14°$ is in equilibrium under the temperature dynamics (2.12). Figure 2.2 shows the change in temperature $\mathrm{d}T_t \,/\, \mathrm{d}t$ as a function of temperature T_t when CO_2 concentration is at its pre-industrial level (left panel), 50% larger (middle panel) and twice as large (right panel). As in the model without feedback (2.2), if CO_2 concentration is at its pre-industrial level (left panel) there is a low temperature equilibrium $T_t \to 0°$ (low temperature grey marker). In addition to this, a new high temperature equilibrium is formed $T_t \to 5°$ (high temperature black marker). In this high temperature equilibrium, the feedback has trapped temperature into a high regime, for example, an ice free world with low albedo. As CO_2 concentration increases

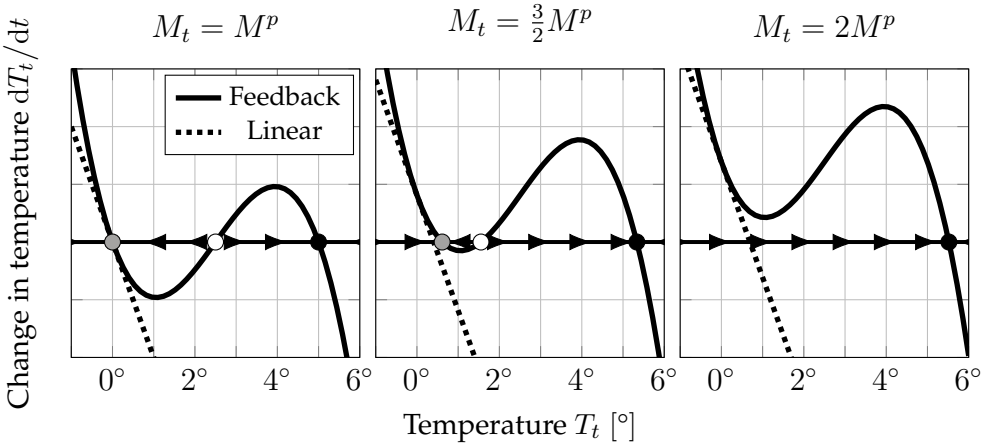

Figure 2.2: Change in temperature over time $\frac{\mathrm{d}T_t}{\mathrm{d}t}$ as a function of temperature T_t for three levels of atmospheric CO_2 concentration, under the feedback dynamics (2.12). Markers represent the equilibria and the arrows indicate whether temperature is increasing (right) or decreasing (left). As reference, the graph is shown for the linear dynamics (2.2) as a dotted line.

(from left to right panel), the average temperature rises. Unlike the case without feedback, the system goes through a tipping point[2]: the low temperature equilibrium vanishing triggers a sudden rise in temperature as it converges to the higher equilibrium. This phenomenon poses two problems from the point of view of the social planner trying to evaluate the trade-off between current emissions and future climate damages.

First, consider the social planner trying to solve the optimal emission stopping problem (2.8) and suppose the planner does not know whether the true model is linear (2.2) or has a tipping point (2.12). The planner hopes to learn which model it is facing by observing the temperature trajectory T_t. Let T_t^{f} and T_t^{l} be the trajectory of temperature under the feedback (2.12) and linear model (2.2) respectively. Notice that temperature has the same initial condition $T_0^{\mathrm{l}} = T_0^{\mathrm{f}} = T_0$, and they grow at the same rate $\dot{T}_0^{\mathrm{l}} = \dot{T}_0^{\mathrm{f}}$. Because of this the two trajectories are equal up to a quadratic term

$$T_t^{\mathrm{l}} - T_t^{\mathrm{f}} = \mathcal{O}(t^2). \tag{2.14}$$

The two paths of temperature are illustrated in Figure 2.2. The two trajectories are similar, until the tipping point is reached. After crossing it, they rapidly diverge. If the planner does not perfectly observe the trajectory, for example due to noise in the observations, or they use a linear model to infer which path temperature they are on, disentangling the two is only possible if the CO_2 concentration relative to its initial level $\log M_t/M_0$ grows large. Despite being presented in a highly stylised form here, this is the core mechanism behind the complications that arise when trying to predict tipping points in more complex climate models (Ben-Yami et al.,

[2]This is a saddle-node bifurcation.

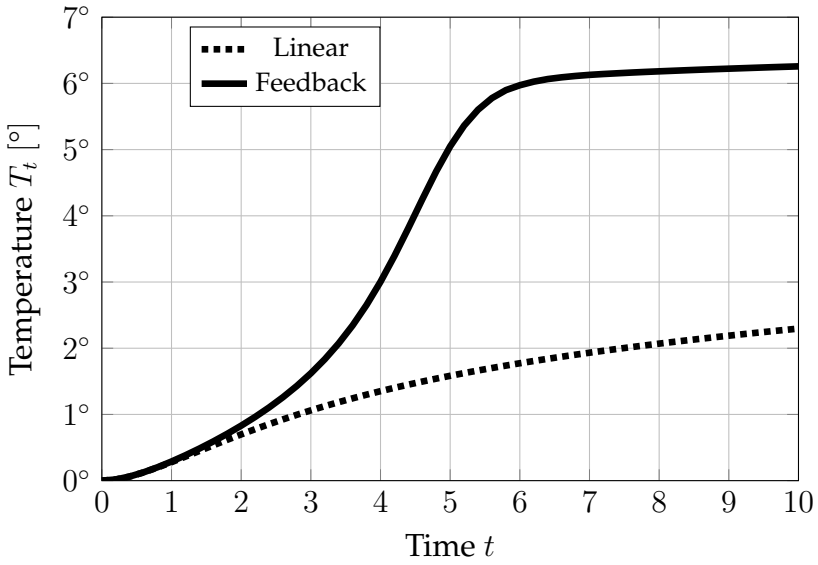

Figure 2.3: Path of temperatures given by (2.12), assuming constant emissions $E_t = \bar{E}$ with (solid) and without (without) temperature feedback.

2024; Ditlevsen and Johnsen, 2010).

The second reason the presence of tipping points can hinder computing optimal stopping times is its irreversibility. Consider a marginalist social planner trying to solve the optimisation problem (2.8) by computing the marginal continuation benefit $J'(t)$ (2.9), at each moment in time t and postponing the end of emissions as long as $J'(t) > 0$. Then $\tau = \inf\{t \geq 0 : J'(t) = 0\}$. As illustrated in Figure 2.2, initially, the equilibrium temperature $\bar{T}_\tau$ reached after emitting for a period τ can take two possible value: a high or a low one. As a consequence, the social utility $J(\tau)$ of emitting for a period τ can also take two values depending on which temperature regimes the planner is in. The objective J for different stopping times τ in the low and high temperature regimes is shown in grey and black, respectively, in Figure 2.4. These equilibria correspond to the grey and black marker in Figure 2.2. Initially, the objective is given by the higher, lighter branch in Figure 2.4. Hence, initially, delaying abating emissions is optimal, as the marginal benefits of emissions are greater than the marginal climate damages, that is, J' is positive. Yet by delaying abatement, the climate is brought closer to a tipping point (grey marker in Figure 2.4). At that point, the marginal damages from an increase in emissions are not defined as any further delay to stopping emissions pushes the climate into the high temperature regime. The costs of tipping are hence not internalised in the marginal continuation benefits J'. Once the system has tipped, the objective is given by the lower, darker line in Figure 2.4. After tipping, the system cannot be reverted to the low temperature equilibrium, as CO_2 concentration remains constant. In this new regime (black line) the planner wish they had stopped as early as possible, as J' is always negative and J is maximised at $\tau = 0$.

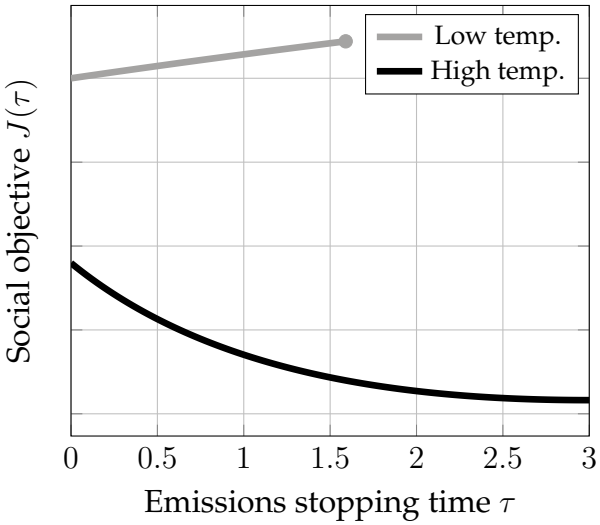

Figure 2.4: Social objective J (2.8) as a function of the emission stopping time τ, under a climate with temperature feedback (2.12).

This simple example illustrates how tipping points can hinder the computation of optimal emissions via their marginal benefits and marginal damages. When the tipping point is hard to predict the path of marginal damages is uncertain. When the tipping point leads to a new irreversible regime the marginal damages are ill-defined along some of these paths.

2.3 Model

The following extends the stylised model from Section 2.2. First, Section 2.3.1 introduces carbon sinks and their interaction with atmospheric CO_2 concentration. Then, Section 2.3.1 presents the temperature dynamics and the feedback mechanism. Section 2.3.1 discusses the resulting tipping point in the climate system. Finally, Section 2.3.2 presents the Harrod–Domar economy.

2.3.1 Climate

Carbon sinks and CO_2 concentration

Emissions from human economic activity E_t increase the atmospheric concentration of CO_2 M_t. Part of this, in turn, decays into natural sinks. I denote by N_t the $\mathrm{Gt}\,\mathrm{CO}_2$ stored in natural sinks. As these sinks saturate, the rate of decay $\delta_m(N_t)$ depends on the stored CO_2 N_t. Hence, the stored CO_2 evolves as

$$\xi_m \,\mathrm{d}N_t = \delta_m(N_t) M_t \;\mathrm{d}t\,, \tag{2.15}$$

where ξ_m converts $\mathrm{Gt\,CO_2}$ to $\mathrm{p.p.m.}$. This assumption on carbon sinks is in line with the empirical evidence (Le Quéré et al., 2007; Shi et al., 2021) and reflects the fact that it is much easier to stabilise atmospheric $\mathrm{CO_2}$ concentration M_t if natural carbon sinks are not saturated, that is, for low values of N_t.

In absence of abatement efforts, atmospheric $\mathrm{CO_2}$ concentration M_t evolves according to the *no-policy* SSP5 scenario (Kriegler et al., 2017). Variables under this scenario are indexed by the superscript np: let E_t^{np} be the no-policy emissions, and M_t^{np} and N_t^{np} the resulting atmospheric $\mathrm{CO_2}$ concentration and the $\mathrm{CO_2}$ stored in natural sinks, respectively. No-policy atmospheric $\mathrm{CO_2}$ concentration M_t^{np} evolves as

$$\frac{\mathrm{d}M_t^{\mathrm{np}}}{M_t^{\mathrm{np}}} = \gamma_t^{\mathrm{np}}\,\mathrm{d}t + \sigma_m\,\mathrm{d}W_{m,t}\,, \tag{2.16}$$

where the no-policy expected growth rate of atmospheric $\mathrm{CO_2}$ concentration is given by

$$\gamma_t^{\mathrm{np}} \coloneqq \xi_m \frac{E_t^{\mathrm{np}}}{M_t^{\mathrm{np}}} - \delta_m(N_t^{\mathrm{np}}) \tag{2.17}$$

and $W_{m,t}$ is a Brownian motion. The Brownian motion accounts for possible uncontrolled and unexpected shocks to $\mathrm{CO_2}$ concentration. The variance σ_m^2 of these shocks is assumed to be constant. The no-policy scenario describes an energy intensive future, in which fossil fuel usage develops rapidly and little to no abatement takes place. Using the projected $\mathrm{CO_2}$ concentration path M_t^{np} from Kriegler et al. (2017), I calibrate the implied growth rate of $\mathrm{CO_2}$ concentration γ_t^{np}. The calibration is described in Appendix B.3 and its results are shown in Figure 2.5. The left figure

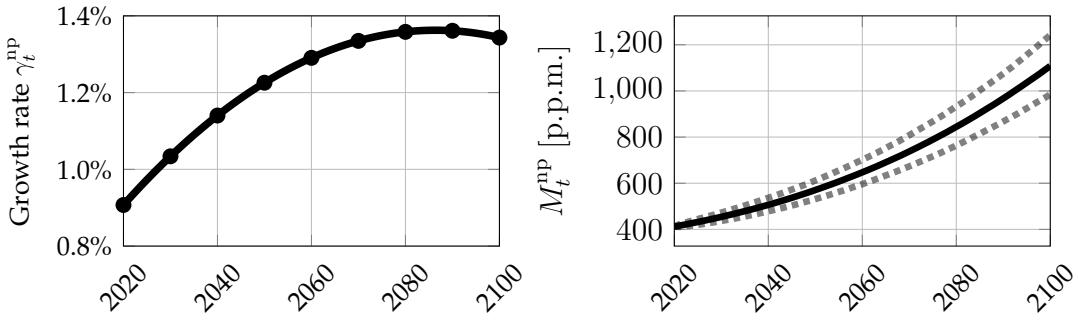

Figure 2.5: Growth rate of $\mathrm{CO_2}$ concentration in the no-policy scenario γ_t^{np} and median path (solid) of no-policy $\mathrm{CO_2}$ concentration M_t^{np} (2.16) with 95% simulation interval (dashed).

shows the path of the no-policy growth rate γ_t^{np} and the right figure shows the resulting growth of $\mathrm{CO_2}$ concentration M_t^{np}. The $\mathrm{CO_2}$ concentration in this scenario grows at an increasingly fast rate until 2080, when the growth rate peaks just below 1.4%. Thereafter, the growth rate starts declining.

Abatement efforts α_t lower the growth rate of $\mathrm{CO_2}$ concentration M_t vis-à-vis the

no-policy scenario M_t^{np}. That is, the growth rate of CO_2 concentration M_t satisfies

$$\frac{\mathrm{d}M_t}{M_t} = \left(\gamma_t^{\text{np}} - \alpha_t\right)\mathrm{d}t + \sigma_m\,\mathrm{d}W_{m,t}\,, \tag{2.18}$$

where $\mathrm{d}W_{m,t}$ is a standard Brownian motion. Not implementing any abatement policy $\alpha_t = 0$ corresponds to a no-policy scenario $M_t = M_t^{\text{np}}$, while implementing a net-zero abatement policy $\alpha_t = \gamma_t^{\text{np}}$ stabilises CO_2 concentration. Any abatement policy α_t can be linked to the implied level of emissions by introducing the emission reduction rate $\varepsilon(\alpha_t)$, which keeps tracks of what fraction of emission is being abated

$$E_t = \left(1 - \varepsilon(\alpha_t)\right)E_t^{\text{np}}. \tag{2.19}$$

In this chapter I assume that the planner does not have access to carbon capture technology. This is implemented by imposing that abatement cannot exceed the growth rate of CO_2 concentration and the carbon decay into natural sinks, namely

$$\alpha_t \leq \gamma_t^{\text{np}} + \delta_m(N_t). \tag{2.20}$$

Hence, if no abatement efforts are undertaken $\alpha_t = 0$, no emissions are abated, so $\varepsilon(0) = 0$ and $E_t = E_t^{\text{np}}$. If abatement efforts are maximal $\alpha_t = \gamma_t^{\text{np}} + \delta_m(N_t)$, net-zero emissions are reached, so $\varepsilon(\gamma_t^{\text{np}} + \delta_m(N_t)) = 1$ and $E_t = 0$.

Temperature

Average world temperature T_t [°] in deviation from pre-industrial levels is determined by a zero dimensional energy balance model (McGuffie and Henderson-Sellers, 2014, p.85). Earth's energy balance, in its simplest form, prescribes that temperature tends towards balancing incoming solar radiation S and outgoing longwave radiations $\eta\sigma T_t^4$, where σ is the Stefan-Boltzmann constant and η is an emissivity rate. Due to the presence of greenhouse gases, certain wavelengths are scattered and not radiated outwards (Ghil and Childress, 2012). This introduces an additional radiative forcing G which yields the balance equation $S = \eta\sigma T_t^4 - G$. Focusing on the role of increased CO_2, as opposed to other greenhouse gases, the greenhouse radiative forcing term G can be decomposed into a constant component G_0 and a component which depends on the equilibrium level of CO_2 concentration in the atmosphere M_t relative to the pre-industrial level M^p, such that

$$G = G_0 + G_1 \log(M_t/M^p). \tag{2.21}$$

I now introduce the main positive temperature feedback by assuming that the

absorbed incoming solar radiation is increasing in temperature. This choice can be seen as a stylised model of the ice-albedo feedback (McGuffie and Henderson-Sellers, 2014), but other temperature feedback mechanisms yield a similar interpretation. Solar radiation S is then decomposed as $S_0\,(1-\lambda(T_t))$ where the function $\lambda(T_t)$ transitions from a higher level λ_1 to a lower level $\lambda_1 - \Delta\lambda$ via a smooth transition function $L(T_t - T^c)$, that is,

$$\lambda(T_t) \coloneqq \lambda_1 - \Delta\lambda\, L(T_t - T^c), \tag{2.22}$$

where T^c $[^\circ]$ is the critical level of temperature at which the feedback comes into effect and the transition function satisfies $L(T_t - T^c) \to 1$ as $T_t \to \infty$. The functional form of L used to calibrate (2.22) is discussed in Appendix B.3.3.

The threshold temperature T^c determines the horizon at which the feedback comes into effect. Best estimates from climate sciences are that the most relevant transitions occur for temperatures T_t between $1.5°$ and $2.5°$ over pre-industrial levels (Armstrong McKay et al., 2022; Seaver Wang et al., 2023). A large body of literature has focused on estimating critical thresholds associated with climate tipping points (e.g. Boulton, Allison and Lenton 2014; Van Westen, Kliphuis and Dijkstra 2024). Yet, for many climate feedback mechanisms the uncertainty is too large to compute their tipping points (Ben-Yami et al., 2024; Ditlevsen and Johnsen, 2010; Wagener, 2013). In line with this, to avoid attaching an arbitrary prior over a critical threshold T^c, in this chapter I choose to focus on two extreme scenarios: one in which the tipping point is *imminent* $T^c = 1.5°$ and one in which it is *remote* $T^c = 2.5°$. The parameter $\Delta\lambda$ and the transition function L are calibrated by matching the equilibrium climate sensitivity, that is, the expected equilibrium temperature of doubling pre-industrial level CO_2 concentration, of $4°$, which is the upper end of the range deemed "very likely" in AR6 (IPCC, 2023) (see Appendix B.3.3 for more details on the calibration). Figure 2.6 shows the transition function (2.22) under these two scenarios. Despite being a highly stylised and reduced form representation of a complex and spatially heterogeneous process, λ captures the core mechanism behind feedback processes in the temperature dynamics that can generate tipping points (McGuffie and Henderson-Sellers, 2014). This simple temperature feedback allows us to discuss and estimate the costs of transitioning to a new regime, what ought to be (or not be) done to prevent this from happening, and the cost associated with the uncertainty around the tipping point.

Putting these processes together we can write down the two determinants of temperature dynamics: radiative forcing, which only depends on temperature,

$$r(T_t) \coloneqq S_0\,(1-\lambda(T_t)) - \eta\sigma T_t^4 \tag{2.23}$$

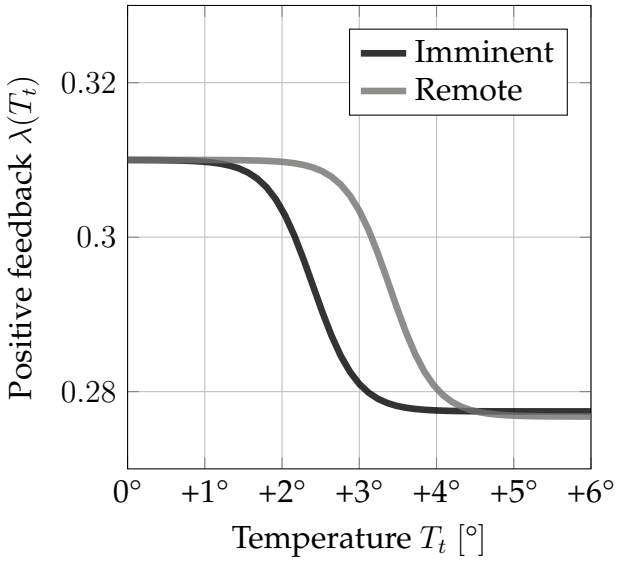

Figure 2.6: Temperature feedback $\lambda(T_t)$ in the *imminent* and the *remote* scenario.

and the greenhouse gas effect, which only depends on atmospheric CO_2 concentration

$$g(M_t) \coloneqq G_0 + G_1 \log(M_t/M^p). \tag{2.24}$$

Under these two drivers, temperature changes are given by

$$\epsilon \, \mathrm{d}T_t = r(T_t) \, \mathrm{d}t + g(M_t) \, \mathrm{d}t + \sigma_T \ \mathrm{d}W_{T,t} \,, \tag{2.25}$$

where ϵ is Earth's total heat capacity and $W_{T,t}$ is a Brownian motion. The total heat capacity ϵ is estimated to match the time necessary for temperature deviations T_t to convergence to equilibrium after an increase in CO_2 concentration M_t. The Brownian motion models uncontrolled shocks to temperature, which are assumed to have constant variance σ_T^2 and be uncorrelated with shocks $W_{m,t}$ to atmospheric CO_2 concentration.

Temperature Dynamics with a Tipping Point

The presence of the temperature feedback λ introduces a tipping point. In the following, I discuss the consequences of this tipping point for the temperature dynamics.

For a given level of CO_2 concentration M_t, the temperature T_t tends towards an equilibrium $\bar{T}$ that gives radiative balance in the temperature dynamics (2.25), that is,

$$r(\bar{T}) = -g(M_t). \tag{2.26}$$

As in the stylised example of Section 2.2, the temperature feedback yields a climate with multiple regimes. That is, for the same level of atmospheric CO_2 concentra-

tion M_t, there are multiple levels of temperature $\bar{T}$ which are in equilibria. Figure 2.7 shows the equilibria of temperature (solid lines) as a function of CO_2 concentration when the tipping point is imminent (darkest line), when it is remote (lighter line), and when there is no temperature feedback λ (lightest line). Without tem-

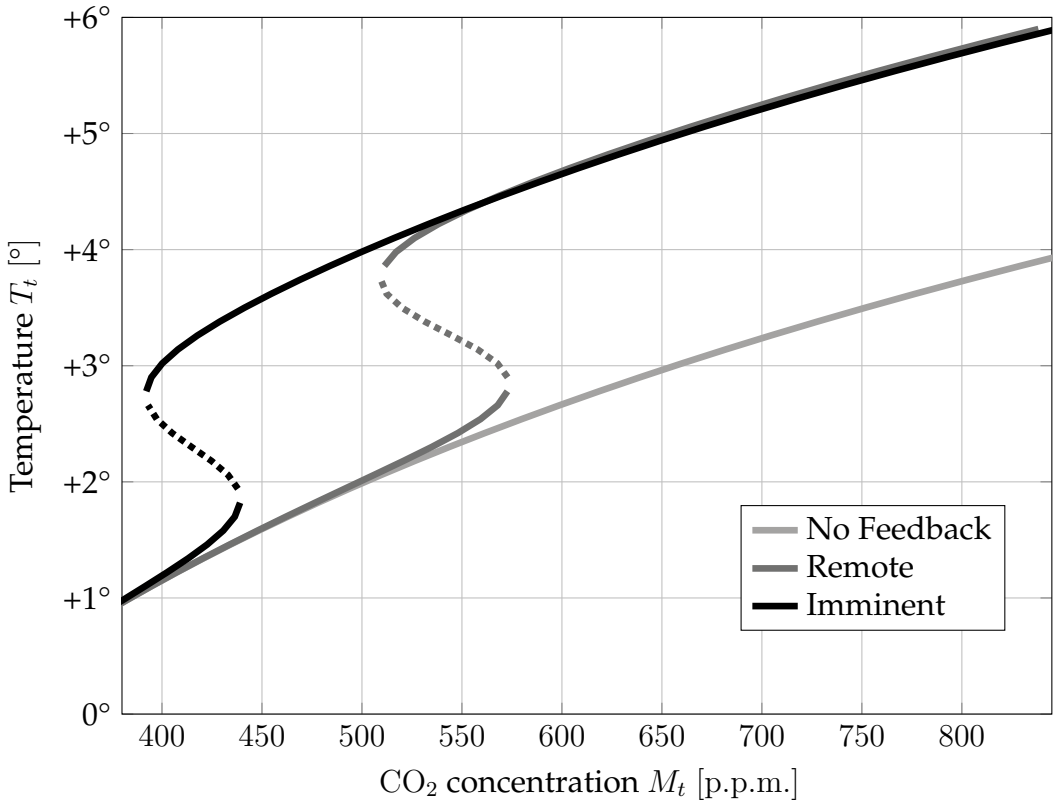

Figure 2.7: Equilibria of temperature T_t and CO_2 concentration M_t for an imminent tipping point, a remote tipping point, and no temperature feedback λ. The solid lines show attracting equilibrium and the dashed line repelling equilibrium.

perature feedback λ, the equilibrium temperature is unique, and it rises with the logarithm of CO_2 concentration. If a temperature feedback λ is introduced there are two possible temperature regimes: a low temperature regime and a high temperature regime. These coexist for some values of CO_2 concentration (dashed lines). If the temperature feedback comes into effect imminently (imminent scenario), for a CO_2 concentration M_t approximately between 390 p.p.m. and 440 p.p.m. both a low temperature regime, with $T_t < 1.9°$, and a high temperature regime, with $T_t > 2.9°$ are possible. In the remote scenario this phenomenon occurs only when CO_2 concentration M_t is approximately between 510 p.p.m. and 570 p.p.m.. As illustrated in Section 2.2, in a low temperature regime, the presence of a high temperature regime is hard to detect as the relationship between temperature T_t and CO_2 concentration M_t is not affected before the temperature feedback comes into effect.[3] If CO_2

[3]Technically, the stable manifold of the model without feedback is a contact element of the stable manifold in the cases with feedback.

concentration M_t increases and temperature crosses the critical threshold T^c, the old stable and low temperature regime is no longer feasible and only a high stable temperature regime remains. Then, any increase of CO_2 concentration pushes the system past a tipping point and temperature rapidly converges to the high temperature equilibrium. Crucially, to revert the system back to the lower temperature regime, it is not sufficient to remove just the carbon that caused the system to tip, but it is necessary to remove all carbon until the only stable equilibrium is the low temperature. In Figure 2.7, this would amount to bringing carbon back to $M_t < 390$, where the high temperature regime (solid black line) vanishes.

We can now ask what path of temperature T_t we would observe in these three scenario if no abatement is implemented and CO_2 concentration M_t grows at the no-policy rate γ_t^{np}, such that $M_t = M_t^{\text{np}}$. This is illustrated in Figure 2.8, which shows the median path over 60 years (marked line) of temperatures T_t and CO_2 concentration M_t under no-policy for the model without feedback and the two tipping scenarios, remote and imminent. Each marker is the outcome at the beginning of each decade. The simulations are overlaid onto the equilibria from Figure 2.7. In the case of an imminent tipping point (darkest line), the temperature deviates from the other scenarios by 2030. In case the tipping point is remote, temperature growth cannot be distinguished from a situation without temperature feedback and the two diverge only starting in 2050. In both cases, once the tipping point is crossed and the only regime is the high temperature one, temperature grows quickly to the new equilibrium. After that, the linear relationship between temperature and (log) CO_2 concentration is restored.

