

DISS. ETH NO. 28780

ESSAYS ON ENVIRONMENTAL GOVERNANCE
AND CLIMATE IMPACT UNCERTAINTY

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH
(Dr. sc. ETH Zurich)

presented by

Jere Lehtomaa

M.Sc. in Agriculture and Forestry, University of Helsinki,
born on 20 July 1990,
citizen of Finland,

accepted on the recommendation of
Prof. Dr. Antoine Bommier
Prof. Dr. Lucas Bretschger
Prof. Dr. Lars Grüne

2022

Thesis Summary

Chapter 1 studies the links between international trade agreements, environmental regulation, and welfare in the presence of politically influential firms. The motivation comes from the observation that after decades of gradual trade liberalization, import tariffs are already low in most sectors, with little room for further cuts. As a result, modern trade agreements increasingly focus on smoothing out other trade-restricting regulatory differences across countries. Such *deep trade integration* is often highly controversial due to concerns of a regulatory race to the bottom. This chapter builds a theoretical two-country model to analyze the strategic interactions between producers, policymakers, and industry lobbies. The results formalize how private-sector interference can lead to welfare-reducing trade agreements, especially if