Irreversible Regime Changes versus Barro Disasters

Before turning to the economic model, in what follows I illustrate why the climate system presented here provides a more adequate representation of tipping points than the Barro disasters (Barro, 2009; Pindyck and Wang, 2013) commonly employed in the literature. When modelling a tipping point as a Barro disaster, one assumes that at each moment in time, with some arrival rate $\pi(T_t)$, there is a temperature shock of size $q(T_t)$, both of which are increasing in temperature (e.g. Hambel, Kraft and Schwartz 2021; Lemoine and Traeger 2016*a*; Van der Ploeg and De Zeeuw 2018). Formally, under this assumption, equation (2.25) of the climate model presented here becomes

$$\epsilon \, \mathrm{d}T_t = r_l(T_t) \, \mathrm{d}t + g(M_t) \, \mathrm{d}t + \sigma_T \, \mathrm{d}W_{T,t} + q(T_t) \, \mathrm{d}J_t \, (T_t), \tag{2.27}$$

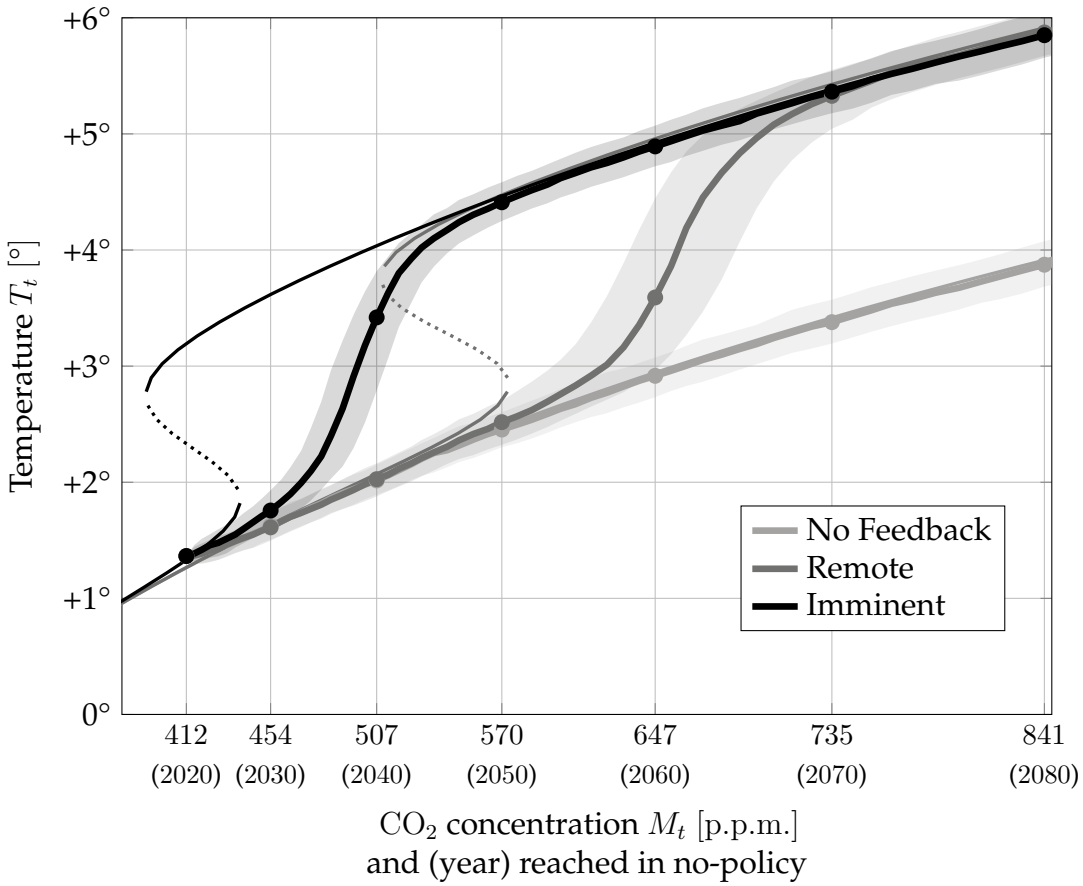

Figure 2.8: Median temperature T_t and CO_2 concentration M_t at time t of 100,000 simulations in case of no feedback (lightest line), a remote tipping point (darker line), and an imminent tipping point (darkest line). The simulations are generated under the temperature dynamics (2.25) and the no-policy CO_2 concentration (2.16). Each marker is the value at the beginning of each decade. The thinner lines represent the equilibrium relationships as in Figure 2.7.

where J_t is a Poisson process with arrival rate $\pi(T_t)$ and r_l, unlike r (2.23), contains no positive feedback mechanism. For comparison, Figure 2.9 shows the distribution of temperature generated by the model of this chapter (2.25) and a model with Barro disasters (2.27), as calibrated in Hambel, Kraft and Schwartz (2021). Both models capture a key feature of tipping points: as CO_2 concentration rises, we expect warmer shocks to be more persistent and hence the equilibrium temperature distribution to be skewed (IPCC, 2023; Weitzman, 2014). Yet, the model with Barro disasters commonly employed in the literature does not display multiple temperature regimes (lower figure). Any shock to temperature is transient: in absence of a new shock, temperature converges back to the unique equilibrium level. This model is hence not suited to answer questions about irreversible tipping points, such as those present in the climate. On the other hand, the model (2.25) used in this chapter (upper figure) generates multiple temperature regimes while still producing a skewed temperature distribution. As discussed in the stylised example in Section 2.2, a climate with multiple regimes can lead to vastly different optimal

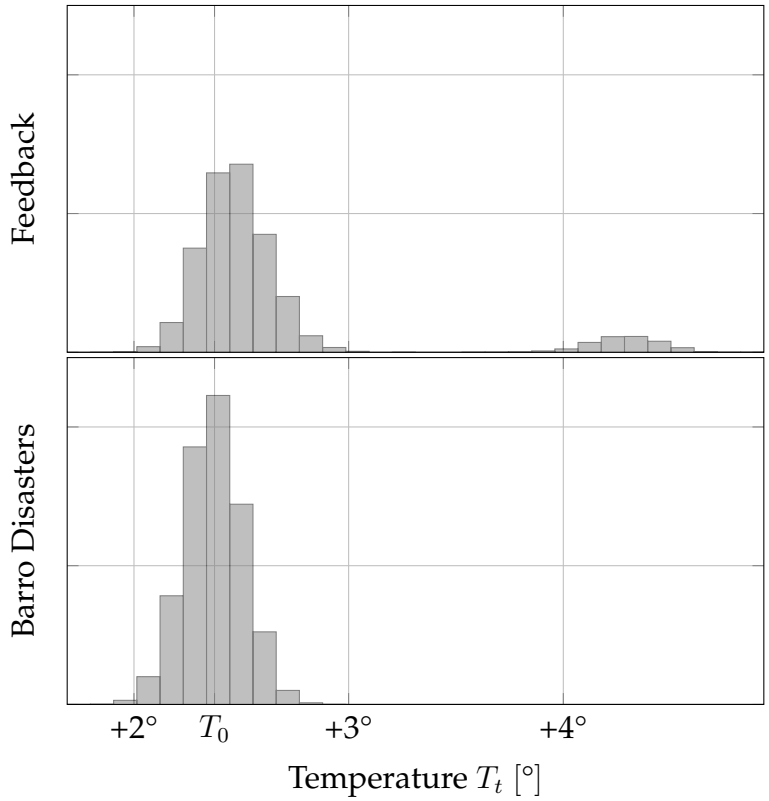

Figure 2.9: Distribution of temperature T_t simulated under (2.25) (Feedback model) and (2.27) (Barro Disasters model) for 1000 years at a constant CO_2 concentration level $M = 540$ p.p.m. starting at an initial temperature $T_0 = 1.14°$. See Appendix B.4 for details on the calibration of the Barro disaster.

abatement paths and costs of tipping.

2.3.2 Economy

This section introduces the economy. The model assumes a Harrod-Domar economy, following Pindyck and Wang (2013), and Hambel, Kraft and Schwartz (2021).

Capital Accumulation and Climate Damages

Output Y_t, in trillion of US \$ per year, is the product of the capital stock K_t and its productivity A_t, that is,

$$Y_t = A_t K_t. \tag{2.28}$$

Productivity A_t grows at a constant rate ϱ. Output Y_t can be used for investment in capital I_t, abatement expenditures B_t, or consumption C_t. This implies the budget constraint

$$Y_t = I_t + B_t + C_t. \tag{2.29}$$

Capital K_t depreciates at a rate δ_k but can be substituted by capital investments I_t, which are subject, along with abatement expenditure B_t, to quadratic adjustment

costs

$$\frac{\kappa}{2}\left(\frac{I_t + B_t}{K_t}\right)^2 K_t. \tag{2.30}$$

Climate change affects the economy by lowering capital growth via damages $d(T_t)$ which are increasing in temperature T_t. Following Weitzman (2012), I assume the damage function to take the form

$$d(T_t) := \xi\, T_t^v. \tag{2.31}$$

The calibrated damage function is displayed in Figure 2.10. This stylised form cap-

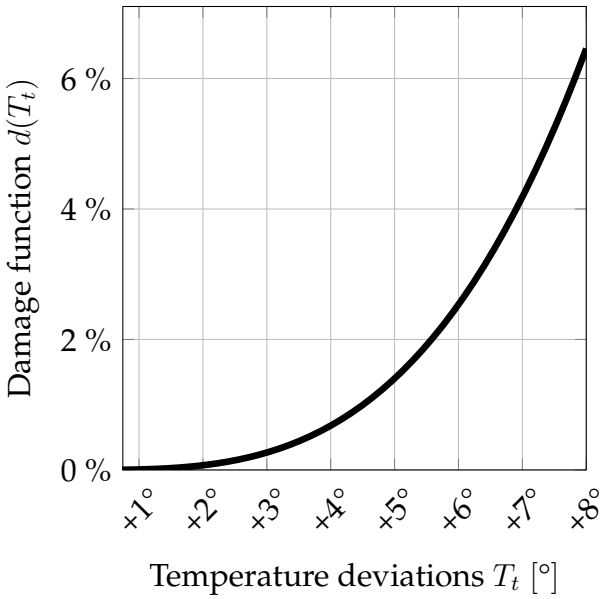

Figure 2.10: Damage function with the Weitzman (2012) calibration.

tures the empirical evidence that under higher temperature levels some forms of capital, particularly in the agricultural (Dietz and Lanz, 2019) and manufacturing sectors (Dell, Jones and Olken, 2009, 2012), become more expensive or harder to substitute. A common alternative in the literature is to assume that higher temperatures wipe out part of the capital stocks (Nordhaus, 1992). A thorough comparison of these two assumptions can be found in Hambel, Kraft and Schwartz (2021).

Combining the endogenous growth of capital and the climate damages, the growth rate of capital satisfies

$$\frac{\mathrm{d}K_t}{K_t} = \left(\frac{I_t}{K_t} - \delta_k - \frac{\kappa}{2}\left(\frac{I_t + B_t}{K_t}\right)^2\right)\mathrm{d}t - d(T_t)\,\mathrm{d}t + \sigma_k\,\mathrm{d}W_{k,t}\,, \tag{2.32}$$

where $W_{k,t}$ is a standard Brownian motion.

In the following, I link the abatement costs B_t with the abatement rate α_t, introduced in the previous section. Define $\beta_t := B_t/Y_t$ to be the fraction of output

devoted to abatement. I assume this to be a quadratic function of the fraction of abated emissions $\varepsilon(\alpha_t)$ (2.19), namely,

$$\beta_t(\varepsilon(\alpha_t)) = \frac{\omega_t}{2}\,\varepsilon(\alpha_t)^2. \tag{2.33}$$

Under this assumption, no abatement is free, as $\beta_t(0) = 0$. At a fixed time t, a higher abatement rate α_t and hence a higher emission reduction $\varepsilon(\alpha_t)$ vis-à-vis the no-policy scenario, becomes increasingly costly at linear rate

$$\beta_t'(\varepsilon(\alpha_t)) = \omega_t \varepsilon(\alpha_t). \tag{2.34}$$

It is common in the literature to assume the marginal abatement costs to be proportional to output and linear in abatement efforts (Baker, Clarke and Shittu, 2008; Dietz and Venmans, 2019; Nordhaus, 1992, 2017). As noted by Dietz and Venmans (2019, p.112-113), the proportionality with output arises because higher output growth increases energy demand. This must be satisfied with low-carbon energy technology which display decreasing marginal productivity. The linearity in abatement efforts, meanwhile, matches the empirical estimates provided in the IPCC Fifth Assessment Report (2023).

As time progresses, so does abatement technology and a given abatement objective becomes cheaper, as a fraction of output. This is modelled by letting the exogenous technological parameter ω_t decrease exponentially over time

$$\omega_t = \omega_0 e^{-\rho_\omega t}, \tag{2.35}$$

at a rate ρ_ω. Following (Nordhaus, 2017), I assume that full abatement $\varepsilon(\alpha_t) = 1$ costs 11% of GDP in 2020 and 2.7% of GDP in 2100. These estimates are in line with a large literature estimating marginal abatement curves (see the meta-analysis by Kuik, Brander and Tol, 2009). The resulting marginal abatement curves are displayed in Figure 2.11.

Finally, let

$$\chi_t := \frac{C_t}{Y_t} \tag{2.36}$$

be the fraction of output devoted to consumption. Using the budget constraint (2.29) and the two controlled rates, the abatement efforts α_t and the consumption rate χ_t, the growth rate of capital (2.32) can be rewritten, in terms of growth rates, as

$$\frac{\mathrm{d}K_t}{K_t} = \Big(\underbrace{\phi_t(\chi_t)}_{\text{growth}} - \overbrace{A_t\beta_t(\varepsilon(\alpha_t))}^{\text{abatement}} - \underbrace{d(T_t)}_{\text{climate}}\Big)\,\mathrm{d}t + \sigma_k\,\mathrm{d}W_{k,t}\,, \tag{2.37}$$

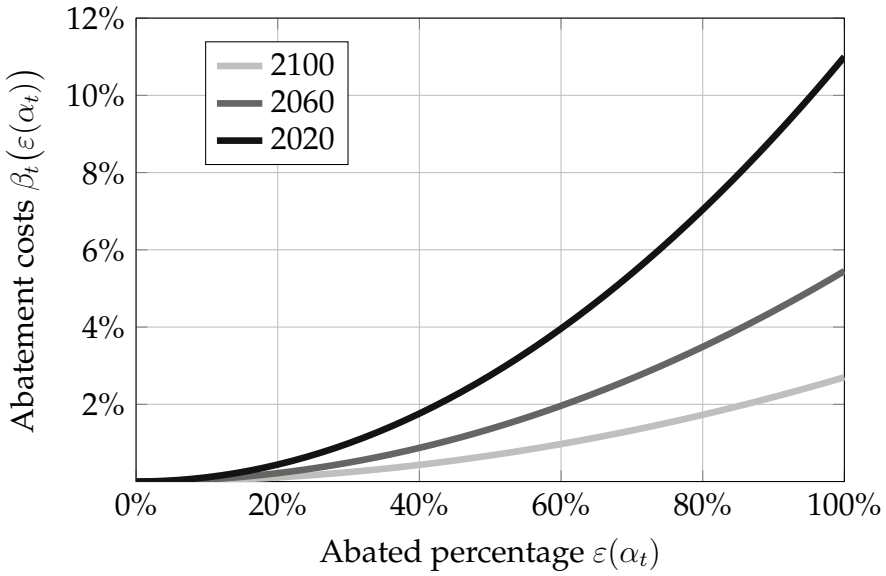

Figure 2.11: Calibrated marginal abatement curves $\beta_t'(\varepsilon) = \omega_t \varepsilon$ at different times t.

where

$$\phi_t(\chi_t) \coloneqq A_t(1-\chi_t) - \frac{\kappa}{2} A_t^2 (1-\chi_t)^2 - \delta_k \tag{2.38}$$

is an endogenous growth component, $\beta_t(\varepsilon(\alpha_t))$ is the fraction of output allocated to abatement for an abatement rate α_t, and $d(T_t)$ are the damages from climate change. As output Y_t is the product of capital K_t and productivity A_t, its growth rate differs from that of capital by the growth rate of productivity

$$\frac{\mathrm{d}Y_t}{Y_t} = \varrho + \frac{\mathrm{d}K_t}{K_t}. \tag{2.39}$$

This formulation models the trade-off between climate abatement and economic growth. On the one hand, devoting fewer resources to abatement (lower α_t and, hence, $\beta_t(\varepsilon(\alpha_t))$) to pursue higher capital growth and output growth, yields higher future temperature and can put the economy in a lower growth path altogether (higher $d(T_t)$). On the other hand, more ambitious abatement lowers future temperature and climate damages (lower $d(T_t)$), which boosts future economic growth, but costs economic growth today (higher α_t and, hence, $\beta_t(\varepsilon(\alpha_t))$).

2.4 Optimal Abatement

Given the climate and economic dynamics described in the previous section, I introduce the objective of the social planner and the resulting optimal policies. At time t given the state of temperature, CO_2 concentration, carbon in sinks, and output

$X_t := (T_t, M_t, N_t, Y_t)$, societal utility is recursively defined as

$$V_t(X_t) = \sup_{\chi,\alpha} \mathbb{E}_t \int_t^\infty f(\chi_s(X_s)Y_s, V_s(X_s))\,\mathrm{d}s\,, \tag{2.40}$$

where χ and α are continuous policies over time and the state space, and f is the Epstein-Zin aggregator

$$f(C,V) := \rho\frac{(1-\theta)\,V}{1-1/\psi}\left(\left(\frac{C}{((1-\theta)V)^{\frac{1}{1-\theta}}}\right)^{1-1/\psi} - 1\right). \tag{2.41}$$

Consumption is integrated into a utility index by means of the Epstein-Zin aggregator (Duffie and Epstein, 1992). This aggregator plays a dual role. First, it disentangles the role of relative risk aversion θ, elasticity of intertemporal substitution ψ, and the discount rate ρ in determining optimal abatement paths. Second, it circumvents the known paradoxical result that abatement policies become less ambitious as society becomes more risk averse (Pindyck and Wang, 2013).

As in Hambel, Kraft and Schwartz (2021), I set $\rho = 1.5\%$, $\theta = 10$ and $\psi = 0.75$. There is no consensus around the calibration of these parameters. I refer the reader to Pindyck and Wang (2013, Section 2), for a more detailed discussion on the calibration of the preference parameters, and to Hambel, Kraft and Schwartz (2021) for their role in determining optimal abatement. Furthermore, the formulation assumes that preferences are stable over time, both in the parameters θ, ψ, and ρ and in their depending only on the consumption stream C_t. Future society might exhibit different preferences, for example, there might be an inherent utility from lower CO_2 concentration or temperature levels. For a discussion on incorporating this possibility in the analysis I refer the reader to Le Kama and Schubert (2004). In Appendix B.5, I test the results of the paper to a different coefficient of relative risk aversion.

The social planner problem (2.40) is solved numerically. The details of the numerical procedure are outlined in Appendix B.1.

As mentioned above, we consider two possible dynamics for temperature T_t: one with an imminent and one with a remote tipping point. Let $\underline{\alpha}_t$ and $\overline{\alpha}_t$ be the optimal abatement in case of an imminent and a remote tipping point respectively. Figure 2.12 shows the resulting optimal policies. Each panel shows the fraction $\varepsilon_t := \varepsilon(\alpha_t)$ of abated emissions compared to the no-policy as a function of the current CO_2 concentration M_t, in the imminent (left) and remote (right) tipping point. As the tipping point generates multiple temperature regimes for a given CO_2 concentration level M_t, each panel shows two curves. The lighter curve represents the optimal abatement before tipping, in a low temperature regime. The darker curve represents the optimal abatement after tipping, in a high temperature regime. For a

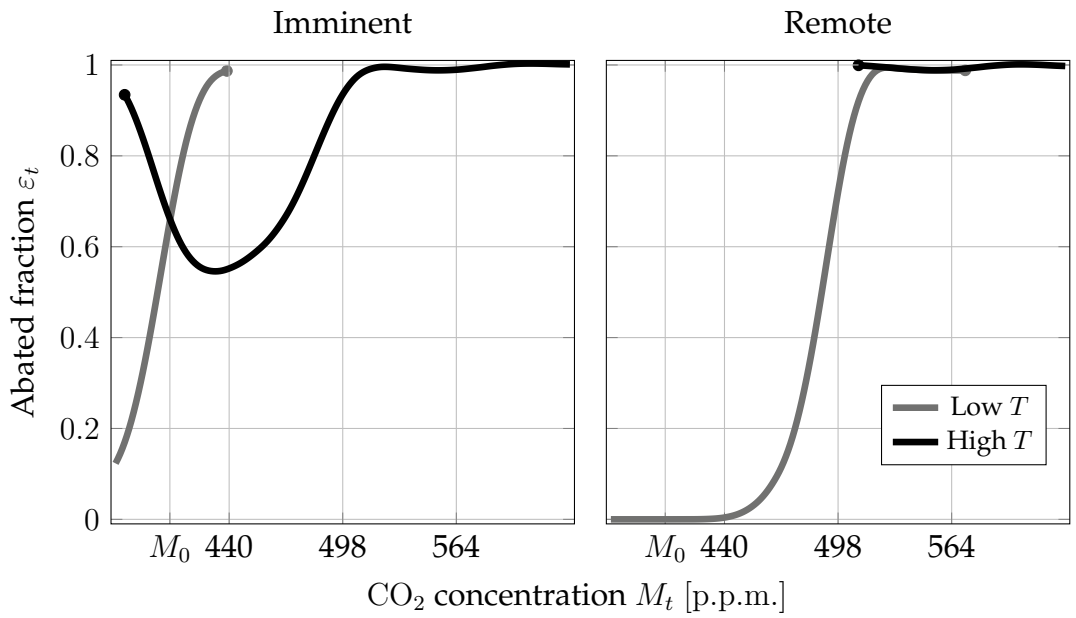

Figure 2.12: Fraction of abated emissions ε_t at different levels of CO_2 concentration M_t, for an imminent (left panel) and remote (right panel) point. To aid illustration, all other dimensions (T_t, N_t, Y_t) are set to their equilibrium value. As for some values of M_t there are two equilibria of temperature T_t, two curves are shown. The current level of CO_2 concentration is indicated with M_0.

given level of CO_2 concentration M_t, the optimal policy is not unique and switches in the event of tipping. In case the tipping point is imminent (left panel) but has not yet being crossed (lighter curve), the fraction of abated emission ε_t increases rapidly to 1 as CO_2 concentration rises. Once the tipping point is crossed abatement efforts are scaled back (darker curve) as the climate is irreversibly in a regime of high temperature. This is a strong policy result: abatement has a dual role, reducing first order climate damages and preventing the climate from tipping. Once the tipping point is crossed, the latter motive vanishes, and optimal abatement fails to undo the change to a higher temperature regime. This aspect of optimal abatement policies with tipping points is not captured by the social cost of carbon. Finally, notice that, if the climate has tipped, but CO_2 concentration is low, it is still worth abating a large fraction of emissions (dark curve is u-shaped). The incentive to remain at a low CO_2 concentration is that a random negative shock in temperature can push the system back to a low-temperature regime. Once the CO_2 concentration has increased, this becomes increasingly rare, and it is hence not worth pursuing large abatement measures on the chance it will happen. At the other extreme, if the tipping point is remote (right panel) it does not affect optimal policy. Abatement efforts are increasing in CO_2 concentration and full abatement is reached, in both regimes, at around 500 p.p.m.. In this case, direct climate damages are sufficiently large to warrant reaching net-zero before there is any risk of crossing the tipping point.

These policies, when implemented, yield a path of temperature and, as a consequence, temperature damages. All results are presented as median paths of each variable, which are computed as follows. First 100,000 paths are simulated under the temperature T_t (2.25), CO_2 concentration M_t (2.18), and output Y_t (2.39) dynamics with optimal controls α and χ. This also implies a path of abated fraction of emissions ε_t. Then at each moment t, the median value of the variable is computed. This sequence of median values for each t is the median path of the variable. Figure 2.13 shows the median path of the fraction ε_t of abated emission (solid line) and the resulting temperature path T_t with 95% simulation intervals (dotted lines), for an imminent (left panel) and a remote (right panel) tipping point.

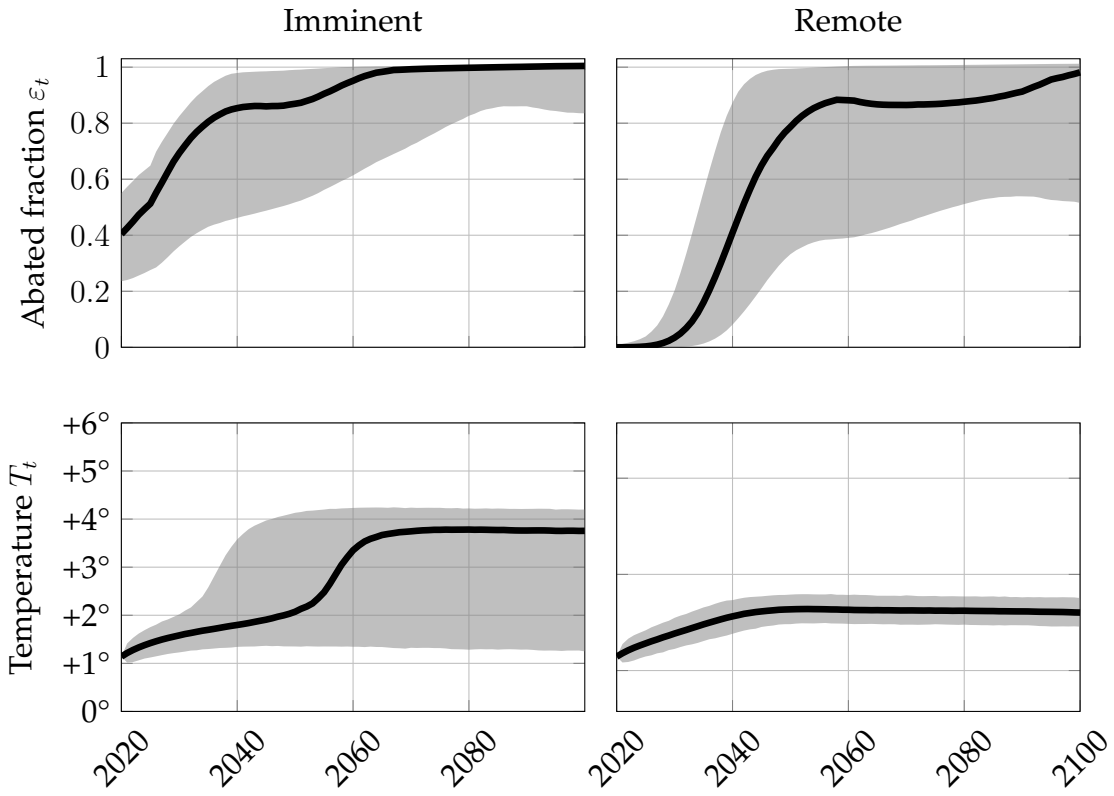

Figure 2.13: Median (solid) and 95 % simulation intervals (shaded area) abated fraction of emissions ε_t and temperature T_t at time t of 100,000 simulations in an imminent tipping point (left column) and a remote tipping point (right column). The simulations are generated under the temperature (2.25), CO_2 concentration (2.18), and output (2.39) dynamics with optimal controls α and χ.

When the tipping point is imminent, optimal abatement is promptly ramped up. In the median case 40 % of the no-policy emissions are abated immediately. Thereafter, abatement ramps up and net-zero is attained by 2060. These large efforts are not sufficient to completely prevent tipping as by 2040 the climate system has tipped with 5% probability and by 2060 with 50% probability. In case the tipping point is remote, the planner has more time to postpone abatement, which is cheaper in the future, and is able to stabilise temperature at around 2° without tipping. When

choosing abatement policies, the planners are balancing two sources of societal costs: abatement expenditures $\beta_t(\varepsilon_t)$, plus their adjustment costs, against climate damages $d(T_t)$. Figure 2.14 shows the average costs in each decade, as a fraction of output Y_t, in the median path, broken down into these three cost sources.

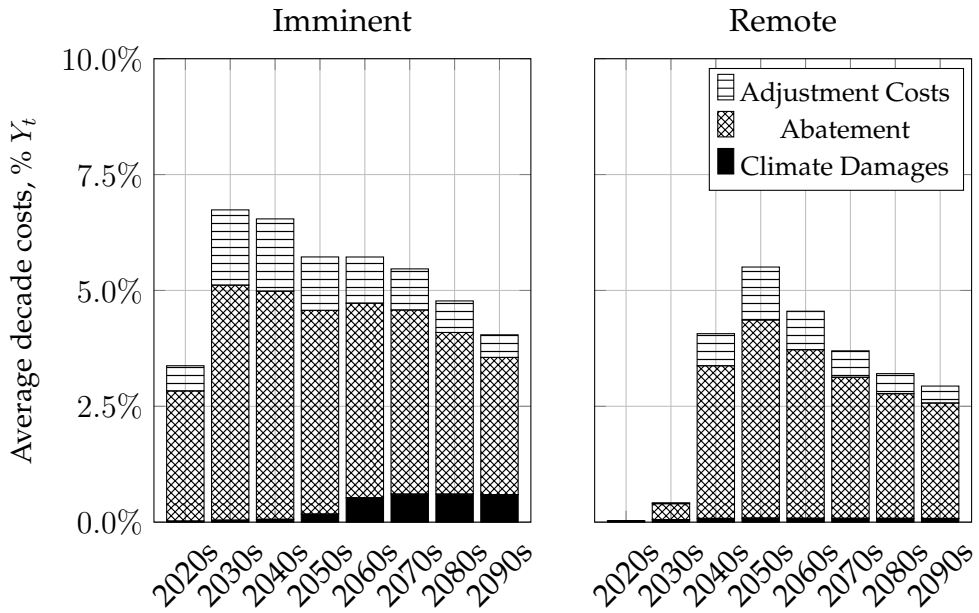

Figure 2.14: Costs, as a fraction of output Y_t, in the median path illustrated in Figure 2.13, broken down in adjustment costs, abatement and climate damages.

If the tipping point is imminent (left panel), optimal abatement costs quickly build up over the next two decades to around 5% of output before stabilising at 4% of output. Because abatement occurs quickly, society incurs large adjustment costs of up to 2% of output, which then fall over time due to technological progress ω_t. Despite these efforts, climate damages stabilise at around 1% of output. In case of a remote tipping point, the planner can scale abatement expenditures more slowly, thereby incurring in moderate adjustment costs and benefiting from improved technology. In both cases, the cost of the abatement required to stabilise the climate shrinks thanks to technological improvement ω_t (2.35), while climate damages are persistent as tipping points are irreversible.

These simulations provide a description of future costs, but to directly compare them it is necessary to take into account societal risk aversion θ, intertemporal elasticity of substitution ψ, and the discount rate ρ. To compute a cost of the tipping point that internalises these preferences and the uncertainty around output growth, first, I compute the net present value under optimal policies which satisfy the maximisation (2.40) at time $t = 0$. Denote this net present value by $\overline{V}_0$, when the tipping point is remote, and $\underline{V}_0$, when the tipping point is imminent. These two values can be translated into dollars by computing the corresponding certainty equiva-

lent. The certainty equivalent is the amount of output that society is willing to accept today to "shut down" climate change (see Appendix B.2 for more details). As climate change lowers the growth of output, the certainty equivalent is always lower than current output. The gap between the certainty equivalent and the current output is a measure of the cost of tipping points which internalises all future climate damages weighted by societal preferences. In case of a remote tipping point, the certainty equivalent $\overline{CE}$ is 70.845 trillion \$ per year (93.463% of current output), while in case of an imminent tipping point, the certainty equivalent $\underline{CE}$ is 65.575 trillion \$ per year (86.51% of current output). Note that, both certainty equivalents are significantly lower than current output Y_0: the cost of a tipping point can be as large as 13.5% of current output. It is important to notice that these values are derived by looking at two extreme scenarios on tipping points: an imminent and a remote one. The true certainty equivalent will hence lie somewhere in between these two quantities. The cost of the temperature feedback being triggered imminently, at $T^c = 1.5°$, rather than remotely, $T^c = 2.5°$, is $\overline{CE} - \underline{CE} = 15.27$ trillion \$ per year or 6.953% of current output. This quantity measures how much society should be willing to pay to push the tipping point back by $1°$.

2.5 The Cost of Tipping Point Uncertainty

The previous section discussed the optimal abatement policies of a planner that knows whether the tipping point is imminent or remote. Yet, tipping points are often hard to predict. This section gives an upper bound on the cost of this unpredictability. To do so, I consider two extreme scenarios. First, I consider a *wishful thinker* planner, denoted by w, who erroneously assumes the tipping point to be remote, while it is in fact imminent. Once the temperature feedback comes into effect, the planner detects the tipping point and switches to the optimal abatement strategy. In other words, letting

$$\tau := \inf_t \{T_t \geq T^c\} \tag{2.42}$$

be the time at which the feedback effects comes into effect, the abatement strategy employed by w is given by

$$\alpha_t^w := \begin{cases} \overline{\alpha}_t & \text{if } t < \tau, \\ \underline{\alpha}_t & \text{if } t \geq \tau. \end{cases} \tag{2.43}$$

Second, I consider a *cautious* planner, denoted c, who erroneously assumes the tipping point to be imminent, while it is in fact remote. Again, in case the feedback

loop is triggered, the planner switches to the optimal strategy, such that

$$\alpha_t^c := \begin{cases} \underline{\alpha}_t & \text{if } t < \tau, \\ \overline{\alpha}_t & \text{if } t \geq \tau. \end{cases} \tag{2.44}$$

Table 2.1 summarises the possible tipping points and the strategies employed before the temperature feedback is triggered, $t < \tau$.

	Tipping Point	
	Imminent	Remote
$\underline{\alpha}$	Optimal	Cautious c
$\overline{\alpha}$	Wishful thinker w	Optimal

Table 2.1: Strategies used before tipping $t < \tau$ for the four scenarios. After tipping $t \geq \tau$ optimal strategies are used in all scenarios.

The wishful thinker planner w incurs the largest climate damages. The cautious planner c incurs the largest abatement and adjustment costs. As society can always act cautiously, the difference in the costs incurred by the wishful thinker and the cautious planner is an upper bound on the cost of the uncertainty around the tipping point.

Focusing first on the wishful thinker w, Figure 2.15 shows the median path of the abated fraction of emissions ε_t (left panel) and the resulting temperature T_t (right panel) if α^w (2.43) is used in face of an imminent tipping point.

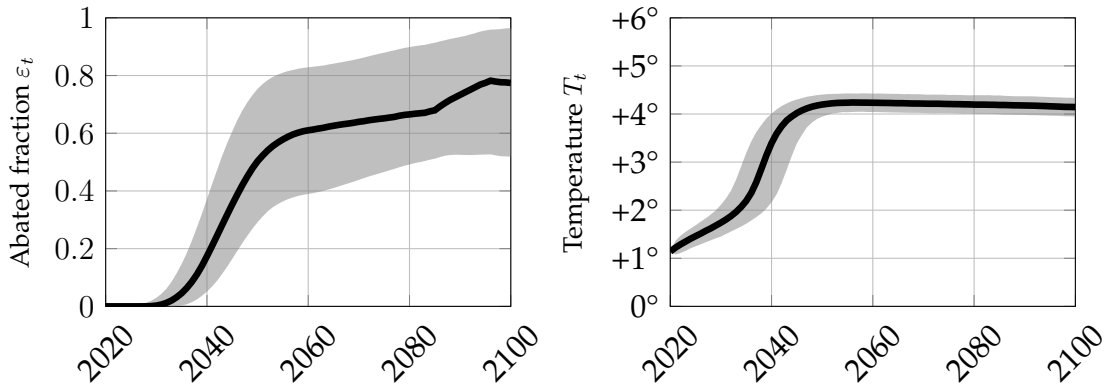

Figure 2.15: Median (solid) and 95% simulation intervals (shaded) of the abated fraction of emissions ε_t (left panel) and the resulting the temperature T_t (right panel). Calculated from 100,000 simulations using policy (2.43).

In the first decade, the planner erroneously believes the tipping point to be remote, and hence postpones abatement efforts, as the median $\varepsilon_t \approx 0$ for $t < 2030$. Once the temperature feedback comes into effect, the planner "slams the brake" and quickly ramps up abatement, achieving net-zero in the 2040s, in most cases. Despite the

quick abatement ramp-up, by 2050 the climate has tipped to a high temperature regime with 95% probability. When compared to the optimal abatement path under an imminent tipping point (left column of Figure 2.13), the abatement followed by w clearly results in much larger climate damages, as the climate tips sooner and with large probability, and in larger adjustment costs, as abatement measures need to be ramped up more quickly. Figure 2.16 shows a breakdown of the cost incurred by the wishful thinker w.

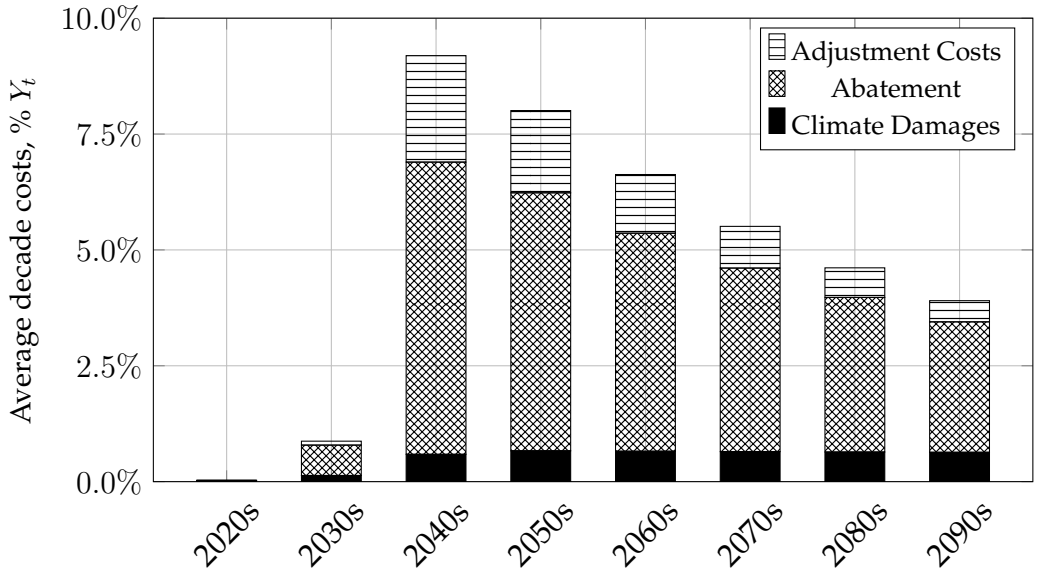

Figure 2.16: Costs in the median scenario of a wishful thinker w planner as a fraction of output Y_t, broken down in adjustment costs, abatement and climate damages.

In the first two decades, the planner implements little abatement. Yet, in the 2040s as the temperature feedback comes into effect, abatement expenditures ramp up quickly. Rapidly ramping up abatement incurs large adjustment costs, which peak at more than 3% of output. This delay in abatement is not sufficient to prevent tipping, which brings about large and, most importantly irreversible, climate damages of around 1% of output. This scenario gives an upper bound on the possible climate damages society can incur if it delays optimal abatement until after the feedback comes into effect. Due to technological progress in abatement technology ρ_ω (2.35), abatement expenditures shrink exogenously. On the contrary, climate damages persist.

I then do an analogous experiment for the cautious planner c. Figure 2.17 displays the path of the abated fraction of emissions ε_t (left panel) and temperature T_t (right panel) resulting from abatement α^c (2.44) whenever the tipping point is remote.

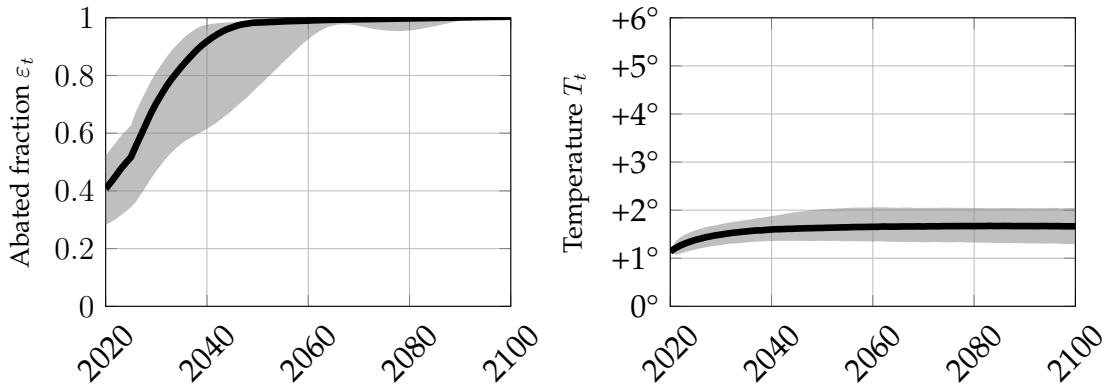

Figure 2.17: Median and 95% simulation intervals of the abated fraction of emissions ε_t (left panel) and the resulting the temperature T_t (right panel). Calculated from 100,000 simulations using policy (2.44).

The cautious planner fully abates emissions already by the early 2040s and quickly stabilises temperature at around $1.9°$. Figure 2.18 breaks down the cost incurred by the cautious c planner.

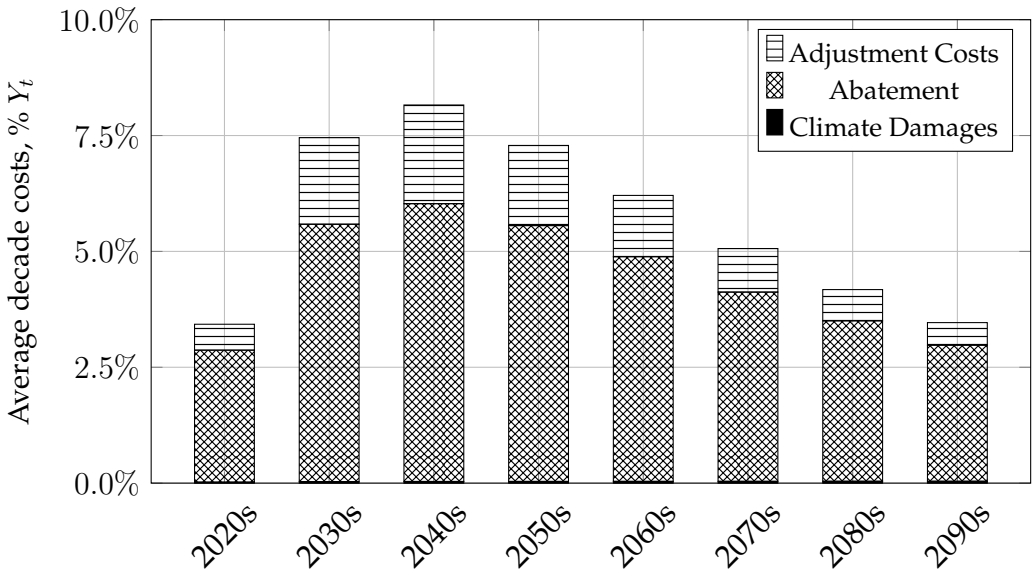

Figure 2.18: Costs in the median scenario of a cautious c planner as a fraction of output Y_t, broken down in adjustment costs, abatement and climate damages.

The cautious planner sustains abatement costs of around or more than 6% of output between 2030 and 2070. This fast and aggressive abatement schedule accumulates large adjustment costs that can be as large as 2% of output. In return, the climate damages are kept below 0.1% of output. This scenario gives an upper bound on the abatement costs society can incur if it abates quickly assuming a tipping point is imminent.

The breakdowns presented so far (Figures 2.16 and 2.18) give upper bounds on the climate damages (w) and abatement cost (c). Yet, as above, these are just nominal

costs that do not internalise societal preference. Hence, as in the optimal scenario presented in the previous section, I compute certainty equivalents CE^w and CE^c, in $\text{trillionUS\$/year}$, for the wishful thinker and cautious case respectively (more detail on how they are computed is given in Appendix B.2). These are summarised in Table 2.2.

	Tipping Point	
	Imminent	Remote
$\underline{\alpha}$	$\underline{CE} = 65.57\text{trUS\$/y}\ (86.51\%Y_0)$	$CE^c = 65.58\text{trUS\$/y}\ (86.513\%Y_0)$
$\overline{\alpha}$	$CE^w = 58.52\text{trUS\$/y}\ (77.2\%Y_0)$	$\overline{CE} = 70.85\text{trUS\$/y}\ (83.46\%Y_0)$

Table 2.2: Certainty equivalents for the four scenarios in$\text{trillionUS\$/year}$. In brackets, the certainty equivalent is expressed as a percentage of current output Y_0.

Table 2.2 has the same structure as Table 2.1. Each cell contains the certainty equivalent of the corresponding scenario, in trillion\$ per year and as a percentage of current output Y_0. These quantities allow to bound the cost of the tipping point. On the one hand, a cautious strategy c incurs in suboptimal abatement and adjustment costs. These are $\overline{CE} - CE^c = 5.27$ trillion\$ per year or 6.95% of current output. On the other hand, a wishful thinker strategy w incurs large and permanent climate damages. This can have costs up to $\underline{CE} - CE^w = 7.06$ trillion\$ per year or 9.31% of current output. Furthermore, albeit costly, a cautious strategy c can be up to $(\overline{CE} - CE^c) - (\underline{CE} - CE^w) = 1.79$ trillion\$ per year or 2.36% of current output cheaper than a wishful thinker strategy w. As caution is always possible, this quantity is also a measure of how costly uncertainty around climate tipping points can be. The robustness of this figure to different relative risk aversion and critical threshold parameters is explored in Appendix B.5.

2.6 Discussion

This chapter estimates the economic cost of uncertain and irreversible climate tipping points. Temperature feedback is integrated in a state-of-the-art integrated assessment model, which is then calibrated. To correctly represent the current uncertainty on the tipping point, I consider two limit scenarios consistent with the climate literature: that of an *imminent* and a *remote* tipping point. Using these two scenarios, I estimate an upper bound on the cost of tipping points and then compute the costs of a cautious strategy in the face of uncertainty, compared to a wishful thinking strategy, postponing abatement until the tipping point is identified.

The approach employed here does not capture the full picture. There is no single tipping point; multiple climate tipping points exist, and they interact. Furthermore,

I only provide upper bounds on the costs tipping. Nevertheless, the model studied in this chapter patches one of the many "inconsistencies between how leading economic models of climate change represent climate dynamics" (Dietz et al., 2021*a*, p. 3) by introducing a realistic tipping point and studying its economic consequences.

Ignoring irreversible climate tipping points could result in costs as high as 13.5% of current world output. This figure demonstrates how feedback mechanisms in the climate system cannot be ignored when considering optimal abatement policies: by abating, we are not only reducing future climate damages directly but also lowering the likelihood of crossing a tipping point. Measures of marginal benefits and costs, such as the social cost of carbon, primarily internalise direct future climate damages and may mislead policymakers. Finally, I show that, in the face of uncertainty, a precautionary approach, albeit costly, can be 2.35% of output cheaper than postponing abatement. The intuition is straightforward: it is better to pay for abatement, which we know is becoming more affordable, than to gamble with the risk of crossing a tipping point.

Chapter 3

Optimal Emission Abatements with Regional Tipping Points

Climate tipping points are a significant concern when evaluating emission abatement policies. Some climate tipping points have different effects on different regions of the world, creating different incentives to mitigate emissions. In this chapter, I study how a regional tipping point can affect optimal emission abatement in two regions. I develop, calibrate, and solve a model in which two regions, the OECD and the rest of the world (RoW), choose emission abatement efforts to mitigate climate change. Countries experience economic damage due to increasing local temperatures. The temperature in the RoW is subject to a tipping point, which, if crossed, causes local temperatures to surge to a new, hotter regime. I show that an imminent tipping point in the RoW allows the OECD to reduce abatement efforts, forcing the RoW to take on a larger fraction of the emission abatement necessary to stabilise temperature. This equilibrium generates a large welfare loss at a global level, as any dollar spent on emission abatement in the RoW, as opposed to OECD countries, has a much higher opportunity cost, given that the RoW output growth is faster than that of the OECD. Furthermore, I show that a large part of the loss in output growth in the RoW is due to the laxer abatement policy from the OECD rather than from the tipping points itself. This suggests that in a non-cooperative world, in which the OECD strategically chooses abatement policies, a climate tipping point can represent a hurdle to emission abatement rather than an incentive to mitigate emissions at a global scale.

The rest of the chapter is structured as follows. Section 3.1 discusses the climate and economic literature on regional tipping points, extending the review provided in Section 2.1. Section 3.2 introduces the model in three parts. First, Section 3.2.1 extends the climate model presented in Section 2.3.1 to an environment with two regions and illustrates the consequences of a regional tipping point under a no-policy

scenario. Second, Section 3.2.2 extends the economy from Section 2.3.2 to an environment with two regions. Thirdly, Section 3.2.3 introduces the "emission abatement game", that is, the objective of the two regions and their possible strategies. Section 3.3 presents the optimal abatement policies in the emission abatement game and attributes the costs incurred by the RoW to either the climate or the adversarial behaviour of the OECD. Finally, Section 3.4 discusses the results and draws policy conclusions.

3.1 Related Literature

Environmental regime shifts, including climate tipping points, have been the subject of growing interest in economics, as shown in Section 2.1. Naturally, many researchers have extended the study of environmental regime shifts to an environment with multiple agents behaving strategically (Heijnen and Wagener, 2013; Wagener and De Zeeuw, 2021, among others). In parallel, a large body of literature has focused on introducing multiple emitters in the study of transboundary pollution and climate change, both conceptually (Dockner and Van Long, 1993) and within integrated assessment modelling (Nordhaus and Yang, 1996). Building on this literature Nkuiya (2015) considers a transboundary pollution game with possible regime shift in damages. They show that multiple equilibria might realise, in which the threat of a regime shift can either raise or lower emissions. Yet, they consider a regime shift which, if triggered, leads instantaneously to higher damages and affects all regions' uniformly. In this chapter, I take a complementary approach. I relax this assumption and bring the model closer to climate dynamics. The climate tipping points, and subsequent regime shift, are endogenous and calibrated. Furthermore, the regime shift has a differential impact on different regions. The price to pay for the added complexity is that I only consider a single equilibrium.

Moreover, this chapter contributes to the debate around the differential regional consequences of climate change. First, optimal regional climate policies are expected to differ in a world with idiosyncratic regional impacts of climate change (Hillebrand and Hillebrand, 2019). Second, regions that differ in their future growth and fiscal sustainability have different opportunity costs of emission abatement measures (D'Autume, Schubert and Withagen, 2016). In this chapter, I show that climate tipping points can act as a further wedge between a social planner abatement policy and optimal regional climate policies.

3.2 Model

The following extends the climate and economy model from the previous chapter, as outlined in Section 2.3, to an environment with multiple emitters.

3.2.1 Climate

Carbon sinks and CO_2 concentration

Economic activity in region $i \in \{\mathrm{OECD}, \mathrm{RoW}\}$ generates $\mathrm{Gt\,CO}_2$ emissions $E_{i,t}$. These increase the atmospheric CO_2 concentration M_t $\mathrm{p.p.m.}$. As in the previous chapter's Section 2.3.1, a fraction $\delta_m(N_t)$ of the atmospheric concentration of CO_2 M_t decays into natural sinks and the rate of decay $\delta_m(N_t)$ is a function of the $\mathrm{Gt\,CO}_2$ stored in natural sinks N_t.

In absence of abatement efforts, the atmospheric CO_2 concentration M_t evolves according to (2.16), where the growth rate γ_t^{np} is calibrated to match the *no-policy* SSP5 scenario (Kriegler et al., 2017). The global growth rate γ_t^{np} is the sum of the growth rate of CO_2 concentration in the two regions, namely,

$$\gamma_t^{\mathrm{np}} = \gamma_{\mathrm{OECD},t}^{\mathrm{np}} + \gamma_{\mathrm{RoW},t}^{\mathrm{np}}, \tag{3.1}$$

where

$$\gamma_{i,t}^{\mathrm{np}} \coloneqq \xi_m \frac{E_{i,t}^{\mathrm{np}}}{M_t^{\mathrm{np}}} - \frac{1}{2}\delta_m(N_t^{\mathrm{np}}) \tag{3.2}$$

for $i \in \{\mathrm{OECD}, \mathrm{RoW}\}$. Using the regional projected emissions under SSP5, I compute the implied regional growth rates of atmospheric CO_2 concentration $\gamma_{i,t}^{\mathrm{np}}$. Figure 3.1 shows the cumulative calibrated growth rate $\int \gamma_t^b \ \mathrm{d}t$ of CO_2 concentration in the no-policy scenario in each decade until the end of the century. The growth rates are decomposed into those driven by emissions in the OECD and those in the RoW. Most of the growth in CO_2 concentration in this scenario is driven by emissions in the RoW.

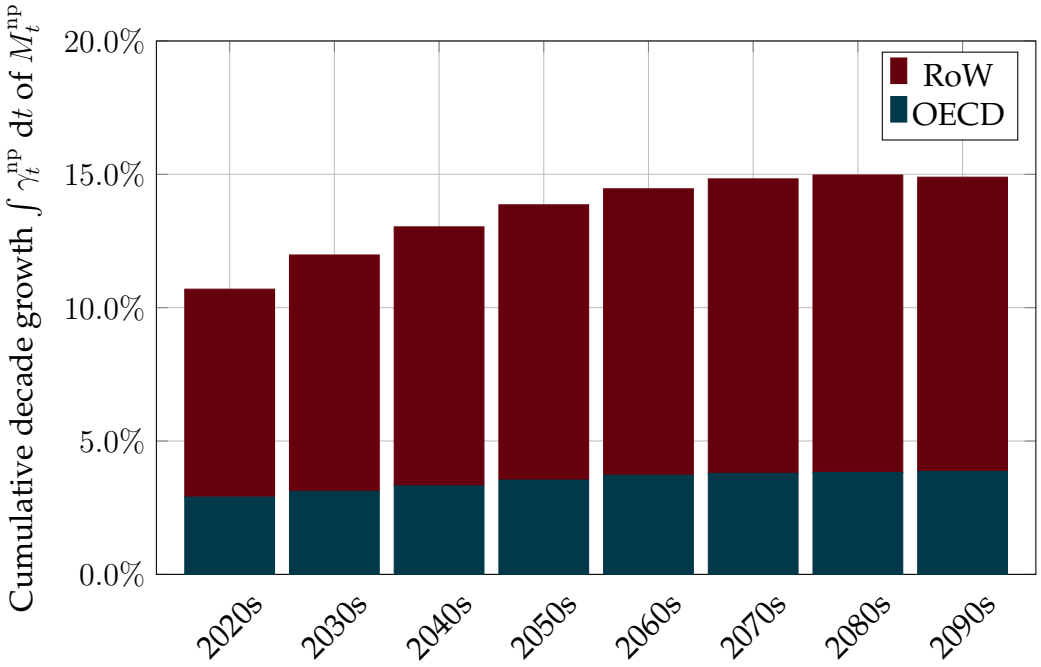

Figure 3.1: Cumulative growth rate of CO_2 concentration in the SSP5 no-policy scenario for each decade attributed to the OECD and the RoW.

The resulting CO_2 concentration in the no-policy scenario is the same as in Chapter 2, as displayed in Figure 2.5.

Regional Temperature

Some tipping points, such as the Atlantic meridional overturning circulation and the Boreal permafrost, affect regional temperatures (Seaver Wang et al., 2023). Furthermore, such effects are strong in poorer regions (Dietz et al., 2021*b*). In this chapter, I capture this asymmetry by extending the climate model introduced in the previous chapter. To keep things tractable, I assume the only interaction channel between temperature in the two regions is the CO_2 concentration. To do so, I assume average regional temperature $T_{i,t}$ in the two regions $i \in \{\text{OECD}, \text{RoW}\}$ is determined by a zero dimensional energy balance model analogous to that of Section 2.3.1 (McGuffie and Henderson-Sellers, 2014, p.85). The only difference between the two regions is the presence of positive feedback in the temperature dynamics in the RoW, namely

$$\epsilon \, \mathrm{d}T_{\text{OECD},t} = \bar{r}(T_{\text{OECD},t}) \, \mathrm{d}t + g(M_t) \, \mathrm{d}t + \sigma_T \; \mathrm{d}W_{T,\text{OECD},t} \,, \tag{3.3}$$

$$\epsilon \, \mathrm{d}T_{\text{RoW},t} = r(T_{\text{RoW},t}) \, \mathrm{d}t + g(M_t) \, \mathrm{d}t + \sigma_T \; \mathrm{d}W_{T,\text{RoW},t} \tag{3.4}$$

where the radiative forcing in the OECD region has no positive feedback

$$\bar{r}(T_{\text{OECD},T,t}) \coloneqq S_0 \left(1 - \bar{\lambda}\right) - \eta\sigma T_t^4, \tag{3.5}$$

while in the RoW it does

$$r(T_{\text{RoW},t}) \coloneqq S_0 \left(1 - \lambda(T_{\text{RoW},t})\right) - \eta\sigma T_t^4. \tag{3.6}$$

Assuming temperature in the OECD region is subject to weak positive feedback, rather than no feedback, would not change our results qualitatively. The regional model is perfectly analogous to the world average model (2.25) illustrated in the previous chapter and retains the same calibration. For a more detailed discussion see Section 2.3.1. The critical threshold parameter T^c, that is, the temperature at which the positive feedback is triggered, is not calibrated, as there is large uncertainty around the location of tipping points (Ben-Yami et al., 2024; Ditlevsen and Johnsen, 2010; Wagener, 2013). Here, I will consider three scenarios consistent with the climate literature (Seaver Wang et al., 2023), that is, $T^c \in \{1.8^\circ, 2^\circ, 2.5^\circ\}$, as well as a benchmark scenario in which the temperature in the RoW is not subject to any feedback mechanism, referred to henceforth as "No Tipping" scenario. The calibrated feedback function λ in the former three scenarios is shown in Figure 3.2.

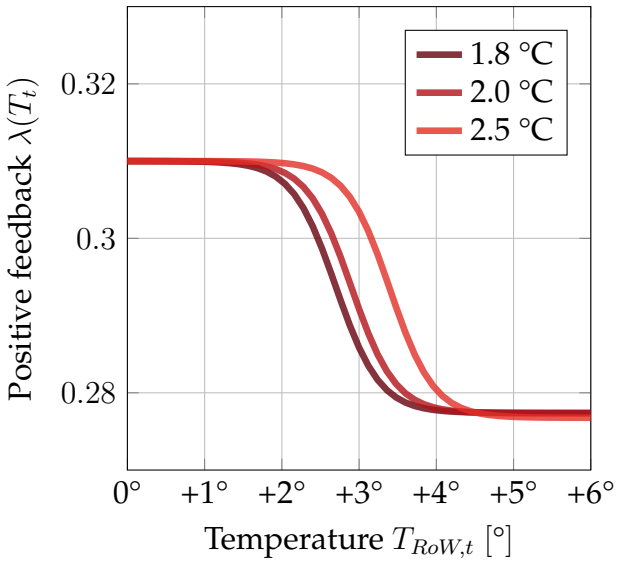

Figure 3.2: Calibrated feedback in the RoW for the three scenarios.

Figure 3.3 shows the temperature paths, for the four scenarios and the two regions, resulting from the temperature dynamics (3.3), under the assumption of a no-policy CO_2 concentration path M_t^{np} (2.16). In case of no tipping element in the RoW (bottom right panel), the temperature evolution is identical in both regions. As the feedback is introduced and the critical threshold T^c is lowered, the temperature

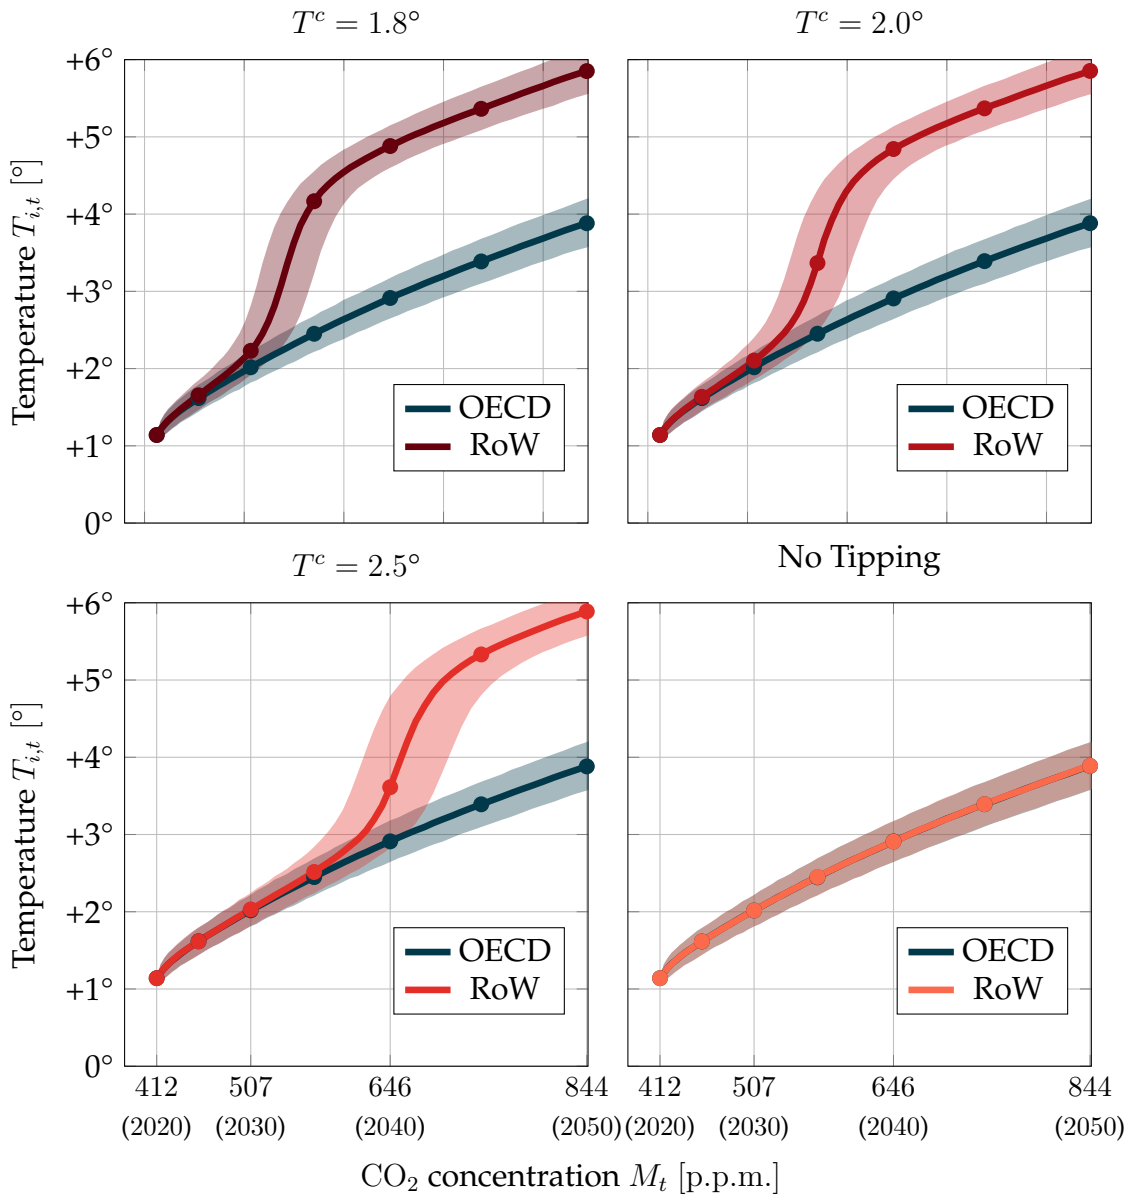

Figure 3.3: Path of temperature $T_{i,t}$ in the two regions for the four possible RoW climates under assumption of a no-policy CO_2 concentration path M_t^{np}. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

path of the RoW diverges sooner with high probability from that of the OECD.

3.2.2 Economy

This section introduces the economy of the two regions. As in the previous chapter (Section 2.3.2), these are assumed to be Harrod-Domar economies. Output in both regions is given by $Y_{i,t} = A_{i,t}K_{i,t}$ where productivity $A_{i,t}$ grows at a constant region specific rate ϱ_i. Each region can allocate output to investment in capital $I_{i,t}$, emission abatement expenditures $B_{i,t}$, or consumption $C_{i,t}$. As in the previous chapter, abatement $B_{i,t}$ and investment are subject to quadratic adjustment costs (2.30),

with common marginal investment friction κ. Climate change affects the economy by lowering capital growth via the damage function $d(T_{i,t})$ as illustrated in Figure 2.10. Here, I assume that while the temperature dynamics differ between the two regions, the damages from climate change $d(T)$ are the same at a given temperature level T. This simplifying assumption allows to isolate the effect of the regional tipping point in determining abatement strategies. Finally, I assume the abatement costs $\beta_{i,t} \coloneqq B_{i,t}/Y_{i,t}$ to be quadratic in the fraction of abated emissions $\varepsilon(\alpha_{i,t})$ (Baker, Clarke and Shittu, 2008; Dietz and Venmans, 2019; Nordhaus, 1992, 2017), such that,

$$\beta_{i,t}(\varepsilon(\alpha_{i,t})) = \frac{\omega_t}{2}\,\varepsilon(\alpha_{i,t})^2, \tag{3.7}$$

where the joint marginal technology $\omega_t \coloneqq \omega_0 e^{-\rho_\omega t}$ is exogenously decreasing over time. The marginal abatement curves are displayed in 2.11. Putting all these components together, the regional growth of capital, analogous to (2.37), is given by

$$\frac{\mathrm{d}K_{i,t}}{K_{i,t}} = \Big(\underbrace{\phi_{i,t}(\chi_{i,t})}_{\text{growth}} - \overbrace{A_{i,t}\beta_{i,t}(\varepsilon(\alpha_{i,t}))}^{\text{abatement}} - \underbrace{d(T_{i,t})}_{\text{climate}} \Big)\,\mathrm{d}t + \sigma_k\,\mathrm{d}W_{k,t}\,, \tag{3.8}$$

where

$$\phi_{i,t}(\chi_{i,t}) \coloneqq A_{i,t}(1-\chi_{i,t}) - \frac{\kappa}{2}A_{i,t}^2(1-\chi_{i,t})^2 - \delta_k \tag{3.9}$$

is an endogenous growth component and $\chi_{i,t} \coloneqq C_{i,t}/Y_{i,t}$ is the fraction of output allocated to consumption. This yields output dynamics

$$\frac{\mathrm{d}Y_{i,t}}{Y_{i,t}} = \varrho_{i,t}\,\mathrm{d}t + \frac{\mathrm{d}K_{i,t}}{K_{i,t}}. \tag{3.10}$$

The outcome of the regional calibration is shown in Appendix B.3. The crucial difference in the calibration of the two regions is that the OECD starts off as richer than the RoW $Y_{\mathrm{OECD},0} > Y_{\mathrm{RoW},0}$, yet, the growth of productivity is faster in the RoW $\varrho_{\mathrm{OECD},0} < \varrho_{\mathrm{RoW},0}$. This implies that abatement expenditures in the RoW, as opposed to the OECD, have the same outcome on emissions, but they have larger opportunity cost, as they cause larger output losses in the future. This mechanism drives welfare losses when the two regions choose their abatement efforts independently.

3.2.3 Emission Abatement Game

Given the economic dynamics introduced in Section 3.2.2, this section sets up the emission abatement game. The two regions' utility is given by the recursive utility

function

$$V_{i,t} = \sup_{\alpha_{i,t},\chi_{i,t}} \mathbb{E} \int_t^\infty f(C_{i,s}, V_{i,s})\, \mathrm{d}s\,, \tag{3.11}$$

where f is a common Epstein-Zin aggregator (2.41), subject to the dynamics of N_t, M_t, $T_{i,t}$ and $Y_{i,t}$ introduced above. I assume players to use elementary feedback strategies as defined in Sîrbu (2014). Let

$$X_t \coloneqq (N_t, M_t, T_{\mathrm{OECD},t}, T_{\mathrm{RoW},t}, Y_{\mathrm{OECD},t}, Y_{\mathrm{RoW},t}) \tag{3.12}$$

be the state vector at time t. Players formulate piecewise constant strategies $\bar{\alpha}_{i,t_{i,k}}$ and $\bar{\chi}_{i,t_{i,k}}$ at discrete times $t \in \{t_{i,0}, t_{i,1}, t_{i,2} \ldots\}$ having observed the evolution of the state vector X_s for $s \leq t$, such that,

$$\alpha_{i,t} = \bar{\alpha}_{i,t_{i,k}} \text{ for } t_{i,k} \leq t < t_{i,k+1} \text{ and} \tag{3.13}$$

$$\chi_{i,t} = \bar{\chi}_{i,t_{i,k}} \text{ for } t_{i,k} \leq t < t_{i,k+1}. \tag{3.14}$$

Intuitively, regions formulate abatement and consumption strategies which depend on time and the state path and take the other region's policies as fixed. I then solve the game by means of a discrete Markov chain approximation, as outlined in employ the solution mechanism introduced by Kushner (2007). See Appendix B.1.3 for further details.

3.3 Results

This section illustrates the results of the emission abatement game introduced in Section 3.2.3.

3.3.1 Optimal Abatement Policies

First, I discuss optimal abatement efforts undertaken by the two regions and their consequences for the climate and the respective economies.

Figure 3.4 shows the fraction of abated emissions $\varepsilon_i(\alpha_{i,t})$ in each region. In absence of tipping points, OECD countries start to abate around 2035, steadily increasing the amount of emission abatement through 2080. The RoW only chooses to abate emissions starting from 2040. This discrepancy is due the difference in opportunity cost of emission abatement in the two regions. As the RoW has a lower initial level of output Y_0 and a higher rate of growth of productivity ϱ_{RoW}, every dollar spent in emission abatement in the RoW generates larger output losses in the future. If a tipping point is introduced, at a critical threshold $T^c = 2.5°$ the

OECD undertakes a less ambitious emission abatement policy which results in a more ambitious and prompt emission abatement in the RoW. Lowering the tipping point threshold to $T^c = 2°$ exacerbates this dynamic between the two regions. If the tipping point is further lowered at $T^c = 1.8°$ the RoW moves emission abatement even earlier in time, but the long run scope of the emission abatement is lower. This occurs because an excessively low critical threshold T^c makes it impossible for the RoW to prevent tipping. Hence, once the climate has tipped, the additional incentive of emission abatement to prevent tipping are lost.

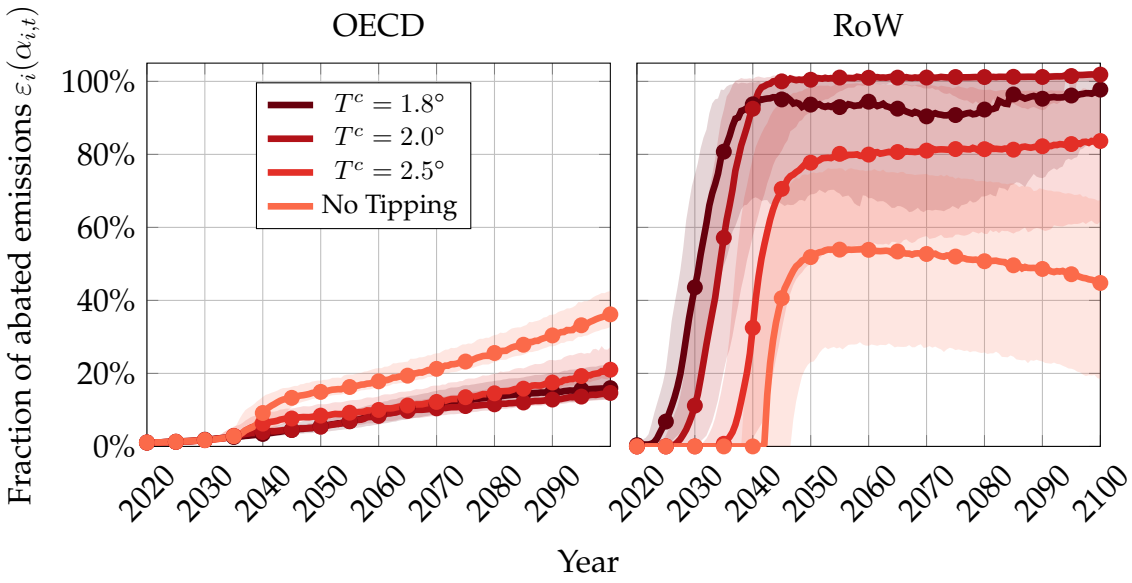

Figure 3.4: Fraction of abated emissions $\varepsilon_i(\alpha_{i,t})$ in the two regions (OECD and RoW) in the scenario without tipping and the three critical threshold scenarios $T^c \in \{1.8^\circ, 2^\circ, 2.5^\circ\}$. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

Figure 3.5 shows the growth of CO_2 concentration $\gamma_t^b - \alpha_{1,t} - \alpha_{2,t}$, given the abatement policies of the two regions. In all scenarios, the growth rate of CO_2 concentration initially grows before falling to net-zero. In absence of tipping points, it peaks around 2040 and net-zero is reached by 2050. If a tipping point affecting the RoW is introduced at a critical threshold $T^c = 2.5^\circ$, the more ambitious abatement schedule in the RoW makes emission fall earlier, yet net-zero is still reached only by 2050. If the tipping point occurs at a lower critical threshold $T^c = 2^\circ$ then emission reduction starts early, driven by the RoW abatement efforts, but net-zero is reached much later due to the OECD procrastination. If the tipping point is imminent, at critical threshold $T^c = 1.8^\circ$, then the reduced post-tipping incentive to abate by the RoW further postpones net-zero.

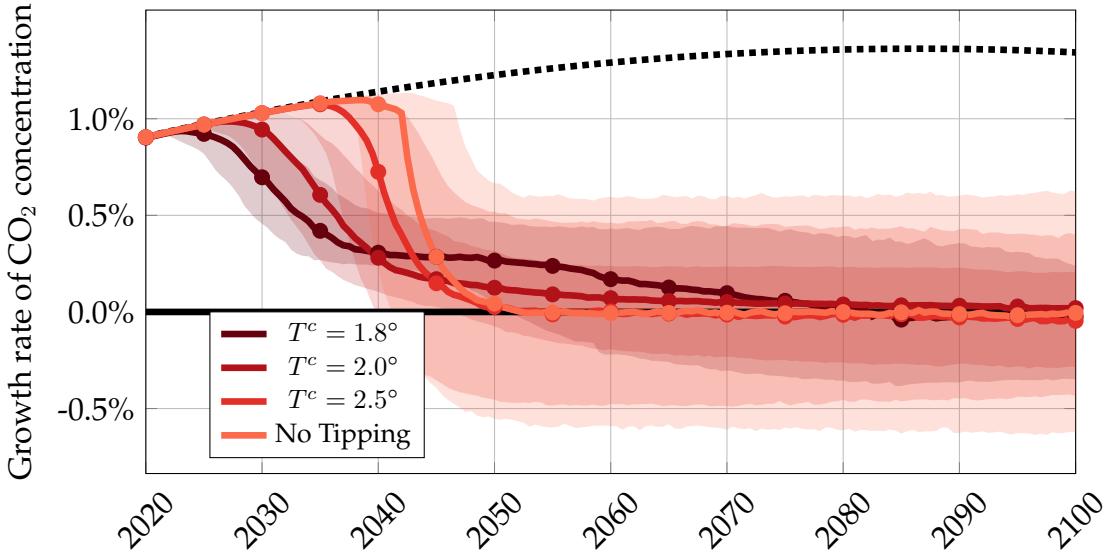

Figure 3.5: Growth rate of CO_2 concentration $\gamma_t^b - \alpha_{1,t} - \alpha_{2,t}$ in the scenario without tipping and the three critical threshold scenarios $T^c \in \{1.8^\circ, 2^\circ, 2.5^\circ\}$. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

These scenarios for CO_2 concentration growth lead to the temperature dynamics illustrated in Figure 3.6. In a climate without a tipping point, with a remote critical threshold at $T^c = 2.5^\circ$, or with a closer critical threshold at $T^c = 2^\circ$, the temperature T_t stabilises around 2°. Yet, as illustrated above, the burden of stabilising the climate is moved towards the RoW as the critical threshold is lowered from $T^c = 2.5^\circ$ to $T^c = 2^\circ$. In the latter case, the stabilisation is close to a critical threshold, hence there is a sizable probability that the climate in the RoW tips before the end of the century. If the critical threshold is imminent, at $T^c = 1.8^\circ$ the rest of the world fails to prevent a tipping point and the climate tips in more than 90% of the cases.

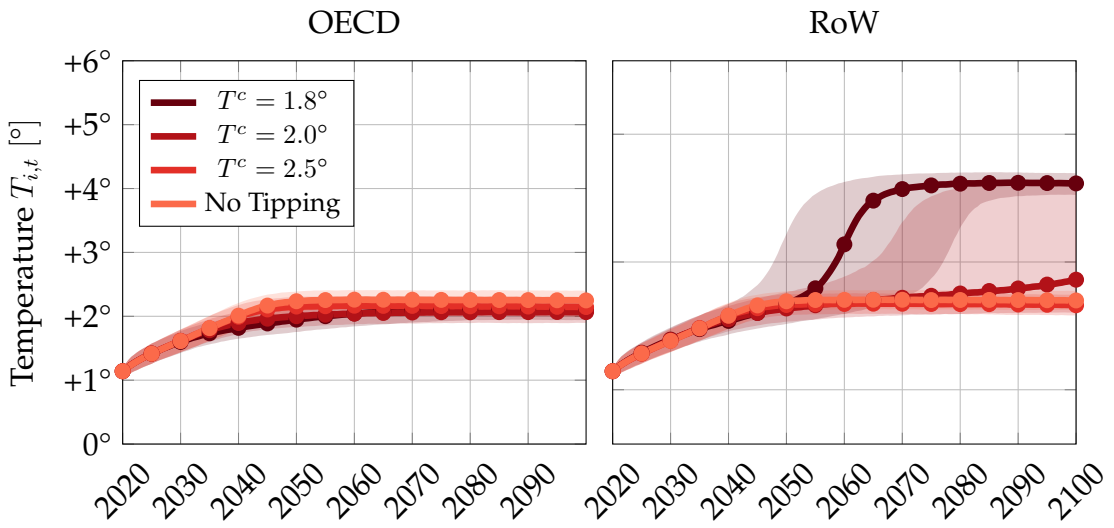

Figure 3.6: Path of temperature $T_{i,t}$ in the scenario without tipping and the three critical threshold scenarios $T^c \in \{1.8^\circ, 2^\circ, 2.5^\circ\}$. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

3.3.2 Cost of an Adversarial OECD for the RoW

This section discusses the economic impacts of the different scenarios on the output of the RoW. Figure 3.7 illustrates the path of output $Y_{\text{RoW},t}$ under the four climate scenario. As expected, the introduction of a tipping element in the RoW's lowers the output by the end of the century. This effect is exacerbated as the critical threshold is lowered. In the limit case of an imminent threshold $T^c = 1.8^\circ$, the presence of a tipping point costs the RoW around 40 trillion$ per year by the end of the century.

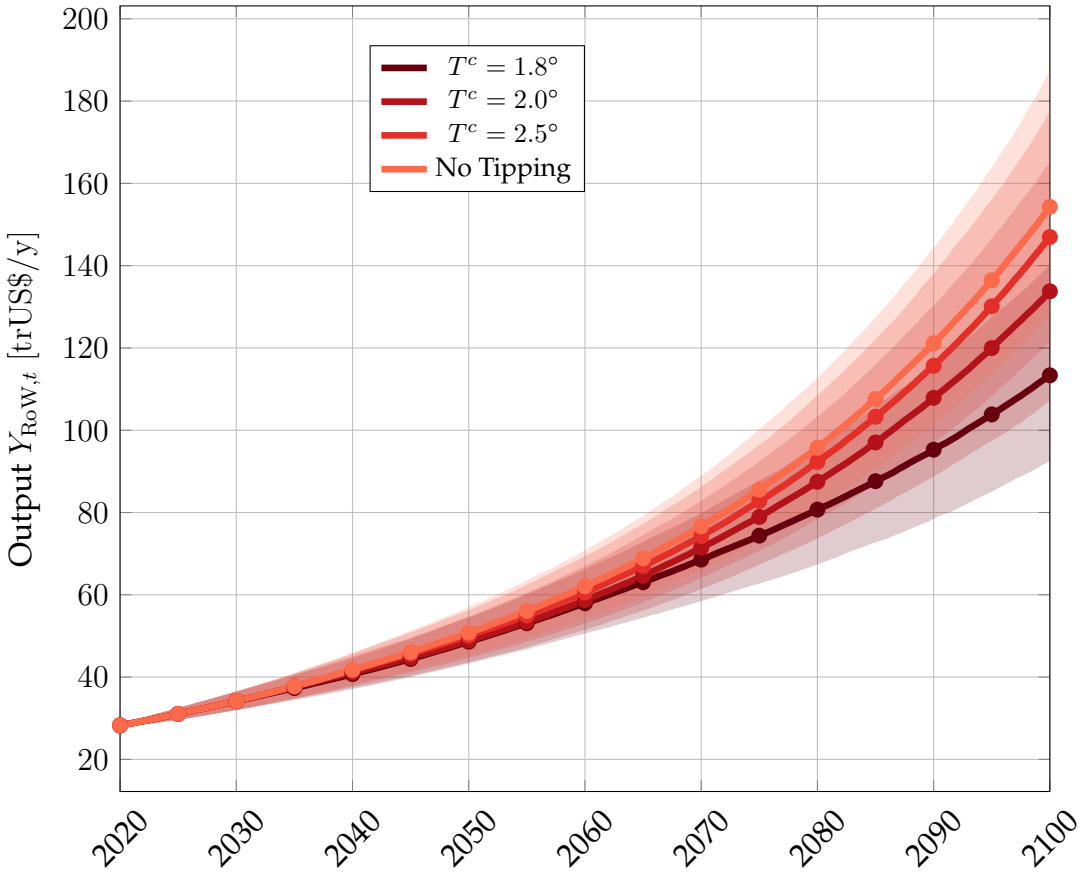

Figure 3.7: Path of RoW's output $Y_{\mathrm{RoW},t}$ in the scenario without tipping and the three critical threshold scenarios $T^c \in \{1.8^\circ, 2^\circ, 2.5^\circ\}$. The dashed line is the no-policy growth rate γ_t^{np}. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

The output loss suffered by the RoW raises a natural question: how much of the loss is due to the adversarial emission abatement policies by OECD countries, as opposed to simply caused by having to face a worse climate condition? To answer this question I consider the following counterfactual scenario. How would the output of the RoW evolve in a climate with a tipping point if the OECD behaved as if there were no tipping point? This experiment amounts to making the behavioural assumption that the OECD employs the emission abatement policy $\alpha_{\mathrm{OECD},t}$ as if there was no tipping in the RoW. This, in turn, sticks to the optimal abatement policy $\alpha_{\mathrm{RoW},t}$ for the tipping scenario. The setup imposes a further behavioural assumption on the RoW, as the region is not internalising that the OECD is behaving favourably. Denote by superscript f the state variables and policies in the non-adversarial scenario. In particular, we are interested in the evolution of the counterfactual output gap $Y^f_{\mathrm{RoW},t} - Y_{\mathrm{RoW},t}$ trillion\$ per year, illustrated in Figure 3.8. First, regardless of the scenario, the magnitude of the effect is approximately 0.4 trillion\$ per year, that is, two orders of magnitude smaller than the total cost of introducing a tipping point in the RoW. This suggests that most of the cost to the RoW is driven

by the worst climate rather than the emission abatement policies of the OECD countries. Second, as expected, there is a gap only starting from 2040, when the policies start diverging. Third, in many scenarios a non-adversarial OECD might lead to lower output for the RoW. In the event that the RoW crosses a tipping point, a non-adversarial OECD does not internalise that the abatement incentives for the RoW have decreased and does hence not scale-up its emission abatement. Lastly, if the tipping point is remote and hence the probability of tipping is low on an optimal path, the adversarial abatement policies of the OECD cost the RoW more than 0.2 trillion$ per year by the end of the century.

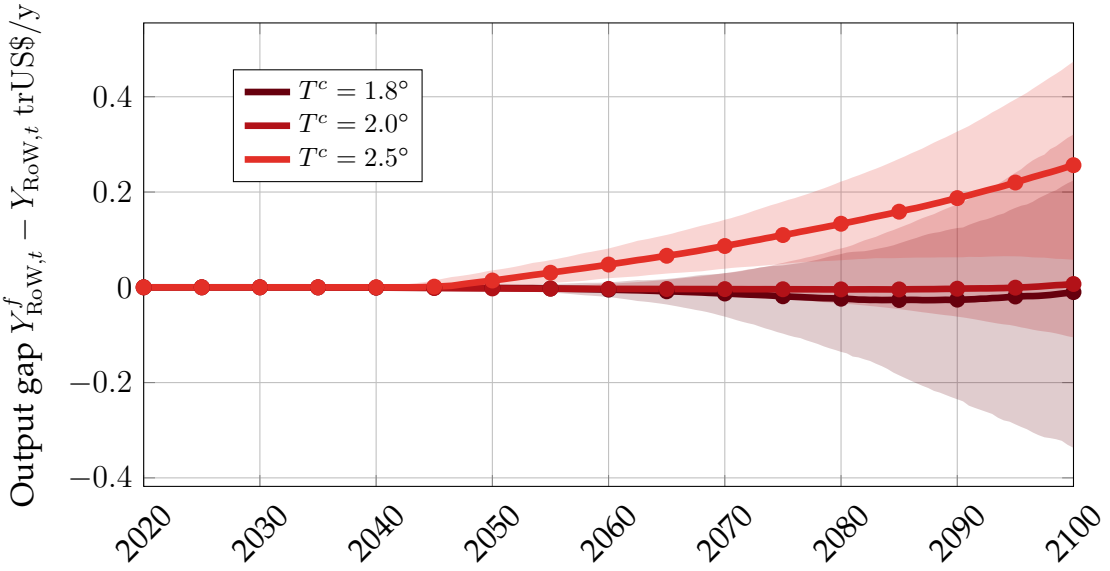

Figure 3.8: Output gap $Y^f_{\mathrm{RoW},t} - Y_{\mathrm{RoW},t}$ output for the RoW when faced with a non-adversarial OECD rather than an adversarial OECD in the three scenarios with tipping points. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

Figure 3.9 displays the same quantity for the OECD region $Y^f_{\mathrm{OECD},t} - Y_{\mathrm{OECD},t}$. In this case, a negative value implies that the output is larger if the OECD behaves adversarially, that is, $Y_{\mathrm{OECD},t} > Y^f_{\mathrm{OECD},t}$. Hence, $Y^f_{\mathrm{OECD},t} - Y_{\mathrm{OECD},t}$ can be seen as the loss deriving from not behaving adversarially. Unlike for the RoW, the benefits for the OECD of its adversarial abatement strategies are positive and large in all cases as $Y_{\mathrm{OECD},t} > Y^f_{\mathrm{OECD},t}$. For example, if the critical threshold is imminent $T^c = 1.8^\circ$, the benefits for the OECD of employing an adversarial emission abatement strategies amount to around 4 trillion$ per year in the median scenario.

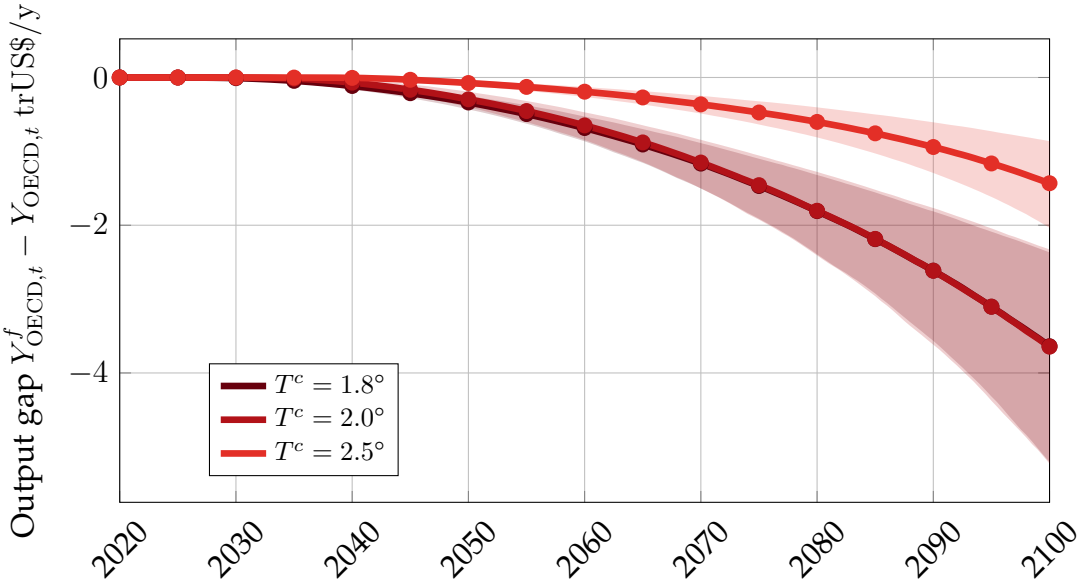

Figure 3.9: Output gap $Y^f_{\mathrm{OECD},t} - Y_{\mathrm{OECD},t}$ output for the OECD when employing a non-adversarial policy rather than an adversarial one in the three scenarios with tipping points. The solid line is the median of 100,000 simulations with 90% simulation interval shaded. The marks represent the value of the median at the beginning of each decade.

3.4 Discussion

This chapter studies the economic consequences of a regional climate tipping point. Two regions, calibrated to the OECD and the rest of the world (RoW), choose their optimal emission abatement policies. Emissions in both regions contribute to the global level of atmospheric CO_2 concentration. However, CO_2 concentration has different impacts on the local temperature and, consequently, on climate damage in the two regions. In particular, the temperature in the RoW presents a tipping point, a critical threshold that, if crossed, pushes the RoW climate to a new regime characterised by higher temperature levels.

I show that if the tipping point is triggered at a lower temperature, this generates more ambitious and rapid emission abatement policies in the RoW. Consequently, the OECD engages in laxer emission abatement efforts, partly offsetting the ambitious abatement policies of the RoW. This adversarial behaviour by the OECD delays net-zero. This effect is not monotonic in the critical temperature threshold: if the threshold is low, it becomes hard for the RoW to prevent a tipping point in its region. In this case, abatement efforts do not have the additional benefit of avoiding a tipping point; hence, the RoW chooses to delay them. In all cases, a tipping point in the RoW's temperature generates significant output losses for the RoW. Part of these are due to higher climate damages, and part are due to the added abatement efforts necessary to compensate for the laxer abatement efforts in the OECD. To quantify

these, I run an experiment where the OECD behaves as if the tipping point is not present to estimate to what extent its adversarial abatement policy can be blamed for the output loss in the RoW. I show that an adversarial OECD policy can cost the RoW up to 400 $\mathrm{billionUS\$}$ annually. On the other hand, the adversarial OECD policy can save the OECD almost $4\mathrm{trillionUS\$}$ annually.

In the model used in this chapter, the two regions interact only via the CO_2 concentration and the OECD is insulated from the tipping point being triggered in the RoW. Yet, in reality, many factors will imply significant spillovers of such a catastrophic event between the two regions, both environmental (e.g. loss of biodiversity, temperature spillovers) and social (e.g. trade, mass migration). Nevertheless, the model can be seen as a cautionary tale on the consequences of not coordinating emission abatement and behaving adversarially whenever one region is susceptible to climate catastrophes and the other believes these do not affect it. A tipping point in the climate system, which would be a strong incentive to speed up abatement efforts in a collaborative world, can be an additional hurdle to net-zero if regions choose their abatement policies strategically and selfishly.

Summary: Economic Consequences of Environmental Catastrophes

In this thesis, I study the economic consequences of climate catastrophes.

The first chapter focuses on the possible repercussions of correlated climate shocks, such as hurricanes or droughts, on supply chains. I model the endogenous formation of supply chains in the presence of correlated disruptions. The incentives of firms to diversify the supply chain risk are concave in the correlation between the disruption events among producers of their input goods. This concavity has consequences for the endogenous formation of the supply chain. If upstream producers are highly diversified, their disruption risk might be correlated, reducing diversification incentives for downstream firms. Because of this mechanism, a slight increase in the correlation of risk among upstream producers due to, for example, climate disruptions to economic activities can generate under-diversification throughout the production network. This creates significant welfare losses. Finally, I show that firms gaining more information on their supply chain risk exacerbates such losses.

The second chapter gives a macroeconomic perspective on the risk of climate tipping points, particularly their unpredictability. Climate tipping points are abrupt, irreversible shifts in the climate system, potentially locking the world into a high-temperature regime that is difficult or impossible to reverse. This chapter examines the economic consequences of such tipping points, focusing on the costs associated with their unpredictability. Using an integrated assessment model with a climate system exhibiting a tipping point, I compute optimal abatement policies, assuming the tipping point is imminent or remote. To place a bound on the economic cost of the uncertainty in the tipping point, I compare two scenarios: a wishful thinker planner erroneously assumes the tipping point is remote and delays abatement, and another where a cautious planner assumes incorrectly that the tipping point is imminent. The uncertainty around irreversible tipping points can cost up to 2.36% of world output. Moreover, I show that proceeding with caution, paying the certain, increasingly affordable costs of abatement today, is cheaper than gambling on the

risk of crossing a tipping point.

The third chapter extends the second chapter by relaxing the idea of a "first best" world. I consider a model in which two regions, the OECD and the Rest of the World, are both subject to the consequences of climate change, but only the latter is exposed to a climate tipping point. This asymmetry allows the OECD region to postpone its abatement effort, forcing the Rest of the World to undertake more ambitious abatement policies to avoid a tipping point. The OECD's strategic behaviour creates significant welfare losses, as the abatement expenditures in the Rest of the World prevent it from capitalising on its faster productivity growth. This equilibrium casts a shadow on the hope that a tipping point or large catastrophes can be a force of international coordination and collaboration in tackling climate change.

Economic actors need to adapt in anticipation of a warming climate. In this thesis, I study how agents adapt in three domains. Firms anticipating climate shocks under-diversify their input sourcing, leading to fragile supply chains. Policymakers generate significant welfare losses by gambling on the uncertainty around climate tipping points and lowering emission abatement targets. World regions less affected by climate tipping points leverage this advantage by forcing poorer regions to bear a large share of the emission abatement burden. These results show how a warming climate, which requires a coordinated response and resilient economic systems, tends to fragment economic agents' behaviour, inducing fragility.

Samenvatting: Economische Gevolgen van Klimaatcatastrofes

In dit proefschrift bestudeer ik de economische gevolgen van klimaatcatastrofes. Het eerste hoofdstuk richt zich op de mogelijke gevolgen van gecorreleerde klimaatschokken, zoals orkanen of droogtes, voor toeleveringsketens. Ik modelleer de endogene vorming van toeleveringsketens in de aanwezigheid van gecorreleerde verstoringen. De prikkels voor bedrijven om risico's in de toeleveringsketen te spreiden zijn concaaf in de correlatie tussen verstoringsevenementen bij producenten van hun inputgoederen. Deze concaviteit heeft gevolgen voor de endogene vorming van de toeleveringsketen. Als producenten stroomopwaarts sterk gespreid zijn, kan hun verstoringsrisico gecorreleerd raken, waardoor bedrijven stroomafwaards minder geneigd zijn om te diversifiëren. Door dit mechanisme kan een lichte toename in de correlatie van risico's onder producenten stroomopwaarts, bijvoorbeeld door klimaateffecten op economische activiteiten, leiden tot onder-diversificatie in het hele productienetwerk. Dit leidt tot aanzienlijke welvaartsverliezen. Ten slotte laat ik zien dat het verschaffen van meer informatie aan bedrijven over hun risico's in de keten deze verliezen juist kan vergroten.

Het tweede hoofdstuk biedt een macro-economisch perspectief op het risico van klimaatkantelpunten, in het bijzonder hun onvoorspelbaarheid. Klimaatkantelpunten zijn abrupte, onomkeerbare veranderingen in het klimaatsysteem, die de wereld kunnen doen belanden in een hoog-temperatuurregime dat moeilijk of onmogelijk terug te draaien is. Dit hoofdstuk onderzoekt de economische gevolgen van dergelijke kantelpunten, waarbij de focus ligt op de kosten van hun onvoorspelbaarheid. Aan de hand van een geïntegreerd beoordelingsmodel met een klimaatsysteem dat een kantelpunt bevat, bereken ik optimale beleidsscenario's voor emissiereductie, ervan uitgaande dat het kantelpunt nabij respectievelijk ver weg is. Om een bovengrens te vinden voor de economische kosten van de onzekerheid rond het kantelpunt, vergelijk ik twee scenario's: een optimistische beleidsmaker gaat ten onrechte uit van een kantelpunt ver in de toekomst en stelt reductiemaatregelen uit, terwijl een voorzichtige beleidsmaker ten onrechte aanneemt dat het kantelpunt

nabij is. De onzekerheid rond onomkeerbare kantelpunten blijkt tot 2,36% van de wereldoutput te kunnen kosten. Bovendien laat ik zien dat het verstandiger is om voorzichtig te handelen en vandaag de zekere, steeds betaalbaardere kosten van emissiereductie te dragen, dan te gokken onder onzekerheid omtrent een kantelpunt.

Het derde hoofdstuk breidt het tweede uit door het idee van een "eerste-beste" wereld los te laten. Ik beschouw een model waarin twee regio's, de OESO en de Rest van de Wereld, beide worden blootgesteld aan de gevolgen van klimaatverandering, maar alleen de laatste aan een klimaatkantelpunt. Deze asymmetrie stelt de OESO in staat om haar mitigatie-inspanningen uit te stellen, waardoor de Rest van de Wereld gedwongen wordt ambitieuzere reductiemaatregelen te nemen om een kantelpunt te vermijden. Het strategische gedrag van de OESO belemmert de Rest van de Wereld om via emissiereductie te profiteren van haar productiviteitsgroei, wat leidt tot welvaartsverlies. Dit evenwicht werpt een schaduw over de hoop dat een kantelpunt of grote rampen kunnen leiden tot internationale coördinatie en samenwerking in de strijd tegen klimaatverandering. Economische actoren moeten zich aanpassen in anticipatie van een opwarmend klimaat.

In dit proefschrift bestudeer ik hoe actoren zich aanpassen op drie terreinen. Bedrijven die klimaatschokken voorzien, spreiden hun inkoop onvoldoende, wat leidt tot kwetsbare toeleveringsketens. Beleidsmakers veroorzaken welvaartsverliezen door te gokken op de onzekerheid rond klimaatkantelpunten en reductiedoelen uit te stellen. Wereldregio's die minder getroffen worden door kantelpunten benutten dit voordeel door armere regio's een onevenredig groot deel van de emissiereductielast op te leggen. Deze resultaten tonen aan hoe een opwarmend klimaat, dat juist om een gecoördineerde aanpak en veerkrachtige economische systemen vraagt, leidt tot fragmentatie van het gedrag van economische actoren en kwetsbaarheid in de hand werkt.

Acknowledgements

Without the following people, this thesis would not exist.

Thank you, Cees and Florian, for your patience in listening to my thoughts and for your guidance in turning them into research. Thank you to the committee members, Professors Van der Ploeg, Tuinstra, Douenne, Schubert, and Dietz, for agreeing to read and consider my work. Thank you, Philippa, for being by my side. Grazie Mamma e Madda per i vostri sacrifici ed aver sopportato la nostra lontananza. Grazie, Ferruccio, per il costante supporto. Grazie Luca, Francesco, Andy per essere stati e venuti fino a qui. Thank you to my colleagues and friends, in no particular order, Sebastian, Gabriela, Katya, Tommaso, Floris, Maddalena, Viola, Alessandro, Sindri, Taco, Josha, Oliver, and Martin for the support and help. Thank you to all mathematicians, economists, and scientists whose work I build on. Thank you to B. Magee, B. Suits, G. Manganelli, L. Meneghello, and G. Morselli for the perspective.

Bibliography

Acemoglu, Daron, and Alireza Tahbaz-Salehi. 2020. "Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions." National Bureau of Economic Research Working Paper 27565. Cited on page 6

Acemoglu, Daron, and Pablo Daniel Azar. 2020. "Endogenous Production Networks." *Econometrica*, 88(1): 33–82. Cited on page 6

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica*, 80(5): 1977–2016. Cited on page 15

Ackerman, Frank, Elizabeth A. Stanton, and Ramón Bueno. 2013. "Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy?" *Environ Resource Econ*, 56(1): 73–84. Cited on page 122

Amelkin, Victor, and Rakesh Vohra. 2023. "Yield uncertainty and strategic formation of supply chain networks." *Networks*, 83(1): 131–152. Cited on pages 6, 7

Armstrong McKay, David I., Arie Staal, Jesse F. Abrams, Ricarda Winkelmann, Boris Sakschewski, Sina Loriani, Ingo Fetzer, Sarah E. Cornell, Johan Rockström, and Timothy M. Lenton. 2022. "Exceeding 1.5°C Global Warming Could Trigger Multiple Climate Tipping Points." *Science*, 377(6611): eabn7950. Cited on pages 33, 36, 45

Austin, Tim. 2015. "Exchangeable Random Measures." *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 51(3). Cited on page 102

Baker, Erin, Leon Clarke, and Ekundayo Shittu. 2008. "Technical Change and the Marginal Cost of Abatement." *Energy Economics*, 30(6): 2799–2816. Cited on pages 52, 71

Baqaee, David Rezza. 2018. "Cascading Failures in Production Networks." *Econometrica*, 86(5): 1819–1838. Cited on pages 1, 6

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica*, 87(4): 1155–1203. Cited on page 6

Barro, Robert J. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*, 99(1): 243–264. Cited on pages 35, 48

Barrot, Jean-Noël, and Julien Sauvagnat. 2016. "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks." *The Quarterly Journal of Economics*, 131(3): 1543–1592. Cited on page 9

Ben-Yami, Maya, Andreas Morr, Sebastian Bathiany, and Niklas Boers. 2024. "Uncertainties Too Large to Predict Tipping Times of Major Earth System Components from Historical Data." *Sci. Adv.*, 10(31): eadl4841. Cited on pages 33, 35, 40, 45, 69

Bierkens, Joris, Paul Fearnhead, and Gareth Roberts. 2019. "The Zig-Zag Process and Super-Efficient Sampling for Bayesian Analysis of Big Data." *Ann. Statist.*, 47(3). Cited on pages 36, 113, 121

Bondarev, Anton A., and Alfred Greiner. 2018. "Global Warming and Technical Change: Multiple Steady-States and Policy Options." Cited on page 124

Boulton, Chris A., Lesley C. Allison, and Timothy M. Lenton. 2014. "Early Warning Signals of Atlantic Meridional Overturning Circulation Collapse in a Fully Coupled Climate Model." *Nature Communications*, 5(1): 5752–5752. Cited on page 45

Byers, Edward, Volker Krey, Elmar Kriegler, Keywan Riahi, Roberto Schaeffer, Jarmo Kikstra, Robin Lamboll, Zebedee Nicholls, Marit Sandstad, Chris Smith, Kaj van der Wijst, Alaa Al -Khourdajie, Franck Lecocq, Joana Portugal-Pereira, Yamina Saheb, Anders Stromman, Harald Winkler, Cornelia Auer, Elina Brutschin, Matthew Gidden, Philip Hackstock, Mathijs Harmsen, Daniel Huppmann, Peter Kolp, Claire Lepault, Jared Lewis, Giacomo Marangoni, Eduardo Müller-Casseres, Ragnhild Skeie, Michaela Werning, Katherine Calvin, Piers Forster, Celine Guivarch, Tomoko Hasegawa, Malte Meinshausen, Glen Peters, Joeri Rogelj, Bjorn Samset, Julia Steinberger, Massimo Tavoni, and Detlef van Vuuren. 2022. "AR6 Scenarios Database." Cited on page 125

Cai, Yongyang, Timothy M. Lenton, and Thomas S. Lontzek. 2016. "Risk of Multiple Interacting Tipping Points Should Encourage Rapid CO2 Emission Reduction." *Nature Clim Change*, 6(5): 520–525. Cited on page 35

Carvalho, Vasco M., and Alireza Tahbaz-Salehi. 2019. "Production Networks: A Primer." *Annual Review of Economics*, 11(1): 635–663. Cited on pages 1, 6, 7

Carvalho, Vasco M., Makoto Nirei, Yukiko Saito, and Alireza Tahbaz-Salehi. 2020. "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake." *Quarterly Journal of Economics*, 136(2): 1255–1321. Cited on page 6

Cousins, Paul D., and Bulent Menguc. 2006. "The Implications of Socialization and Integration in Supply Chain Management." *Journal of Operations Management*, 24(5): 604–620. Cited on page 9

Crost, Benjamin, and Christian P. Traeger. 2013. "Optimal Climate Policy: Uncertainty versus Monte Carlo." *Economics Letters*, 120(3): 552–558. Cited on page 122

D'Autume, Antoine, Katheline Schubert, and Cees Withagen. 2016. "Should the Carbon Price Be the Same in All Countries?" *Journal of Public Economic Theory*, 18(5): 709–725. Cited on page 66

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2009. "Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates." *American Economic Review*, 99(2): 198–204. Cited on page 51

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics*, 4(3): 66–95. Cited on page 51

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot. 2020. "Trade and Domestic Production Networks." *The Review of Economic Studies*, 88(2): 643–668. Cited on page 15

Diaconis, P., and D. Freedman. 1980. "Finite Exchangeable Sequences." *The Annals of Probability*, 8(4). Cited on page 15

Diaconis, Persi, and Svante Janson. 2007. "Graph Limits and Exchangeable Random Graphs." arXiv 0712.2749. Cited on page 15

Dietz, Simon, and Bruno Lanz. 2019. "Growth and Adaptation to Climate Change in the Long Run." IRENE Institute of Economic Research IRENE Working Papers 19-09. Cited on page 51

Dietz, Simon, and Frank Venmans. 2019. "Cumulative Carbon Emissions and Economic Policy: In Search of General Principles." *Journal of Environmental Economics and Management*, 96: 108–129. Cited on pages 52, 71

Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans. 2021*a*. "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists*. Cited on pages 1, 35, 63

Dietz, Simon, James Rising, Thomas Stoerk, and Gernot Wagner. 2021*b*. "Economic Impacts of Tipping Points in the Climate System." *Proc. Natl. Acad. Sci. U.S.A.*, 118(34): e2103081118. Cited on page 68

di Giovanni, Julian, and Andrei A. Levchenko. 2010. "Putting the Parts Together: Trade, Vertical Linkages, and Business Cycle Comovement." *American Economic Journal: Macroeconomics*, 2(2): 95–124. Cited on page 9

Ditlevsen, Peter D., and Sigfus J. Johnsen. 2010. "Tipping Points: Early Warning and Wishful Thinking." *Geophysical Research Letters*, 37(19): 2010GL044486. Cited on pages 33, 41, 45, 69

Dockner, Engelbert J., and Ngo Van Long. 1993. "International Pollution Control: Cooperative versus Noncooperative Strategies." *Journal of Environmental Economics and Management*, 25(1): 13–29. Cited on page 66

Duffie, Darrell, and Larry G. Epstein. 1992. "Asset Pricing with Stochastic Differential Utility." *Review of Financial Studies*, 5(3): 411–436. Cited on page 54

Elliott, Matthew, Benjamin Golub, and Matthew V. Leduc. 2022. "Supply Network Formation and Fragility." *American Economic Review*, 112(8): 2701–2747. Cited on pages 6, 7, 9, 26

Erol, Selman, and Rakesh Vohra. 2022. "Network formation and systemic risk." *European Economic Review*, 148: 104213. Cited on page 6

Ferguson, Thomas S. 1973. "A Bayesian Analysis of Some Nonparametric Problems." *Annals of Statistics*, 1(2). Cited on page 13

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79(3): 733–772. Cited on page 6

Ghil, Michael, and Stephen Childress. 2012. *Topics in Geophysical Fluid Dynamics: Atmospheric Dynamics, Dynamo Theory, and Climate Dynamics.* Vol. 60, Springer Science & Business Media. Cited on page 44

Greiner, Alfred, and Willi Semmler. 2005. "Economic Growth and Global Warming: A Model of Multiple Equilibria and Thresholds." *Journal of Economic Behavior & Organization*, 57(4): 430–447. Cited on page 36

Guivarch, Céline, Thomas Le Gallic, Nico Bauer, Panagiotis Fragkos, Daniel Huppmann, Marc Jaxa-Rozen, Ilkka Keppo, Elmar Kriegler, Tamás Krisztin, Giacomo Marangoni, Steve Pye, Keywan Riahi, Roberto Schaeffer, Massimo Tavoni, Evelina Trutnevyte, Detlef van Vuuren, and Fabian Wagner. 2022. "Using Large Ensembles of Climate Change Mitigation Scenarios for Robust Insights." *Nat. Clim. Chang.*, 12(5): 428–435. Cited on page 125

Hambel, Christoph, Holger Kraft, and Eduardo Schwartz. 2021. "Optimal Carbon Abatement in a Stochastic Equilibrium Model with Climate Change." *European Economic Review*, 132: 103642. Cited on pages 35, 48, 49, 50, 51, 54, 113, 122, 123, 125, 126

Hameed, Abdul, and Faisal Khan. 2014. "A Framework to Estimate the Risk-Based Shutdown Interval for a Processing Plant." *Journal of Loss Prevention in the Process Industries*, 32: 18–29. Cited on page 9

Heijnen, Pim, and Florian Wagener. 2013. "Avoiding an Ecological Regime Shift Is Sound Economic Policy." *Journal of Economic Dynamics and Control*, 37(7): 1322–1341. Cited on page 66

Hillebrand, Elmar, and Marten Hillebrand. 2019. "Optimal Climate Policies in a Dynamic Multi-Country Equilibrium Model." *Journal of Economic Theory*, 179: 200–239. Cited on page 66

Hulten, Charles R. 1978. "Growth Accounting with Intermediate Inputs." *The Review of Economic Studies*, 45(3): 511–518. Cited on pages 1, 6

IPCC. 2023. *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.* . 1 ed., Cambridge University Press. Cited on pages 45, 49, 52

Jiang, Bomin, Daniel Rigobon, and Roberto Rigobon. 2022. "From Just-in-Time, to Just-in-Case, to Just-in-Worst-Case: Simple Models of a Global Supply Chain under Uncertain Aggregate Shocks." *IMF Economic Review*, 70(1): 141–184. Cited on page 9

Kallenberg, Olav. 2005. *Probabilistic Symmetries and Invariance Principles. Probability and Its Applications*, Springer-Verlag. Cited on pages 13, 15

Kamien, Morton I., and Nancy L. Schwartz. 1971. "Sufficient Conditions in Optimal Control Theory." *Journal of Economic Theory*, 3(2): 207–214. Cited on page 35

Kraft, Holger, and Frank Thomas Seifried. 2014. "Stochastic Differential Utility as the Continuous-Time Limit of Recursive Utility." *Journal of Economic Theory*, 151: 528–550. Cited on page 118

Kriegler, Elmar, Nico Bauer, Alexander Popp, Florian Humpenöder, Marian Leimbach, Jessica Strefler, Lavinia Baumstark, Benjamin Leon Bodirsky, Jérôme Hilaire, David Klein, Ioanna Mouratiadou, Isabelle Weindl, Christoph Bertram, Jan-Philipp Dietrich, Gunnar Luderer, Michaja Pehl, Robert Pietzcker, Franziska Piontek, Hermann Lotze-Campen, Anne Biewald, Markus Bonsch, Anastasis Giannousakis, Ulrich Kreidenweis, Christoph Müller, Susanne Rolinski, Anselm Schultes, Jana Schwanitz, Miodrag Stevanovic, Katherine Calvin, Johannes Emmerling, Shinichiro Fujimori, and Ottmar Edenhofer. 2017. "Fossil-Fueled Development (SSP5): An Energy and Resource Intensive Scenario for the 21st Century." *Global Environmental Change*, 42: 297–315. Cited on pages 7, 43, 67, 124

Kuik, Onno, Luke Brander, and Richard S.J. Tol. 2009. "Marginal Abatement Costs of Greenhouse Gas Emissions: A Meta-Analysis." *Energy Policy*, 37(4): 1395–1403. Cited on page 52

Kushner, Harold J. 2007. "Numerical Approximations for Nonzero-Sum Stochastic Differential Games." *SIAM J. Control Optim.*, 46(6): 1942–1971. Cited on pages 72, 113, 117, 121

Kushner, Harold J., and Paul Dupuis. 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time.* Vol. 24 of *Stochastic Modelling and Applied Probability*, New York, NY:Springer New York. Cited on pages 36, 113, 117, 119

Lafrogne-Joussier, Raphael, Julien Martin, and Isabelle Mejean. 2022. "Supply Shocks in Supply Chains: Evidence from the Early Lockdown in China." *IMF Economic Review*. Cited on page 9

Le Kama, Alain Ayong, and Katheline Schubert. 2004. "Growth, Environment and Uncertain Future Preferences." *Environmental and Resource Economics*, 28(1): 31–53. Cited on page 54

Lemoine, Derek, and Christian P. Traeger. 2016*a*. "Ambiguous Tipping Points." *Journal of Economic Behavior & Organization*, 132: 5–18. Cited on pages 35, 48

Lemoine, Derek, and Christian P. Traeger. 2016*b*. "Economics of Tipping the Climate Dominoes." *Nature Clim Change*, 6(5): 514–519. Cited on page 35

Lemoine, Derek, and Christian Traeger. 2014. "Watch Your Step: Optimal Policy in a Tipping Climate." *American Economic Journal: Economic Policy*, 6(1): 137–166. Cited on page 35

Leontief, Wassily. 1936. "Quantitative Input and Output Relations in the Economic Systems of the United States." *The Review of Economics and Statistics*, 18(3): 105. Cited on page 6

Le Quéré, Corinne, Christian Rödenbeck, Erik T. Buitenhuis, Thomas J. Conway, Ray Langenfelds, Antony Gomez, Casper Labuschagne, Michel Ramonet, Takakiyo Nakazawa, Nicolas Metzl, Nathan Gillett, and Martin Heimann. 2007. "Saturation of the Southern Ocean CO2 Sink Due to Recent Climate Change." *Science*, 316(5832): 1735–1738. Cited on page 43

Li, Chuan-Zhong, Anne-Sophie Crépin, and Therese Lindahl. 2024. "The Economics of Tipping Points: Some Recent Modeling and Experimental Advances." *IRERE*, 18(4): 385–442. Cited on pages 1, 35

Lin, Xu, and Sweder Van Wijnbergen. 2023. "The Social Cost of Carbon under Climate Volatility Risk." Tinbergen Institute Tinbergen Institute Discussion Papers 23-032/IV. Cited on page 35

Lontzek, Thomas S., Yongyang Cai, Kenneth L. Judd, and Timothy M. Lenton. 2015. "Stochastic Integrated Assessment of Climate Tipping Points Indicates the Need for Strict Climate Policy." *Nature Climate Change*, 5(5): 441–444. Cited on pages 35, 122

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–2945. Cited on page 9

Mäler, Karl-Göran, Anastasios Xepapadeas, and Aart de Zeeuw. 2003. "The Economics of Shallow Lakes." *Environ Resource Econ*, 26(4): 603–624. Cited on page 36

McGuffie, Kendal, and Ann Henderson-Sellers. 2014. *The Climate Modelling Primer.* . 4. ed ed., Chicester:Wiley Blackwell. Cited on pages 35, 44, 45, 68

Mogensen, Patrick Kofod, and Asbjørn Nilsen Riseth. 2018. "Optim: A Mathematical Optimization Package for Julia." *Journal of Open Source Software*, 3(24): 615. Cited on page 121

Nævdal, Eric, and Michael Oppenheimer. 2007. "The Economics of the Thermohaline Circulation—A Problem with Multiple Thresholds of Unknown Locations." *Resource and Energy Economics*, 29(4): 262–283. Cited on page 35

Nkuiya, Bruno. 2015. "Transboundary Pollution Game with Potential Shift in Damages." *Journal of Environmental Economics and Management*, 72: 1–14. Cited on page 66

Nordhaus, William D. 1992. "An Optimal Transition Path for Controlling Greenhouse Gases." *Science*, 258(5086): 1315–1319. Cited on pages 51, 52, 71

Nordhaus, William D. 2014. "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches." *Journal of the Association of Environmental and Resource Economists*, 1(1): 273–312. Cited on page 122

Nordhaus, William D. 2017. "Revisiting the Social Cost of Carbon." *Proceedings of the National Academy of Sciences of the United States of America*, 114(7): 1518–1523. Cited on pages 52, 71

Nordhaus, William D. 2019. "Economics of the Disintegration of the Greenland Ice Sheet." *Proc. Natl. Acad. Sci. U.S.A.*, 116(25): 12261–12269. Cited on page 36

Nordhaus, William D., and Zili Yang. 1996. "A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies." *The American Economic Review*, 86(4): 741–765. Cited on page 66

Pindyck, Robert S., and Neng Wang. 2013. "The Economic and Policy Consequences of Catastrophes." *American Economic Journal: Economic Policy*, 5(4): 306–339. Cited on pages 35, 48, 50, 54, 123

Rackauckas, Christopher, and Qing Nie. 2017. "Adaptive Methods for Stochastic Differential Equations via Natural Embeddings and Rejection Sampling with Memory." *Discrete and continuous dynamical systems. Series B*, 22(7): 2731. Cited on page 121

Rackauckas, Christopher, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, and Ali Ramadhan. 2020. "Universal Differential Equations for Scientific Machine Learning." *arXiv preprint arXiv:2001.04385*. Cited on page 125

Riihelä, Aku, Ryan M. Bright, and Kati Anttila. 2021. "Recent Strengthening of Snow and Ice Albedo Feedback Driven by Antarctic Sea-Ice Loss." *Nat. Geosci.*, 14(11): 832–836. Cited on page 38

Seaver Wang, A. Foster, E. A. Lenz, J. Kessler, J. Stroeve, L. Anderson, M. Turetsky, R. Betts, Sijia Zou, W. Liu, W. Boos, and Z. Hausfather. 2023. "Mechanisms and Impacts of Earth System Tipping Elements." *Reviews of Geophysics*. Cited on pages 1, 33, 36, 38, 45, 68, 69

Shi, Hao, Hanqin Tian, Naiqing Pan, Christopher P. O. Reyer, Philippe Ciais, Jinfeng Chang, Matthew Forrest, Katja Frieler, Bojie Fu, Anne Gädeke, Thomas Hickler, Akihiko Ito, Sebastian Ostberg, Shufen Pan, Miodrag Stevanović, and Jia Yang. 2021. "Saturation of Global Terrestrial Carbon Sink Under a High Warming Scenario." *Global Biogeochemical Cycles*, 35(10): e2020GB006800. Cited on page 43

Sîrbu, Mihai. 2014. "Stochastic Perron's Method and Elementary Strategies for Zero-Sum Differential Games." *SIAM J. Control Optim.*, 52(3): 1693–1711. Cited on page 72

Skiba, A. K. 1978. "Optimal Growth with a Convex-Concave Production Function." *Econometrica*, 46(3): 527. Cited on page 36

Smith, Christopher J., Piers M. Forster, Myles Allen, Nicholas Leach, Richard J. Millar, Giovanni A. Passerello, and Leighton A. Regayre. 2017. "FAIR v1.1: A Simple Emissions-Based Impulse Response and Carbon Cycle Model." Cited on page 36

Smolders, Emma J. V., René M. van Westen, and Henk A. Dijkstra. 2024. "Probability Estimates of a 21st Century AMOC Collapse." Cited on page 1

Stott, Peter. 2016. "How Climate Change Affects Extreme Weather Events." *Science*, 352(6293): 1517–1518. Cited on page 1

Tsur, Yacov, and Amos Zemel. 1996. "Accounting for Global Warming Risks: Resource Management under Event Uncertainty." *Journal of Economic Dynamics and Control*, 20(6-7): 1289–1305. Cited on page 35

Turetsky, Merritt R., Benjamin W. Abbott, Miriam C. Jones, Katey Walter Anthony, David Olefeldt, Edward A. G. Schuur, Charles Koven, A. David McGuire, Guido Grosse, Peter Kuhry, Gustaf Hugelius, David M. Lawrence, Carolyn Gibson, and A. Britta K. Sannel. 2019. "Permafrost Collapse Is Accelerating Carbon Release." *Nature*, 569(7754): 32–34. Cited on page 39

Vakil, Bindiya. 2021. "The Latest Supply Chain Disruption: Plastics." *Harvard Business Review*. Cited on page 5

Van den Bremer, Ton S., and Frederick Van der Ploeg. 2021. "The Risk-Adjusted Carbon Price." *American Economic Review*, 111(9): 2782–2810. Cited on page 35

Van der Ploeg, Frederick, and Aart De Zeeuw. 2018. "Climate Tipping and Economic Growth: Precautionary Capital and the Price of Carbon." *Journal of the European Economic Association*, 16(5): 1577–1617. Cited on pages 35, 48

Van Westen, René M., Michael Kliphuis, and Henk A. Dijkstra. 2024. "Physics-Based Early Warning Signal Shows That AMOC Is on Tipping Course." *Sci. Adv.*, 10(6): eadk1189. Cited on page 45

Wagener, Florian. 2013. "Regime Shifts: Early Warnings." In *Encyclopedia of Energy, Natural Resource, and Environmental Economics*. 349–359. Elsevier. Cited on pages 36, 45, 69

Wagener, Florian. 2015. "Economics of Environmental Regime Shifts." In *The Oxford Handbook of the Macroeconomics of Global Warming*. , ed. Lucas Bernard and Willi Semmler, 0. Oxford University Press. Cited on page 36

Wagener, Florian, and Aart De Zeeuw. 2021. "Stable Partial Cooperation in Managing Systems with Tipping Points." *Journal of Environmental Economics and Management*, 109: 102499. Cited on page 66

Weitzman, Martin L. 2012. "GHG Targets as Insurance against Catastrophic Climate Damages." *J Public Economic Theory*, 14(2): 221–244. Cited on page 51

Weitzman, Martin L. 2014. "Fat Tails and the Social Cost of Carbon." *American Economic Review*, 104(5): 544–546. Cited on page 49

Zhao, Ming, and Nickolas K. Freeman. 2019. "Robust Sourcing from Suppliers under Ambiguously Correlated Major Disruption Risks." *Production and Operations Management*, 28(2): 441–456. Cited on page 5

Appendices

Appendix A

Chapter 1

The following serves as an appendix to Chapter 1.

A.1 Notation and Distributions

This appendix introduces standard notation and some definitions that will be used throughout the following appendices.

For $x \in \mathbb{R}$ and $n \in \mathbb{N}$, I denote the rising factorial as

$$x^{\overline{n}} := \underbrace{x(x+1)(x+2)\dots(x+(n-1))}_{n \text{ terms}}. \tag{A.1}$$

For non-integer exponents $n \in \mathbb{R}$, the definition (A.1) can be extended as

$$x^{\overline{n}} := \frac{\Gamma(x+n)}{\Gamma(x)}, \tag{A.2}$$

where Γ is the gamma function.

Two properties of the rising factorial that are used below but not proven are the additive property of the exponent

$$x^{\overline{n+m}} = x^{\overline{n}}(x+m)^{\overline{m}}, \tag{A.3}$$

and the fact that it is strictly increasing in its base

$$\frac{\partial x^{\overline{n}}}{\partial x} > 0. \tag{A.4}$$

A.2 Proofs

This appendix contains the proofs omitted from the chapter.

A.2.1 Proof of Proposition 1.11

The proof is a straightforward application of de Finetti's theorem, which I now state.

Lemma 9.1 (De Finetti's Theorem). *An infinite sequence of Bernoulli random variables $X_1, X_2 \ldots$ is exchangeable if and only if there exists a random variable P over the unit interval such that $(X_1, X_2 \ldots)|P$ are independent and identically distributed.*

Using this, we can prove Proposition 1.11.

Proof. By construction the basal disruption probabilities $\{X_{0,j}\}_j$ are independent and identically distributed conditional on P_0. □

A.2.2 Proof of Proposition 2

The proof is built on four steps. First, in Lemma 9.2 I show that downstream firms, facing exchangeable disruptions among upstream producers, are indifferent between suppliers $\mathcal{S}_{k+1,i}$ and only care about the number of sources $|\mathcal{S}_{k+1,i}|$. Second, in Lemma 9.3, I show that this choice is finite and the same among all firms, $s_{k+1} := |\mathcal{S}_{k+1,i}| < \infty$. Third, in Lemma 9.4, I show that, given the source selection procedure of Assumption 3, firms facing exchangeable disruption among suppliers, experience disruptions which are themselves exchangeable. Finally, by induction on the layer k, using Lemma 9.4 as induction step, all downstream layers are shown to be exchangeable.

Lemma 9.2. *If the disruption $\{X_{k,i}\}_i$ events among upstream firms are exchangeable, then the probability that a downstream firm $(k+1, j)$ is disrupted depends only on the number $|\mathcal{S}_{k+1,j}|$ of suppliers it picks.*

Proof. By assumption the sequence $\{X_{k,i}\}_i$ of disruption events among upstream firms is an exchangeable sequence of random variables.

Let $\mathcal{A} = \{X_{k,a_1}, X_{k,a_2} \ldots X_{k,a_n}\}$ and $\mathcal{B} = \{X_{k,b_1}, X_{k,b_2} \ldots X_{k,b_n}\}$ be two subsequences of finite size n, where $a_i, b_i \in \mathbb{N}$. Now consider a permutation $\sigma : \mathbb{N} \to \mathbb{N}$ satisfying

$$\sigma(\mathcal{A}) = \mathcal{B} \text{ and } \sigma(\mathcal{A}^c) = \mathcal{B}^c, \tag{A.5}$$

that is, elements of $\mathcal{A}$ are mapped into $\mathcal{B}$. The joint probability of $\mathcal{A}$ taking a value

$x \in \{0,1\}^{\mathcal{A}}$ is

$$\begin{aligned}
\mathbb{P}(\mathcal{A} = x) &= \sum_{y \in \{0,1\}^{\mathcal{A}^c}} \mathbb{P}(\mathcal{A} = x \text{ and } \mathcal{A}^c = y) \\
\text{then by exchangeability, } &= \sum_{y \in \{0,1\}^{\mathcal{A}^c}} \mathbb{P}(\sigma(\mathcal{A}) = x \text{ and } \sigma(\mathcal{A}^c) = y) \\
&= \sum_{y \in \{0,1\}^{\mathcal{A}^c}} \mathbb{P}(\mathcal{B} = x \text{ and } \mathcal{B}^c = y) \\
&= \mathbb{P}(\mathcal{B} = x),
\end{aligned} \tag{A.6}$$

hence,

$$\mathcal{A} \sim \mathcal{B}. \tag{A.7}$$

□

Lemma 9.3. *If the disruption $\{X_{k,i}\}_i$ events among upstream firms are exchangeable, each firm chooses the number of suppliers $s_{k+1,i} = |\mathcal{S}_{k+1,i}|$ from the same maximising set, that is, for all i,*

$$s_{k+1,i} \in \{s_1^*, s_2^*, \dots s_n^*\} \coloneqq \arg\sup_{s \in \mathbb{N}} \Pi(s), \tag{A.8}$$

using the procedure of Assumption 3. Importantly the maximising set is nonempty and finite.

Proof. By Lemma 9.2, the probability of failure only depends on the number of suppliers in the sourcing strategy $s_{k+1,i} = |S_{k+1,i}|$, that is,

$$\mathbb{E}\left[X_{k+1,i} || \mathcal{S}_{k+1,i}| = s\right] = p_{k+1}(s) \text{ for all } i \tag{A.9}$$

for some probability mass function p_{k+1}. Notice that

$$p_{k+1}(0) = 1. \tag{A.10}$$

Abusing notation, I write the profit function (1.24) as $\Pi : \mathbb{N} \to \mathbb{R}$

$$\Pi(s) = (1 - p_{k+1}(s))\pi - C(s). \tag{A.11}$$

I now show that the set of maximisers

$$\arg\sup_{s \in \mathbb{N}} \Pi(s) \tag{A.12}$$

is nonempty and contains elements which and independent of the downstream firm $(k+1, i)$.

We first notice that without suppliers, the firm makes no profits $\Pi(0) = -C(0) =$

0. Furthermore, as the number of suppliers grows

$$\lim_{s\to\infty} \Pi(s) = (1 - \lim_{s\to\infty} p_{k+1}(s))\pi - \lim_{s\to\infty} C(s) \to -\infty, \tag{A.13}$$

as p_{k+1} is bounded above by 1. Hence, there exists a finite s^* such that $\Pi(s^*)$ is maximal. Furthermore, notice that this s^* is independent of the downstream firm index i. □

Lemma 9.4. *Suppose all downstream firms $(k+1, i)$ select their sourcing strategy $\mathcal{S}_{k+1,i}$ using the procedure outlined in Assumption 3 and the number of sources chosen are the same, that is, $|S_{k+1,i}| = |S_{k+1,j}| = s < \infty$ for all i and j. Then, if upstream disruptions are exchangeable $\{X_{k,j}\}_j$, downstream disruptions $\{X_{k+1,i}\}_i$ are exchangeable.*

The idea of the proof is to construct a random variable P_{k+1} over the unit interval such that, conditional on this random variable, all the downstream disruptions $X_{k+1,j}$ are i.i.d. Bernoulli trials. The random variable P_{k+1} is constructed to be the probability of a disruption $X_{k+1,i}$. Importantly, this random variable is independent of the downstream firm index i. Then by de Finetti's theorem (Lemma 9.1), the disruptions are exchangeable. This is a simple case of the so called "Noise-Outsourcing" lemma in (Austin, 2015, Lemma 3.1, p.7).

Proof. We seek to construct a random variable P_{k+1} such that

$$(X_{k+1,i} \,|P_{k+1} = p_{k+1}) \sim \text{Bernoulli}(p_{k+1}). \tag{A.14}$$

I do this in two steps. In the first step, I compute the probability that a given upstream firm (k, j) is disrupted. In the second step, I compute the probability that a particular firm (k, j) is among downstream firm suppliers $\mathcal{S}_{k+1,i}$.

The first step follows from the assumption of exchangeability of upstream disruption events. Indeed, the upstream disruption events $\{X_{k,j}\}_j$ are exchangeable, hence, by de Finetti's theorem, there exist a random variable on the unit interval P_k such that

$$(X_{k,j} \,|P_k = p_k) \sim \text{Bernoulli}(p_k). \tag{A.15}$$

We can then write

$$\mathbb{P}(X_{k,j} = 1) = P_k. \tag{A.16}$$

The second step builds on Assumption 3. Let

$$(I_1, I_2, I_3 \ldots I_s) \coloneqq \mathcal{S}_{k+1,j}, \tag{A.17}$$

such that I_i is the i-th index of the upstream supplier selected by the downstream producer, using the procedure outlined in Assumption 3. Recall that I_i is a tuple

taking the form (k, j). Then, the probability that an upstream producer (k, j) is selected as a supplier of firm $(k + 1, i)$ is given by

$$\mathbb{P}\left((k,j) \in \mathcal{S}_{k+1,i}\right) = \sum_{m=1}^{s} \mathbb{P}\left((k,j) = I_m \text{ and } (k,j) \notin I_{1:(m-1)}\right), \tag{A.18}$$

where $I_{1:(m-1)} \coloneqq (I_1, I_2, \ldots I_{m-1})$. Equation A.18 decomposes the probability that firm (k, j) is chosen as a supplier into the probability that the firm is picked first $\mathbb{P}((k, j) = I_1)$, added to the probability that it is picked second but not first $\mathbb{P}((k, j) = I_2$ and $(k, j) = I_1)$, and so on. Using conditional probabilities, each summand can be written as

$$\begin{aligned}\mathbb{P}\left((k,j) = I_m \text{ and } (k,j) \notin I_{1:(m-1)}\right) = \mathbb{P}((k,j) = I_m | I_{1:(m-1)} = ((k,j_l))_{l=1}^{m-1}) \times \\ \mathbb{P}\left(I_{1:(m-1)} = ((k,j_l))_{l=1}^{m-1}\right)\end{aligned} \tag{A.19}$$

These can then be written as

$$\mathbb{P}\left((k,j) = I_m | I_{1:(m-1)} = ((k,j_l))_{l=1}^{m-1}\right) = \frac{\alpha_{k,j}}{1 - \sum_{l=1}^{m-1} \alpha_{k,l}} \tag{A.20}$$

and

$$\mathbb{P}\left(I_{1:(m-1)} = ((k,j_l))_{l=1}^{m-1}\right) = \prod_{l}^{m-1} \frac{\alpha_{k,l}}{1 - \sum_{z=1}^{l-1} \alpha_{k,z}}. \tag{A.21}$$

Importantly, none of these depend on the downstream firm index i. Denote the probability that a downstream firm is picked as a supplier as

$$\omega_{k,j} \coloneqq \mathbb{P}((k,j) \in \mathcal{S}_{k+1,i}). \tag{A.22}$$

Then putting this together we can write

$$P_{k+1} = \prod_{j=1}^{n} \omega_{k,j} P_k. \tag{A.23}$$

Crucially, this probability does not depend on the downstream index i.

Finally, we can show that the downstream disruptions $X_{k+1,j}$, conditional on this random variable P_{k+1} are independent and identically distributed Bernoulli trials. First, the downstream disruption events have law

$$\mathbb{P}(X_{k+1,i} = 1 \,|P_{k+1} = p_{k+1}) = p_{k+1} \tag{A.24}$$

hence, $\{X_{k+1,i}\}_i$ are identically distributed. Second, for two distinct downstream

firms i_1, i_2, their joint probability distribution can be decomposed as

$$\mathbb{P}(X_{k+1,i_1}=1, X_{k+1,i_2}=1 \mid P_{k+1}) = \sum_{j_1,j_2} \mathbb{P}(j_1 \in \mathcal{S}_{k+1,i_1}, j_2 \in \mathcal{S}_{k+1,i_2} \mid P_{k+1}) \times \\ \mathbb{P}\left(X_{k+1,i_1}=1, X_{k+1,i_2}=1 \mid j_1 \in \mathcal{S}_{k+1,i_1}, j_2 \in \mathcal{S}_{k+1,i_2}, P_{k+1}\right) \tag{A.25}$$

First, conditional on P_{k+1}, the sampling procedure is independent, hence

$$\mathbb{P}(j_1 \in \mathcal{S}_{k+1,i_1}, j_2 \in \mathcal{S}_{k+1,i_2} | P_{k+1}) = \mathbb{P}(j_1 \in \mathcal{S}_{k+1,i_1} | P_{k+1}) \mathbb{P}(j_2 \in \mathcal{S}_{k+1,i_2} | P_{k+1}). \tag{A.26}$$

Second, conditional on the sampling and P_{k+1}, the failure probabilities are independent, that is,

$$\mathbb{P}(X_{k+1,i_1}=1, X_{k+1,i_2}=1 \mid j_1 \in \mathcal{S}_{k+1,i_1}, j_2 \in \mathcal{S}_{k+1,i_2}, P_{k+1}) \; = \\ \mathbb{P}\left(X_{k+1,i_1}=1 \mid j_1 \in \mathcal{S}_{k+1,i_1}, P_{k+1}\right) \mathbb{P}\left(X_{k+1,i_2}=1 \mid j_2 \in \mathcal{S}_{k+1,i_2}, P_{k+1}\right). \tag{A.27}$$

This allows us to write the joint disruption probability as

$$\begin{aligned} &\mathbb{P}(X_{k+1,i_1}=1, X_{k+1,i_2}=1 | P_{k+1}) = \\ &\sum_{j_1} \mathbb{P}(j_1 \in \mathcal{S}_{k+1,i_1} | P_{k+1}) \mathbb{P}(X_{k+1,i_1}=1 | j_1 \in \mathcal{S}_{k+1,i_1}, P_{k+1}) \times \\ &\sum_{j_2} \mathbb{P}(j_2 \in \mathcal{S}_{k+1,i_2} | P_{k+1}) \mathbb{P}(X_{k+1,i_2}=1 | j_2 \in \mathcal{S}_{k+1,i_2}, P_{k+1}) = \\ &\mathbb{P}(X_{k+1,i_1}=1 | P_{k+1}) \mathbb{P}(X_{k+1,i_2}=1 | P_{k+1}), \end{aligned} \tag{A.28}$$

hence X_{k+1,i_1} and X_{k+1,i_2} are conditionally independent.

As the downstream disruption events $X_{k+1,j}$ are i.i.d. Bernoulli trials conditional on P_{k+1}, by de Finetti's theorem, they are exchangeable.

□

I now prove Proposition 2.

Proof. The proof is done by induction, with the base case $k = 0$ given by Proposition 1.11. Assume that for some layer $k \geq 0$ the disruption events $\{X_{k,j}\}_i$ are exchangeable. By Lemma 9.2 and Lemma 9.3, all downstream firms choose the same number of sources, that is, $|S_{k+1,i}| = s_{k+1} < \infty$ for all i. Then by Assumption 3 and Lemma 9.4, the disruption events $X_{k+1,i}$ are exchangeable. □

A.2.3 Proof of Proposition 3

Lemma 9.5. *For any two firms i_1 and i_2, the probability that they share a supplier becomes arbitrarily small as $\theta \to \infty$, that is,*

$$\lim_{\theta \to \infty} \mathbb{P}\left(\mathcal{S}_{k+1,i_1} \cap \mathcal{S}_{k+1,i_2} \neq \emptyset\right) = 0. \tag{A.29}$$

Proof. Conditional on the weights α_j, the probability that a firm is not selected among the suppliers of a downstream firm is

$$\begin{aligned} \mathbb{P}((k,j) \notin \mathcal{S}_{k+1,i} \mid \alpha_j) &= \prod_{l=1}^{s} \left(1 - \frac{\alpha_j}{1 - \sum_{(k,l) \in \mathcal{S}_{k+1,i}} \alpha_l}\right) \\ &\geq \prod_{l=1}^{s} \left(1 - \frac{\alpha_j}{1 - s\bar{\alpha}}\right) = \left(1 - \frac{\alpha_j}{1 - s\bar{\alpha}}\right)^s, \end{aligned} \tag{A.30}$$

where $\bar{\alpha} \coloneqq \sup_{j \in \mathbb{N}} \alpha_j$ and $s = |\mathcal{S}_{k+1,j}|$. Hence, the probability that a firm is selected as a supplier $\mathcal{S}_{k+1,i}$ can be bounded above by

$$\begin{aligned} \mathbb{P}((k,j) \in \mathcal{S}_{k+1,i} \mid \alpha_j) &= 1 - \mathbb{P}((k,j) \notin \mathcal{S}_{k+1,i} \mid \alpha_j) \\ &\leq 1 - \left(1 - \frac{\alpha_j}{1 - s\bar{\alpha}}\right)^s \\ &\leq \frac{s\alpha_j}{1 - s\bar{\alpha}}, \end{aligned} \tag{A.31}$$

where we have used the Bernoulli inequality $1 - np \leq (1-p)^n$. Using this we can compute the probability that two firms i_1 and i_2 share suppliers

$$\begin{aligned} \mathbb{P}\left(\mathcal{S}_{k+1,i_1} \cap \mathcal{S}_{k+1,i_2} \neq \emptyset \mid \alpha_j\right) &\leq \sum_j \mathbb{P}\left(j \in \mathcal{S}_{k+1,i_1} \mid \alpha_j\right) \mathbb{P}\left(j \in \mathcal{S}_{k+1,i_2} \mid \alpha_j\right) \\ &\leq \sum_j \left(\frac{s\alpha_j}{1 - s\bar{\alpha}}\right)^2 = \frac{s^2}{(1 - s\bar{\alpha})^2} \sum_j \alpha_j^2. \end{aligned} \tag{A.32}$$

Now we need to show that the unconditional probability goes to 0 as $\theta \to \infty$. First notice that

$$\frac{s^2}{(1 - s\bar{\alpha})^2} \leq \frac{s^2}{(1-s)^2}, \tag{A.33}$$

hence

$$\begin{aligned} \mathbb{P}\left(\mathcal{S}_{k+1,i_1} \cap \mathcal{S}_{k+1,i_2} \neq \emptyset\right) &= \mathbb{E}_{\alpha_j}\left[\mathbb{P}\left(\mathcal{S}_{k+1,i_1} \cap \mathcal{S}_{k+1,i_2} \neq \emptyset \mid \alpha_j\right)\right] \\ &\leq \frac{s^2}{(1-s)^2} \sum_j \mathbb{E}\left[\alpha_j^2\right]. \end{aligned} \tag{A.34}$$

Finally, to tackle the summand, we can use the definition of the weights to

rewrite

$$\mathbb{E}\left[\alpha_j^2\right] = \mathbb{E}[u_j^2] \prod_{l<j} \mathbb{E}[(1-u_l)^2], \tag{A.35}$$

where

$$\mathbb{E}[u_j^2] = \int_0^1 u^2 \theta (1-u)^{\theta-1} \, \mathrm{d}u = \frac{2}{(1+\theta)(2+\theta)} \tag{A.36}$$

and

$$\mathbb{E}[(1-u_l)^2] = \int_0^1 u^2 \theta u^{\theta-1} \, \mathrm{d}u = \frac{\theta}{\theta+2}, \tag{A.37}$$

arriving at

$$\mathbb{E}\left[\alpha_j^2\right] = \frac{2}{(1+\theta)(2+\theta)} \left(\frac{\theta}{\theta+2}\right)^{j-1}. \tag{A.38}$$

Using this we can finally write

$$\begin{aligned} \sum_j \mathbb{E}[\alpha_j^2] &= \frac{2}{(1+\theta)(2+\theta)} \sum_j \left(\frac{\theta}{\theta+2}\right)^{j-1} \\ &= \frac{2}{(1+\theta)(2+\theta)} \frac{1}{1-\frac{\theta}{\theta+2}} = \frac{1}{1+\theta}. \end{aligned} \tag{A.39}$$

Hence, we have

$$\mathbb{P}\left(\mathcal{S}_{k+1,i_1} \cap \mathcal{S}_{k+1,i_2} \neq \emptyset\right) \leq \frac{s^2}{(1-s)^2} \frac{1}{1+\theta} \to 0 \text{ as } \theta \to \infty. \tag{A.40}$$

□

Using Lemma 9.5, we can now prove Proposition 3.

Proof. A firm producing good $k+1$ sources from $s_{k+1} < \infty$ suppliers. A disruption for the downstream firm is defined as

$$X_{k+1,i} = \prod_{(k,j)\in\mathcal{S}_{k+1,j}} X_{k,j}. \tag{A.41}$$

The sequence of upstream disruptions $\{X_{k,j}\}_{(k,j)\in\mathcal{S}_{k+1,j}}$ is a subsequence of an exchangeable sequence $\{X_{k,j}\}_{j\in\mathbb{N}}$, hence it is itself exchangeable. Furthermore, by assumption, conditional on the random variable $P_k \sim \text{BetaPower}(\mu_0, \rho_0, S_k)$, the disruption events $\{X_{k,j}\}_{(k,j)\in\mathcal{S}_{k+1,j}}$ are i.i.d. Bernoulli trials sampled independently.

Using this, we can derive the probability of a downstream disruption

$$\begin{aligned} P_{k+1} \coloneqq \mathbb{E}[X_{k+1,i}] &= \mathbb{E}\left[\prod_{l=1}^{s_k} X_{k,j_l}\right] \\ \text{by the law of iterated expectations, } &= \mathbb{E}\left[\mathbb{E}\left[\prod_{l=1}^{s_k} X_{k,j_l} \middle| P_k\right]\right]. \end{aligned} \tag{A.42}$$

By Lemma 9.5, in the limit $\theta \to \infty$, the probability that two upstream suppliers share sources is zero, $X_{k,j_l}\big|P_k$ are independent and identically distributed Bernoulli trials with probabilities independently sampled from P_k, hence we can rewrite

$$P_{k+1} = \prod_{l=1}^{s_k} \mathbb{E}\left[\mathbb{E}\left[X_{k,j_l}\middle|P_k\right]\right] = \prod_{l=1}^{s_k} \mathbb{E}\left[X_{k,jl}\right] = P_k^{s_k}. \tag{A.43}$$

□

A.2.4 Mapping of risk across layers

Introduce

$$r_0 \coloneqq \frac{\rho_0}{1-\rho_0}. \tag{A.44}$$

Proposition 10. *The expected probability of disruption faced by a firm is given by*

$$\mu_k = \mathbb{E}\Big[P_k\Big] = \frac{B\left(\mu_0\, r_0 + S_k, (1-\mu_0)\, r_0\right)}{B\left(\mu_0\, r_0, (1-\mu_0)\, r_0\right)}, \tag{A.45}$$

where $S_k = \prod_{l=1}^{k} s_l$ and B is the standard beta function defined as

$$B(\alpha,\beta) \coloneqq \int_0^1 p^{\alpha-1}(1-p)^{\beta-1}\,\mathrm{d}p = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}. \tag{A.46}$$

Proof. The disruption probability $P_k \sim \text{BetaPower}(\mu_0, \rho_0, S_k)$, hence it can be written as

$$P_k = P_0^{S_k} \text{ where } P_0 \sim \text{Beta}(\mu_0, \rho_0). \tag{A.47}$$

Therefore, the expected probability of a disruption

$$\mu_k = \mathbb{E}[P_k] = \mathbb{E}\left[P_0^{S_k}\right] \tag{A.48}$$

is the S_k-th moment of P_0. If $S_k = 0$, then the distribution $P_k = 1$, hence $\mu_k = 1$. If $S_k > 0$, μ_k can be derived using the moment generating function of the beta

distribution, which is given by

$$M(t) = \sum_{n=0}^{\infty} \frac{t^n}{n!} \frac{B((1-\mu_0)r_0 + n, \mu_0 r_0)}{B((1-\mu_0)r_0, \mu_0 r_0)}. \tag{A.49}$$

Using this we can obtain

$$\mu_k = \mathbb{E}\left[P_0^{S_k}\right] = M^{(S_k)}(0) = \frac{B((1-\mu_0)r_0 + S_k, \mu_0 r_0)}{B((1-\mu_0)r_0, \mu_0 r_0)}, \tag{A.50}$$

where $M^{(n)}$ is the n-th derivative of M. □

Using this result we can derive the "risk reduction factor" η, if $S_k > 0$, as

$$\eta(s_{k+1}, S_k) \coloneqq \frac{\mu_{k+1}}{\mu_k} = \frac{B((1-\mu_0)r_0 + S_k s_{k+1}, \mu_0 r_0)}{B((1-\mu_0)r_0 + S_k, \mu_0 r_0)}. \tag{A.51}$$

By the definition of the B function in terms of the Γ function, we can rewrite η in terms of a falling factorial

$$\eta(s, S) = \frac{(\mu_0 r_0 + S)^{\overline{S(s-1)}}}{(r_0 + S)^{\overline{S(s-1)}}} \tag{A.52}$$

which satisfies the recursion

$$\eta(s+1, S) = \eta(s, S) \frac{(\mu_0 r_0 + Ss)^{\overline{S}}}{(r_0 + Ss)^{\overline{S}}}. \tag{A.53}$$

Corollary 10.1. *The risk reduction factor $\eta(s, S)$ is strictly decreasing in s.*

Proof. The derivative of η can be expressed in terms of the digamma function

$$\psi(z) \coloneqq \frac{\Gamma'(z)}{\Gamma(z)}. \tag{A.54}$$

In particular,

$$\frac{\partial \eta}{\partial s}(s, S) = \eta(s, S) S \left(\psi(\mu_0 r_0 + Ss) - \psi(r_0 + Ss)\right). \tag{A.55}$$

The digamma function ψ is strictly increasing over $[0, \infty)$, hence $\psi(\mu_0 r_0 + Ss) - \psi(r_0 + Ss) < 0$, as $\mu_0 < 1$. This implies that $\frac{\partial \eta}{\partial s}(s, S) < 0$. □

A.2.5 Limit case $\rho_0 = 0$

To prove Corollary 7.1, I first related the fixed point of the one-dimensional risk propagation map ϕ to the fixed point of its "continuous" counterpart $\tilde{\phi}$, that is, the map that would arise if the firms could choose $s_{k+1} = \tilde{s}_{k+1} \in \mathbb{R}$, as opposed to $s_{k+1} \in \mathbb{N}$.

Definition 5. *Introduce the continuous law of motion*

$$\mu_{k+1} \coloneqq \tilde{\phi}(\mu_k) \coloneqq \mu^{\tilde{s}_{k+1}(\mu_k)}, \tag{A.56}$$

where $\tilde{s}_{k+1}$ is the desired sourcing strategy as in Definition 4.

Lemma 10.1. *A fixed point $\bar{\mu}$ of the law of motion ϕ, is attained if and only if $\tilde{\phi}(\bar{\mu}) \geq \bar{\mu}$.*

Proof. Let $\bar{\mu}$ be a fixed point of ϕ, as defined in equation (1.31). By Definition 5,

$$\tilde{\phi}(\bar{\mu}) = \bar{\mu}^{\tilde{s}(\bar{\mu})}. \tag{A.57}$$

As $\bar{\mu} \leq 1$,

$$\tilde{\phi}(\bar{\mu}) \geq \bar{\mu} \iff \tilde{s}(\bar{\mu}) \in (0, 1]. \tag{A.58}$$

By definition $s(\bar{\mu}) = \lceil \tilde{s}(\bar{\mu}) \rceil = 1$, which implies that $\phi(\bar{\mu}) = \bar{\mu}$. □

Now we can prove Corollary 7.1.

Proof. As shown in Lemma 10.1, a risk level μ satisfying $\tilde{\phi}(\mu) \geq \mu$ implies that μ is a fixed point of the map ϕ. The risk level μ satisfies the property $\tilde{\phi}(\mu) \geq \mu$ if and only if the desired sourcing strategy is $\tilde{s}(\mu) \in (0, 1]$. For this to be the case, it must the marginal profits $\Delta\Pi_{k+1}$ (1.28) intersect the $s_{k+1} = 0$ axis in the interval $(0, 1]$, that is,

$$\Delta\Pi_{k+1}(0) > 0 \text{ and } \Delta\Pi_{k+1}(1) \leq 0.$$

These yield the desired inequality. □

A.2.6 General Case, $\rho_0 > 0$

This appendix proves that the optimal sourcing strategy, shown to exist in Lemma 9.3, is unique. Going from existence to uniqueness is possible, as now, we have imposed a particular structure on the profit function Π_{k+1}, and we have derived a closed form solution for the propagation of risk (A.52).

Proof. It is sufficient to show that the marginal profit $\Delta\Pi_{k+1} : \mathbb{R} \to \mathbb{R}$ is strictly decreasing in its input s.

The marginal profit function can be written in terms of the risk reduction factor (A.52) as

$$\begin{aligned}\Delta\Pi_{k+1}(s) &= \big(\eta(s, S_k) - \eta(s+1, S_k)\big)\pi\mu_k - c\left(s + \frac{1}{2}\right) \\ &= \left(1 - \frac{(\mu_0 r_0 + S_k s)^{\overline{S_k}}}{(r_0 + S_k s)^{\overline{S_k}}}\right)\eta(s, S_k)\mu_k\pi - c\left(s + \frac{1}{2}\right),\end{aligned} \tag{A.59}$$

where we have used the recursion (A.53). Recall that $(\cdot)^{\overline{S_k}}$ is the raising factorial, as defined in equation (A.1). Taking derivatives with respect to s yields

$$\begin{aligned}\Delta\Pi'_{k+1}(s) = \partial_s &\left(1 - \frac{(\mu_0 r_0 + S_k s)^{\overline{S_k}}}{(r_0 + S_k)^{\overline{S_k}}}\right)\eta(s, S_k)\mu_k\pi + \\ &\left(1 - \frac{(\mu_0 r_0 + S_k s)^{\overline{S_k}}}{(r_0 + S_k)^{\overline{S_k}}}\right)\partial_s\eta(s, S_k)\mu_k\pi - c.\end{aligned} \tag{A.60}$$

Considering all terms separately, we have

$$\underbrace{\partial_s\left(1 - \frac{(\mu_0 r_0 + S_k s)^{\overline{S_k}}}{(r_0 + S_k)^{\overline{S_k}}}\right)}_{\leq 0}\underbrace{\eta(s, S_k)\mu_k\pi}_{>0},$$

as $\mu_k, \mu_0 \in [0, 1]$ and by the definition of η (A.52),

$$\underbrace{\left(1 - \frac{(\mu_0 r_0 + S_k s)^{\overline{S_k}}}{(r_0 + S_k)^{\overline{S_k}}}\right)\mu_k\pi}_{\geq 0}\underbrace{\partial_s\eta(s, S_k)}_{<0},$$

by the strictly decreasing property of η in s, proven in Corollary 10.1, and finally $-c < 0$. Putting the three terms together we have $\Delta\Pi'_{k+1}(s) < 0$.

□

A.3 Solution of the Social Planner Problem

First, notice that the terminal condition V_K is linear in P_{K-1}, hence

$$\mathbb{E}\big[V_K(P_{K-1})\big] = V_K\big(\mathbb{E}[P_{K-1}]\big). \tag{A.61}$$

In turn, this implies that V_k is linear for all k. Hence, we can rewrite the value to be a function of the state space S,

$$V_k(S_{k-1}) = \max_s \left\{ \Big(1 - \mathbb{E}\big[P_k\big]\Big)\pi - \frac{c}{2}s^2 + V_k(S_{k-1}\ s_k) \right\}, \tag{A.62}$$

where

$$P_k \sim \text{BetaPower}(\mu_0, \rho_0, S_{k-1}\ s). \tag{A.63}$$

We then find V_k numerically. Let the state space be

$$\Omega_m \coloneqq (1, 2, 3 \dots m) \times (1, 2, 3, \dots m, m+1, \dots m^2, m^2+1, \dots, m^K) \tag{A.64}$$

for some $m \in \mathbb{N}$ sufficiently large. This can be seen as a projection of $\mathbb{N} \times \mathbb{N}$ into Ω_m by simple truncation, where we just make m larger as the firms pair with more suppliers. We are guaranteed $m < \infty$, as the socially optimal sourcing strategy in each layer $s_k < \infty$ for each $k \in \{0, 1, 2, \dots K\}$. The payoff of the planner in each layer k is given by

$$l(s, S_k) \coloneqq \Big(1 - \mathbb{E}\big[P_k\big]\Big)\pi - \frac{c}{2}s^2. \tag{A.65}$$

Denote by $l(\Omega_m)$ the image of l from Ω_m. Then, by means of backward induction we obtain a recursive expression for $W \equiv V_1$, namely

$$\begin{aligned} V_K(S_{K-1}) &= \max_s l(\Omega_m), \\ V_{k-1}(S_{k-2}) &= \max_s l(\Omega_m) + V_K(S_{k-2} \times s), \\ &\vdots \\ V_1(s) &= \max_s l(\Omega_m) + V_2(s). \end{aligned} \tag{A.66}$$

Appendix B

Chapter 2 and 3

The following serves as an appendix to Chapter 2 and 3.

B.1 Numerical Solution of Climate-Economy Model

This appendix illustrates the solution of the maximisation problem (Section 2.4) of the social planner dealing with a global tipping point and the differential game (Section 3.2.3) between the two regions with a regional tipping point.

The numerical solution relies on three steps. First, in Section B.1.1, I make some simplifying assumption on the decay rate of carbon δ_m. This allows to drop the carbon stored in sinks N_t from the state space. Second, in Section B.1.2, I use this assumption to derive a Hamilton-Jacobi-Bellman partial differential equation for the value function of the social planner V (2.40) and a system of Hamilton-Jacobi-Bellman partial differential equations for the value function of the games V_i (3.11). Third, in Section B.1.3, I derive an approximating Markov chain, that is,a discretisation of the state space and a discrete Markov chain over it, which converges to the Hamilton-Jacobi-Bellman equations. This extends the techniques developed by Kushner and Dupuis (2001) and Kushner (2007) to a problem with recursive utility, such as the ones solved in Sections 2.4 and 3.2.3. I further provide a parallelisation technique based on Bierkens, Fearnhead and Roberts (2019).

B.1.1 Simplifying Assumptions on the Decay Rate of Carbon

Following Hambel, Kraft and Schwartz (2021), I make an assumption on the decay rate of carbon. The calibrated carbon decay δ_m, as a function of the carbon stored in sinks N_t, is illustrated in Figure B.1. The calibration assumes a functional form

$$\delta_m(N_t) = a_\delta e^{-\left(\frac{N_t - c_\delta}{b_\delta}\right)^2}, \tag{B.1}$$

for parameters $a_\delta, b_\delta, c_\delta$. I further assume that the amount of carbon sinks present

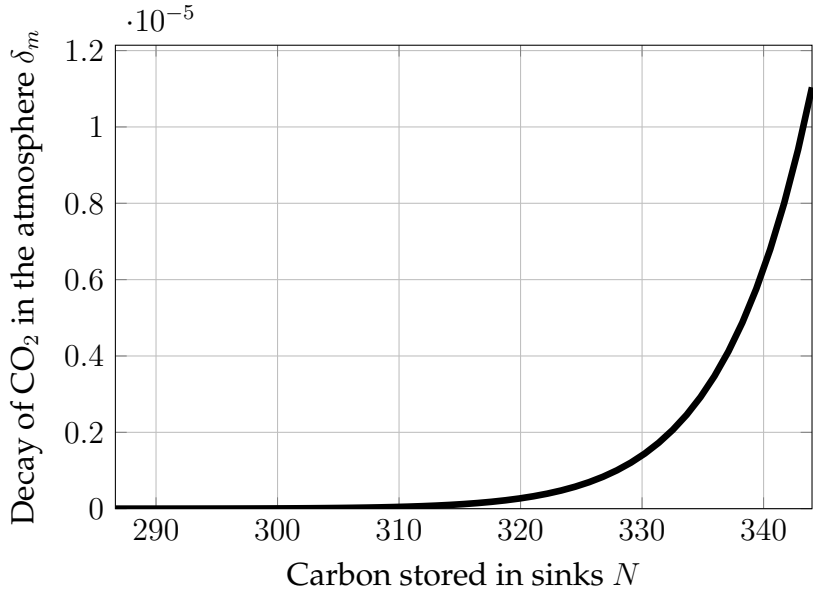

Figure B.1: Estimated decay of carbon δ_m as a function of the carbon stored in sinks N_t.

in the atmosphere is a constant fraction of the concentration in the atmosphere, $N_t = \frac{N_0}{M_0} M_t$. Abusing notation, I henceforth write $\delta_m(M_t)$ for the decay rate. Using this setup, under a no-policy scenario, the decay of carbon follows the path in Figure B.2.

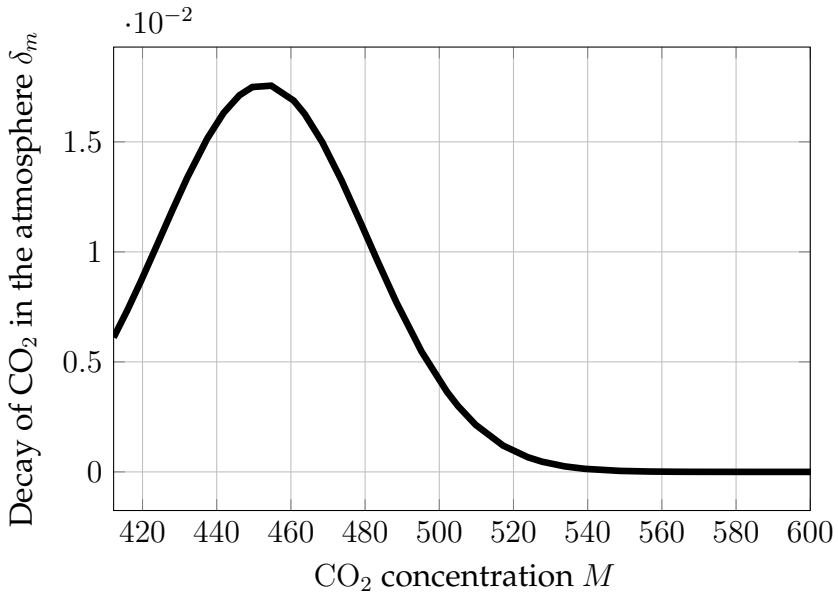

Figure B.2: Estimated decay of carbon δ_m under the no-policy emission scenario M^{np}.

B.1.2 Hamilton-Jacobi-Bellman equation

This section derives the Hamilton-Jacobi-Bellman equation for the social planner problem and the emission game.

Social Planner Problem

The value function of the social planner V (2.40) at time t depends only on temperature T_t, log-CO_2 concentration m_t and output Y_t. This satisfies the Hamilton-Jacobi-Bellman equation

$$\begin{aligned} -\partial_t V = \sup_{\chi,\alpha} & f(\chi Y, V) + \partial_m V \ (\gamma_t^{\text{np}} - \alpha) + \partial_m^2 V \ \frac{\sigma_m^2}{2} \\ & + \partial_Y V(\varrho + \phi(\chi) - d(T) - \beta_t(\varepsilon(\alpha)))Y + \partial_Y^2 V \ \frac{\sigma_Y^2}{2} Y^2 \\ & + \partial_T V \ \frac{r(T) + g(m)}{\epsilon} + \partial_T^2 V \ \frac{(\sigma_T/\epsilon)^2}{2}. \end{aligned} \tag{B.2}$$

Now, I derive an auxiliary, lower dimensional Hamilton-Jacobi-Bellman equation using the ansatz

$$V_t(T, m, Y) \equiv \frac{Y^{1-\theta}}{1-\theta} F_t(T, m). \tag{B.3}$$

First, I compute the Epstein-Zin aggregator at the ansatz $V \equiv F\frac{Y^{1-\theta}}{1-\theta}$. Recall that the Epstein-Zin aggregator is given by

$$f(\chi Y, V) \coloneqq \rho \frac{(1-\theta)\ V}{1 - 1/\psi} \left(\left(\frac{\chi Y}{((1-\theta)V)^{\frac{1}{1-\theta}}} \right)^{1-1/\psi} - 1 \right). \tag{B.4}$$

Notice that

$$(1-\theta)V = FY^{1-\theta} \tag{B.5}$$

$$((1-\theta)V)^{\frac{1}{1-\theta}} = F^{\frac{1}{1-\theta}} Y. \tag{B.6}$$

Plugging back into the aggregator we obtain

$$f(\chi Y, V) = Y^{1-\theta} \ \frac{\rho F}{1 - 1/\psi} \left(\left(\frac{\chi}{F^{\frac{1}{1-\theta}}} \right)^{1-1/\psi} - 1 \right). \tag{B.7}$$

Hence, f is homogenous of degree $1-\theta$ in Y. Second, we can evaluate the derivates

of the value function at our ansatz

$$\begin{aligned}
\partial_Y V &= Y^{-\theta} F, \qquad \partial_Y^2 V = -\theta Y^{-\theta-1} F, \\
\partial_T V &= \frac{Y^{1-\theta}}{1-\theta}\,\partial_T F,\ \partial_T^2 V = \frac{Y^{1-\theta}}{1-\theta}\,\partial_T^2 F, \\
\partial_m V &= \frac{Y^{1-\theta}}{1-\theta}\,\partial_m F,\ \partial_m^2 V = \frac{Y^{1-\theta}}{1-\theta}\,\partial_m^2 F, \\
\partial_t V &= \frac{Y^{1-\theta}}{1-\theta}\,\partial_t F.
\end{aligned}$$

Plugging this into the HJB equation we obtain

$$\begin{aligned}
-\frac{Y^{1-\theta}}{1-\theta}\partial_t F = \sup_{\chi,\alpha}\Bigg\{ & Y^{1-\theta}\frac{\rho F}{1-1/\psi}\left(\left(\frac{\chi}{F^{1/(1-\theta)}}\right)^{1-1/\psi}-1\right) \\
& + \frac{Y^{1-\theta}}{1-\theta}\partial_m F(\gamma_t^{\mathrm{np}}-\alpha) + \frac{Y^{1-\theta}}{1-\theta}\partial_m^2 F\frac{\sigma_m^2}{2} \\
& + Y^{-\theta}F\left(\varrho+\phi(\chi)-d(T)-\beta_t(\varepsilon(\alpha))\right)Y \\
& - \theta Y^{-\theta-1}F\frac{\sigma_Y^2}{2}Y^2 \\
& + \frac{Y^{1-\theta}}{1-\theta}\partial_T F\frac{r(T)+g(m)}{\epsilon} + \frac{Y^{1-\theta}}{1-\theta}\partial_T^2 F\frac{(\sigma_T/\epsilon)^2}{2}\Bigg\}
\end{aligned} \tag{B.8}$$

Dividing both sides by $\frac{Y^{1-\theta}}{1-\theta}$, we obtain

$$\begin{aligned}
-\partial_t F = \sup_{\chi,\alpha}\Bigg\{ & (1-\theta)\frac{\rho F}{1-1/\psi}\left(\left(\frac{\chi}{F^{1/(1-\theta)}}\right)^{1-1/\psi}-1\right) \\
& + (1-\theta)F\left(\varrho+\phi(\chi)-d(T)-\beta_t(\varepsilon(\alpha))-\theta\frac{\sigma_Y^2}{2}\right) \\
& + \partial_m F(\gamma_t^{\mathrm{np}}-\alpha) + \partial_m^2 F\frac{\sigma_m^2}{2} \\
& + \partial_T F\frac{r(T)+g(m)}{\epsilon} + \partial_T^2 F_t\frac{(\sigma_T/\epsilon)^2}{2}\Bigg\}
\end{aligned} \tag{B.9}$$

This can be interpreted as the HJB equation of an auxiliary optimisation problem

$$F_t(T_t, m_t) = (1-\theta)\sup_{\alpha,\chi}\int_0^\infty g(\chi_s, F_s(T_s,m_s); s, T_s)\,\mathrm{d}s \text{ where} \tag{B.10}$$

$$\begin{aligned}
g(\chi, F; t, T) = & \frac{\rho F}{1-1/\psi}\left(\left(\frac{\chi}{F^{1/(1-\theta)}}\right)^{1-1/\psi}-1\right) + \\
& F\left(\varrho+\phi(\chi)-d(T)-\beta_t(\varepsilon(\alpha))-\theta\frac{\sigma_Y^2}{2}\right).
\end{aligned} \tag{B.11}$$

Game Problem

Similarly, the value function V_i (3.11) of region $i \in \{\text{OECD}, \text{RoW}\}$ is a function of time t, the two regions' temperatures T_i, the outputs Y_i, and the common CO_2 concentration m. The value functions satisfy

$$\begin{aligned} -\partial_t V_i = \sup_{\chi_i, \alpha_i} & f(\chi_i Y_i, V_i) + \partial_m V_i \left(\gamma_t^{\text{np}} - \alpha_i - \alpha_j\right) + \partial_m^2 V_i \frac{\sigma_m^2}{2} \\ & + \partial_{Y_i} V_i (\varrho + \phi(\chi_i) - d(T_i) - \beta_t(\varepsilon_i(\alpha_i))) + \partial_{Y_i}^2 V_i \frac{\sigma_Y^2}{2} \\ & + \partial_{T_i} V_i \frac{r_i(T_i) + g(m_i)}{\epsilon} + \partial_{T_i}^2 V_i \frac{(\sigma_T/\epsilon)^2}{2}, \end{aligned} \tag{B.12}$$

where α_j is the other region emission abatement policy. As above, one can check that the ansatz

$$V_{i,t}(T_i, T_j, m, Y_i, Y_j) = \frac{Y_i^{1-\theta}}{1-\theta} F_{i,t}(T_i, m) \tag{B.13}$$

satisfies (B.12). As the other region's policy α_j is assumed to be an elementary feedback strategy, the maximisation problem (B.12) can be solved taking that policy as given.

B.1.3 Approximating Markov Chain

Henceforth, the subscript i is omitted for notational simplicity and reintroduced only whenever the solution method between the social planner problem (B.2) and the regional game (B.12) differ.

The objective is to solve for the function F by adapting the method proposed in Kushner and Dupuis (2001) and Kushner (2007). I discretise the state space of temperature T and log-CO_2 concentration m by a fixed grid of size parametrised by some small step size h. I discretise the time t by computing state dependent intervals $\Delta t(T, m)$. Then I define a Markov chain $\mathcal{M}$ over the discretised space. Finally, I compute a discretised value function $F^h : \mathcal{X} \to \mathcal{X}$ with the property that

$$\sup_{x \in \mathcal{X}} \left| F^h(x) - F(x) \right| \to 0 \text{ as } h \to 0, \tag{B.14}$$

for any compact subset $\mathcal{X} \subset [T^p, \infty) \times [m^p, \infty) \times [0, \infty)$. For more details on convergence, see Chapter 15 of Kushner and Dupuis (2001).

First, I fix a suitable subset of the state space

$$\mathcal{X} := [T^p, T^p + \Delta T] \times [m^p, m^p + \Delta m]. \tag{B.15}$$

In practice ΔT and Δm are chosen such that T in the upper-right corner of the grid,

the drift $\mu(T^p + \Delta T, m^p + \Delta m) \equiv 0$. Then, given a step size $h > 0$ I define the grid over the space $\mathcal{X}$ by

$$\begin{aligned}\Omega_h(\mathcal{X}) = (T^p, T^p + h\Delta T, \ldots, T^p + (1-h)\Delta T, T^p + \Delta T) \times \\ (m^p, m^p + h\Delta m, \ldots, m^p + (1-h)\Delta m, m^p + \Delta m).\end{aligned} \tag{B.16}$$

Using the ansatz (B.3) I now define a discretised value function F_t^h over the grid such that $F_t^h \to F_t$ as $h \to 0$. First, introduce the discrete-time Epstein-Zin aggregator

$$J^h(\chi, F) = \left(\left(1 - e^{-\rho\Delta t}\right) \chi^{1-\frac{1}{\psi}} + e^{-\rho\Delta t} \left(\delta_y(\chi) F \right)^{\frac{1-\frac{1}{\psi}}{1-\theta}} \right)^{\frac{1-\theta}{1-\frac{1}{\psi}}}, \tag{B.17}$$

where

$$\begin{aligned}\delta_y(\chi) &:= \mathbb{E}_t \left[\left(\frac{Y_{t+\Delta t}}{Y_t} \right)^{1-\theta} \right] \\ &= 1 + \Delta t(1-\theta) \left(\varrho + \phi(\chi) - d(T_t) - \frac{\theta}{2}\sigma_k^2 \right) + \mathbb{E}_t \mathcal{O}\left(\Delta t^{\frac{3}{2}}\right).\end{aligned} \tag{B.18}$$

As the aggregator J^h satisfies Assumptions 1 through 4 of Kraft and Seifried (2014, p. 534), we have uniform convergence of the discrete-time operator to the continuous analogue, that is,

$$\sup_{F \in \mathcal{F}, \chi \in \mathcal{C}} \left| \frac{J^h(\chi, F) - J(\chi, F)}{h} - g(\chi, F) \right| \to 0 \text{ as } h \to 0, \tag{B.19}$$

for each compact $\mathcal{F} \subset \mathbb{R}$ and $\mathcal{C} \subseteq [0, 1]$. The discretised value function satisfies the recursion

$$F_t^h(T_t, m_t) = \min_{\chi, \alpha} J^h\left(\chi, \mathbb{E}_{t, \mathcal{M}(\alpha)} F_{t+\Delta t}(T_t, m_t)\right), \tag{B.20}$$

where $\mathbb{E}_{t,\mathcal{M}(\alpha)}$ is the expectation with respect to the Markov chain $\mathcal{M}(\alpha)$ over the grid $\Omega_h(\mathcal{X})$. The Markov chain $\mathcal{M}(\alpha)$, parameterised by h, is constructed as follows. Introduce the normalising factor

$$Q_t(T, m, \alpha) := \left(\frac{\sigma_T}{\epsilon \Delta T} \right)^2 + \left(\frac{\sigma_m}{\Delta m} \right)^2 + h \left| \frac{r(T) + g(m)}{\epsilon \Delta T} \right| + h \left| \frac{\gamma_t^{\text{np}} - \alpha}{\Delta m} \right|. \tag{B.21}$$

Then the probabilities of moving from a point $(T, m) \in \Omega_h(\mathcal{X})$ of the grid to an

adjacent point are given by

$$p(T \pm h\Delta T, m \mid T, m) \propto \frac{1}{2}\left(\frac{\sigma_T}{\epsilon \Delta T}\right)^2 + h\left(\frac{r(T) + g(m)}{\epsilon \Delta T}\right)^{\pm} \text{ and} \tag{B.22}$$

$$p(T, m \pm h\Delta m \mid T, m) \propto \frac{1}{2}\left(\frac{\sigma_m}{\Delta m}\right)^2 + h\left(\frac{\gamma_t^{\text{np}} - \alpha}{\Delta m}\right)^{\pm} \tag{B.23}$$

where $(\cdot)^+ := \max\{\cdot, 0\}$ and $(\cdot)^- := -\min\{\cdot, 0\}$. One can check that this is a well-defined probability measure. Finally, the time step is given by

$$\Delta t = h^2 / Q_t(T, m, \alpha), \tag{B.24}$$

which satisfies $\Delta t \to 0$ as $h \to 0$.

The Markov chain defined above allows us to derive $F_t^h(T_t, m_t)$ from the subsequent $F_{t+\Delta t}^h(T_{t+\Delta t}, m_{t+\Delta t})$. This requires a terminal condition

$$\bar{F}^h(T_\tau, m_\tau) := F_\tau^h(T_\tau, m_\tau). \tag{B.25}$$

To derive this, assume that at some point in a far future $\tau \gg 0$, the abatement is free and all emissions are abated, $\gamma^{\text{np}} \equiv \alpha$ for the social planner such that $\mathrm{d}m \equiv \sigma_m \, \mathrm{d}W_m$. Then I construct an equivalent, control independent, Markov chain $\bar{\mathcal{M}}$ as above for the steady state value function

$$\bar{F}^h(T, m) = \min_{\chi} J^h\left(\mathbb{E}_{t,\bar{\mathcal{M}}} \bar{F}(T, m); \chi\right). \tag{B.26}$$

This is now a fixed point equation for $\bar{F}$ which can be solved by value or policy function iteration. As the abatement control is bounded, $|\alpha_t| \leq \gamma_t^{\text{np}}$ the controlled drift

$$\begin{pmatrix} \mu(T_t, m_t) & \alpha_t - \gamma_t^{\text{np}} \end{pmatrix} \tag{B.27}$$

of the state vector (T_t, m_t) is also bounded, as well as continuous in the state vector. Furthermore, the aggregator J^h is bounded and continuous. This implies that the problem at hand satisfies Assumptions 10.1.1 and 10.1.2 for (weak) convergence given in Kushner and Dupuis (2001, p. 269). Hence, we have convergence of the discretised value function to the continuous analogue, $\bar{F}^h \to \bar{F}$. The steady state value function $\bar{F}$ serves then as a continuous and bounded terminal condition of the finite horizon problem (B.20). Which in turn implies (weak) convergence of $F_t^h \to F_t$ as $h \to 0$.

Markov Chain Approximation for the Game and Elementary Policies

As in the social planner problem, I assume that at some time $\tau \gg 0$ abatement is free and both regions abate all of their emissions $\gamma_i^{\text{np}} \equiv \alpha_i$ for both regions. At this time, the two regions do not interact, and the value function can be derived as in B.26.

For all $t < \tau$, one can define the discretised value function $F_{i,t}^h$ for the game B.13 using the aggregator J^h. The only difference with the social planner problem, is in the construction of the approximating Markov Chain, which will now depend on the abatement policies of both regions α_i, α_j. The approximating Markov Chain $\mathcal{M}_{\alpha_i,\alpha_j}$ for the game is defined by transition probabilities

$$p(T_i \pm h\Delta T, m \mid T, m) \propto \frac{1}{2}\left(\frac{\sigma_T}{\epsilon\Delta T}\right)^2 + h\left(\frac{r_i(T_i) + g(m)}{\epsilon\Delta T}\right)^{\pm}, \tag{B.28}$$

$$p(T_i, m \pm h\Delta m \mid T, m) \propto \frac{1}{2}\left(\frac{\sigma_m}{\Delta m}\right)^2 + h\left(\frac{\gamma_t^{\text{np}} - \alpha_i - \alpha_j}{\Delta m}\right)^{\pm}, \tag{B.29}$$

with time steps $\Delta t_i \propto h^2$ and normalisation factor

$$\begin{aligned} Q_{i,t}(T_i, m, \alpha_i, \alpha_j) \coloneqq & \left(\frac{\sigma_T}{\epsilon\Delta T}\right)^2 + \left(\frac{\sigma_m}{\Delta m}\right)^2 \\ & + h\left|\frac{r_i(T) + g(m)}{\epsilon\Delta T}\right| + h\left|\frac{\gamma_t^{\text{np}} - \alpha_i - \alpha_j}{\Delta m}\right|. \end{aligned} \tag{B.30}$$

As before, $\Delta t_i \to 0$ as $h \to 0$.

As mentioned in Section 3.2.3, each region chooses an elementary policy function α_i over a discretised time grid. Given the Markov chain discretisation introduced in this section, the time steps in which region i decides its policy depend on the sequence of time steps $\{\Delta t_{i,l}\}_{l=1}^n$ given by

$$t_{i,k} \coloneqq \tau - \sum_{l=1}^{n-k} \Delta t_{i,l}. \tag{B.31}$$

Given this discretisation, when choosing an abatement policy α_i for the interval $(t_{k_i-1}, t_{k_i}]$ of length Δt_{k_i}, the firm observes the other region's abatement policy and takes it as given, as illustrated in Figure B.3. As the time steps $\Delta t_{i,k_i}$ are computed independently and the two regions are asymmetric, it is unlikely that the regions would have to choose the policy contemporaneously, that is, often in practice $t_{i,k_i} \neq t_{j,k_j}$ for any two steps k_i and k_j. If this fails and $t_{i,k_i} = t_{j,k_j}$, region i assumes the other has already switched to the future policy at t_{i,k_j+1} when choosing their current policy at t_{i,k_i}, and vice versa.

Assuming elementary policy functions leads to two simplifications. First, within

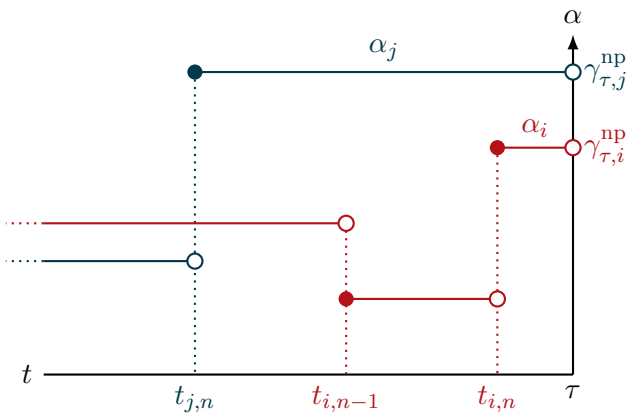

Figure B.3: Illustration of elementary policies as t approaches the terminal time τ.

an interval $(t_{k_i-1}, t_{k_i}]$ the problem of region i is a simple optimisation problem with the growth rate of the stock of CO_2 given by $\gamma_t^{\text{np}} - \alpha_j$, hence the convergence results discussed above remain valid. Furthermore, as the controls α_i take value in a compact set $[0, \gamma_i^{\text{np}}]$, Assumption 2.1 in Kushner (2007) is satisfies, hence F_i is well-defined and $F_i^h \to F_i$ uniformly as $h \to 0$. Second, the discretised game is just a finite sequential game in which players play at discrete times $\{t_{i,1}, t_{i,2} \dots t_{i,n_i}\}$. Hence, the sequence of policies yields a subgame perfect Nash equilibrium. This assumption might seem restrictive, as a region has to "wait its turn" before re-optimising its abatement policies. But notice that Δt_i is increasing in α_j such that, if region i chooses a laxer abatement policy, region j will be called to act earlier.

B.1.4 Parallelisation

When computing the backward recurrence (B.20), conditional on a given strategy α, each grid point $X_i \in \Omega^h(\mathcal{X})$ is assigned a different time step $\Delta t(X_i; \alpha)$, which depends on the curvature of the drift at that state. To parallelise the computation, I leverage the ZigZag algorithm by Bierkens, Fearnhead and Roberts (2019). Given the value function F_t^h, for a step back $t - \min_i \Delta t(X_i; \alpha)$, I construct a directed graph among grid points $\mathcal{X}$ where an edge $X_i \to X_j$ is drawn if there is a positive probability of transitioning from X_i to X_j under the Markov Chain $\mathcal{M}$. This allows to obtain, at each point in time t, sets of points $\mathcal{C}_t \subseteq \mathcal{X}$ which are independent and over which it is possible to parallelise. The parallelisation has been conducted on Snellius, the national high performance computer of the Netherlands. The algorithm is written in the `Julia` programming language and relies on `Optim.jl` developed by Mogensen and Riseth (2018) and the differential equation suit developed by Rackauckas and Nie (2017).

B.2 Certainty Equivalence

This appendix defines the certainty equivalence, which translate the net present values to monetary values in the context of our problem. Let

$$X_t := (T_t, M_t, N_t, Y_t) \tag{B.32}$$

be the state at time t and denote by $\overline{\mu}$ and $\underline{\mu}$ the drift functions of X_t under a remote and an imminent tipping point respectively, such that, in case of remote tipping,

$$\mathrm{d}\overline{X}_t = \overline{\mu}_t(X_t)\,\mathrm{d}t + \Sigma\,\mathrm{d}W_t\,, \text{ where} \tag{B.33}$$

$$\Sigma := \mathrm{diag}(\sigma_T/\epsilon, \sigma_m, 0, \sigma_k) \tag{B.34}$$

and $\mathrm{d}W_t$ is a standard Wiener process. Under these dynamics, one can obtain the net present value in utils of the consumption stream $\overline{V}_0$ and $\underline{V}_0$ satisfying equation (2.40). Similarly, one can compute the net present value of the consumption stream in utils of the wishful thinker (w) and prudent (p) policies as

$$V_0^i = \mathbb{E}_t \int_0^\infty f(Y_t \chi_t^i, V_t^i)\,\mathrm{d}t\,, \tag{B.35}$$

with $i \in \{w, p\}$. For each $V \in \{\overline{V}, \underline{V}, V^w, V^p\}$, the corresponding certainty equivalent CE solves

$$V_0^i = \int_0^\infty f(CE\, e^{\varrho t}, V_0^i)\,\mathrm{d}t\,. \tag{B.36}$$

B.3 Calibration and Parameters

This section summarises the parameters for the preferences, economy, and climate model and discusses the calibration strategy.

B.3.1 Economy and Base Climate Calibration

The following Table B.1 illustrates the preferences parameters used throughout the paper. There is no consensus in the literature on preference parameters. In line with previous literature focusing on recursive preferences, I set relative risk aversion $\theta = 10$ (Ackerman, Stanton and Bueno, 2013; Crost and Traeger, 2013; Lontzek et al., 2015) and the time preference parameter $\rho = 1.5\%$ (Nordhaus, 2014). Following the discussion in Hambel, Kraft and Schwartz (2021), I choose $\psi = 0.75$.

Preferences		
ρ	1.5%	Time preference
θ	10	Relative risk aversion
ψ	0.75	Elasticity of intertemporal substitution

Table B.1: Parameters of preferences.

For the calibration of the economy (Table B.2) and the baseline climate parameters (Table B.4), I follow the calibration suggested by Hambel, Kraft and Schwartz (2021, Section 3). I first calibrate the expected growth rate $\varrho + \phi(\chi_t)$ of output (2.39) in absence of abatement measures to match the DICE model through 2200. As in Hambel, Kraft and Schwartz (2021), as DICE is deterministic, I match the average simulation of my model to the DICE output. Following Pindyck and Wang (2013), the initial level of productivity $A(0) = 0.113$.

Economy		
ω_0	11%	GDP loss required to fully abate today
ρ_ω	2.7%	Rate of abatement cost reduction
ϱ	0.09% $[\mathrm{y}^{-1}]$	Growth of TFP
κ	6.32% $[\mathrm{y}^{-1}]$	Adjustment costs of abatement technology
δ_k	0.0116 $[\mathrm{y}^{-1}]$	Initial depreciation rate of capital
ξ	$2.6\mathrm{e}{-4}$ $[°\mathrm{C}^{-\nu}\mathrm{y}^{-1}]$	Coefficient of damage function
ν	3.25	Exponent of damage function
A_0	0.113	Initial TFP
Y_0	75.8 [trUS\$/y]	Initial GDP
σ_k	0.0162 $[\mathrm{y}^{-1/2}]$	Variance of GDP
τ	500 [y]	Steady state horizon

Table B.2: Parameters of the economic model.

Similarly, the regional economy (3.10) calibration matches the RICE output projection in the absence of abatement efforts. Abatement costs are assumed to be identical in both regions.

OECD		
ϱ	0.052% [y^{-1}]	Growth of TFP
A_0	0.13	Initial TFP
Y_0	47.54 [trUS\$/y]	Initial GDP
Rest of the World		
ϱ	0.4322% [y^{-1}]	Growth of TFP
A_0	0.09	Initial TFP
Y_0	28.25 [trUS\$/y]	Initial GDP

Table B.3: Parameters of the regional economies.

B.3.2 Climate

The parameters of the climate model are chosen to match the median temperature T_t^{np} and CO_2 concentration M_t^{np} in the median SSP5 (Kriegler et al., 2017) scenario in the no-policy scenario.

Climate		
T_0	1.14 [°C]	Initial temperature
M_0	410 [p.p.m.]	Initial CO_2 concentration
M^{p}	280 [p.p.m.]	Pre-industrial CO_2 concentration
N_0	286.66 [p.p.m.]	Initial carbon in sinks
σ_T	1.5844 [$\mathrm{y}^{-1/2}$]	Volatility of temperature
S_0	342 [$\mathrm{W\,m^{-2}}$]	Mean solar radiation
ϵ	15.844 [$\mathrm{J\,m^{-2}\,K\,y}$]	Heat capacity of the ocean
η	5.67e−8	Stefan-Boltzmann constant
G_1	20.5 [$\mathrm{W\,m^{-2}}$]	Effect of CO_2 on radiation
G_0	150 [$\mathrm{W\,m^{-2}}$]	Pre-industrial GHG radiation

Table B.4: Parameters of the climate model.

B.3.3 Feedback Mechanism

The transition function $L(T_t - T^c)$ of the feedback mechanism described in equation (2.22) takes the form

$$L(T_t - T^c) = \frac{1}{1 + \exp\left(L_1\left(T_t - T^c\right) + L_0\right)}. \tag{B.37}$$

See Bondarev and Greiner (2018) for a discussion on the effect of an alternative specification on the results of the optimisation problem.

The calibration of $\Delta\lambda$, L_0, and L_1 matches the upper bound of climate sensitivity in the AR6 WG1 report of 4°C, sourced from Byers et al. (2022). It is important to notice that these ensembles do not yield confidence interval based on statistical uncertainty but on model uncertainty, hence matching the distribution of the ensemble in a reduced model is not correct. See Guivarch et al. (2022) for a more detailed discussion. In line with this, the parameters are matched by simulating the path of temperature as described by (2.25) and estimating parameters via the `SciMLSensitivity.jl` package (Rackauckas et al., 2020). Table B.5 reports the median calibrated parameter. The critical threshold T^c cannot be estimated and is assumed for a particular scenario.

Imminent		
T^c	1.5 [°C]	Critical Temperature
L_0	3.15	Location Parameter
L_1	3.5 [°C^{-1}]	Speed Parameter
λ_1	0.31	Initial level of feedback
$\Delta\lambda$	0.0326	Magnitude of feedback
Remote		
T^c	2.5 [°C]	Critical Temperature
L_0	3.15	Location Parameter
L_1	3.5 [°C^{-1}]	Speed Parameter
λ_1	0.31	Initial level of feedback
$\Delta\lambda$	0.0332	Magnitude of feedback

Table B.5: Median calibrated parameters of feedback mechanism.

B.4 Stochastic Tipping Benchmark Model

This appendix introduces a benchmark model with stochastic tipping. The stochastic tipping model is a widely used in the economic literature to approximate tipping points in the climate dynamics (e.g. Hambel, Kraft and Schwartz 2021). Comparing the model developed in this chapter with the stochastic tipping model allows us to determine if and how the optimal abatement differ and, as a consequence, what the approximation misses.

To establish a meaningful benchmark, I will assume that the contribution of temperature to forcing (2.23) is given by

$$r_T^l(T_t) \coloneqq S_0(1-\lambda_1) - \eta\sigma T_t^4. \tag{B.38}$$

This model has no tipping point as $\lambda(T_t) \equiv \lambda_1$. Stochastic tipping, as commonly modelled in the literature, is introduced as a jump process J_t with arrival rate $\pi(T_t)$ and intensity $\Theta(T_t)$, both increasing in temperature. Intuitively, as temperature rises, the risk of tipping $\pi(T_t)$ and the size of the temperature increase $\Theta(T_t)$ grow. Then temperature dynamics in the Stochastic Tipping model follow

$$\epsilon \, \mathrm{d}T_t = r^l(T_t) \, \mathrm{d}t + g(m) \, \mathrm{d}t + \sigma_T \, \mathrm{d}W_{s,t} + \Theta(T_t) \, \mathrm{d}J_t \,, \tag{B.39}$$

where W_s is a Wiener process. Following Hambel, Kraft and Schwartz (2021), the calibrated arrival rate and temperature increase are calibrated as

$$\pi(T_t) = -\frac{1}{4} + \frac{0.95}{1 + 2.8e^{-0.3325(T_t - T_t^{\mathrm{P}})}} \qquad \text{and} \tag{B.40}$$

$$\Theta(T_t) = -0.0577 + 0.0568(T_t - T_t^{\mathrm{P}}) - 0.0029(T_t - T_t^{\mathrm{P}})^2. \tag{B.41}$$

B.5 Robustness Checks

B.5.1 Certainty Equivalence

This section checks the robustness of the cost of regret $(\overline{CE} - CE^c) - (\underline{CE} - CE^w)$ to the critical threshold T^c in the remote scenario and the coefficient of relative risk aversion θ. A larger relative risk aversion or a lower critical threshold increases the cost of regret. The cost of regret is always positive, hence, a cautious strategy is always cheaper than a wishful thinker one. Furthermore, the order of magnitude of the regret is stable across the different specifications.

		Remote T^c	
		2.5°	3.5°
θ	10	4.06trUS\$/y	3.08trUS\$/y
	2	1.79trUS\$/y	1.36trUS\$/y

Table B.6: Certainty equivalent robustness to different values of θ and the remote critical threshold.

B.5.2 Robustness to Remote Critical Threshold

This section illustrates the outcome of a cautious and wishful thinker scenario when the critical threshold in the remote scenario is chosen to be higher, namely at $T^c = 3.5°$.

Optimal Abatement Policies

First, I illustrate the robustness check on the optimal abatement policy, as defined in Section 2.4.

Figure B.4 shows the comparison between the two policy functions. The interpretation of the policy in the remote scenario is the same as the default calibration (Figure 2.12). Figure B.5 illustrates the resulting simulation. As in Figure 2.13, in the

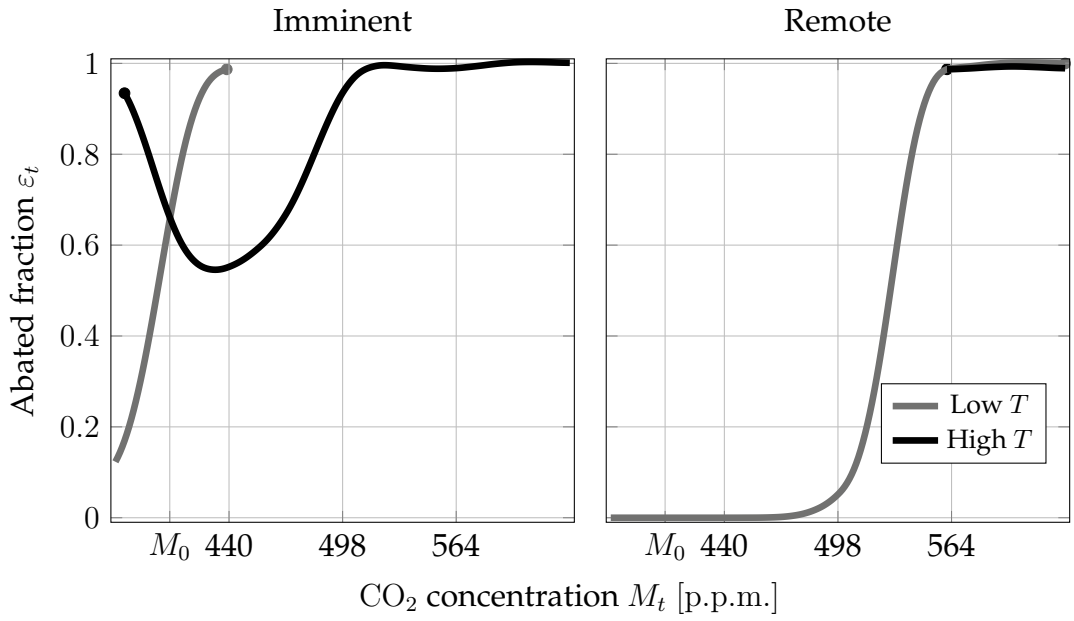

Figure B.4: Fraction of abated emissions ε_t at different levels of CO_2 concentration M_t, for an imminent (left panel) and remote (right panel) point. To aid illustration, all other dimensions (T_t, N_t, Y_t) are set to their equilibrium value. As for some values of M_t there are two equilibria of temperature T_t, two curves are shown. The current level of CO_2 concentration is indicated with M_0.

remote scenario, the tipping point is never reached and temperature is stabilised around 2°

As shown in Figure B.6, a higher critical threshold in the remote scenario is cheaper for society. Abatement measures can be postponed to the 2040s and climate damages remain minimal.

Wishful thinker and cautious planners

Figure B.7 shows the trajectory of the abated fraction of emissions and temperature, analogous to Figure 2.15. A higher critical threshold T^c is not sufficient for the wishful thinker planner to avoid tipping, yet results in a cheaper cost trajectory, as shown in Figure B.8, as the emissions are much lower when the tipping point is reached and the planner is not forced to slam the breaks.

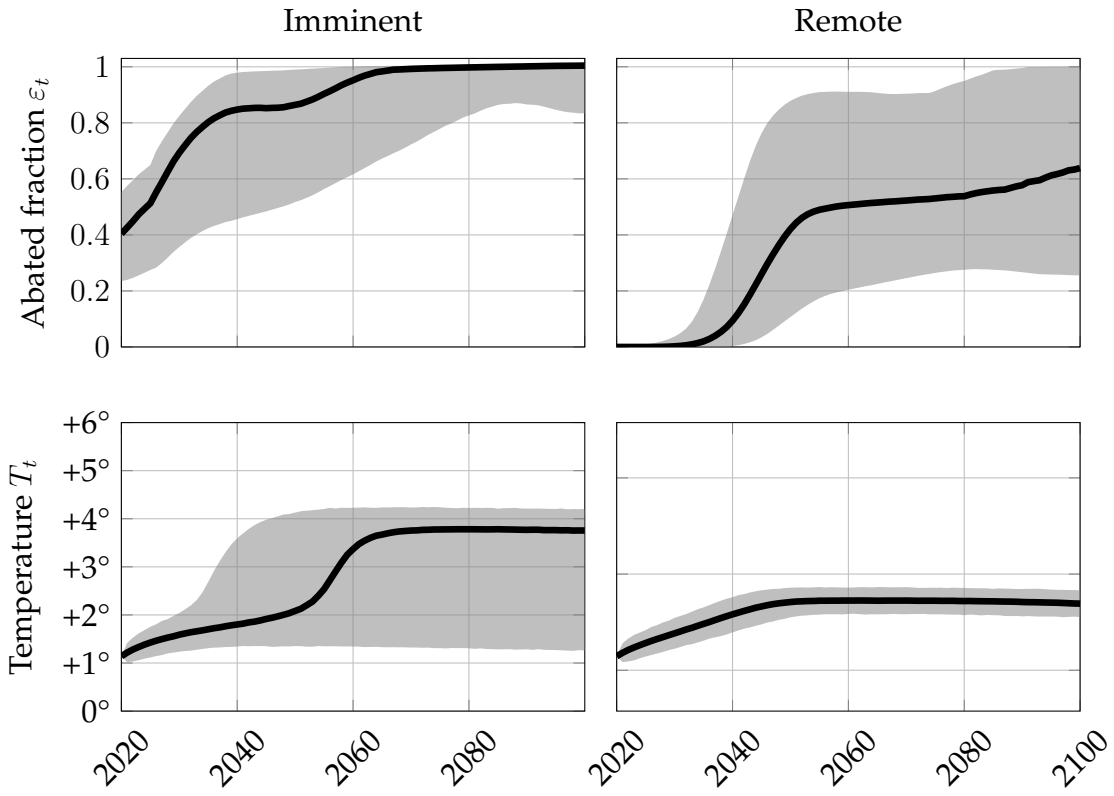

Figure B.5: Median (solid) and 95 % simulation intervals (shaded area) abated fraction of emissions ε_t and temperature T_t at time t of 100,000 simulations in an imminent tipping point (left column) and a remote tipping point (right column). The simulations are generated under the temperature (2.25), CO_2 concentration (2.18), and output (2.39) dynamics with optimal controls α and χ.

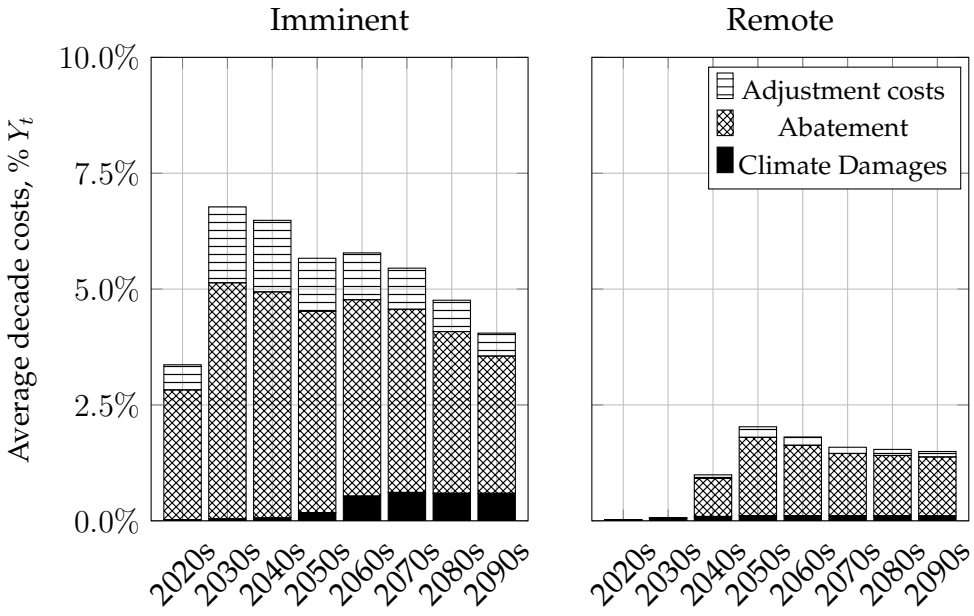

Figure B.6: Costs, as a fraction of output Y_t, in the median path illustrated in Figure 2.13, broken down in adjustment costs, abatement and climate damages.

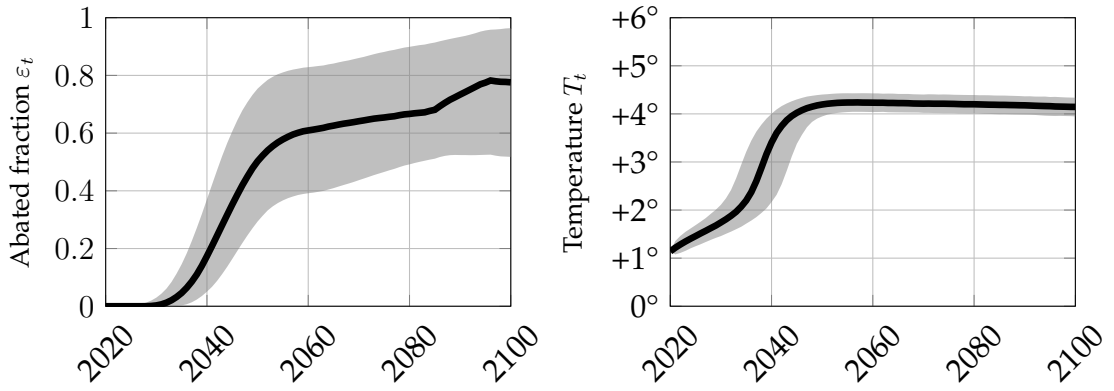

Figure B.7: Median (solid) and 95% simulation intervals (shaded) of the abated fraction of emissions ε_t (left panel) and the resulting the temperature T_t (right panel). Calculated from 100,000 simulations using policy (2.43).

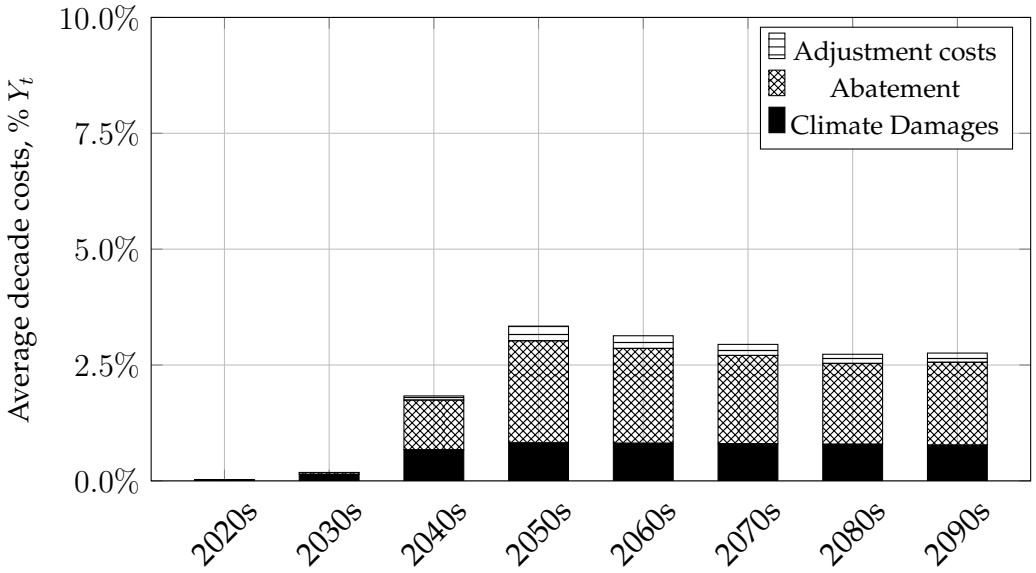

Figure B.8: Costs in the median scenario of a wishful thinker w planner as a fraction of output Y_t, broken down in adjustment costs, abatement and climate damages.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus University Rotterdam, University of Amsterdam and Vrije Universiteit Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. For a full list of PhD theses that appeared in the series we refer to https://www.tinbergen.nl/phd-theses. The following books recently appeared in the Tinbergen Institute Research Series:

823 H.P. LETTERIE, *Essays on the regulation of long-term care in the Netherlands*

824 D.D. PACE, *Essays on the cognitive foundations of human behavior and on the behavioral economics of climate change*

825 J.N. VAN BRUMMELEN, *On the estimation of parameters in observation-driven time series models*

826 Z. CSAFORDI, *Essays on Industry Relatedness*

827 B. VAN OS, *On Dynamic Models: Optimization-Based Methods and Practical Applications*

828 D.T. Ó CEALLAIGH, *Self-control Failures and Physical Inactivity: Measuring, Understanding and Intervening*

829 S.B. DONOVAN, *Ties that bind and fray: Agglomeration economies and location choice*

830 A. SOEBHAG, *Essays in Empirical Asset Pricing*

831 H. YUAN, *Essays in Behavioral Economics*

832 A. LENGYEL, *Essays on Government Bond Markets and Macroeconomic Stabilization*

833 S. KÜTÜK, *Essays on Risk Creation in the Banking Sector*

834 E. VLADIMIROV, *Essays on the Econometrics of Option Pricing*

835 R.E.K. PRUDON, *From the onset of illness to potential recovery. Empirical economic analysis of health, disability and work*

836 K. MOUSSA, *Signal Extraction by the Extremum Monte Carlo Method*

837 D. FAVOINO, *The Adaptation of Firms to Institutional Change*

838 B. WACHE, *Information Frictions in Financial Flows*

839 A. FEHER, *Essays in Law and Economics*

840 Q. WIERSMA, *Dynamic Models for Multi-Dimensional Time Series*

841 R. SILVESTRINJ, *On the Importance of Firm Heterogeneity, Business Dynamism, and Market Power Dynamics in the Macroeconomy*

842 E.S.R. DIJK, *Innovative Start-Ups and Competition Policy – How to Reign in Big Tech*

843 T.D. SCHENK, *Essays in Causal Inference with Panel Data*

844 S. TYROS, *Workers' Skills and (green) Technology Adoption*

845 C.J. GRASER, *Mechanisms for the Evolution of Prosociality*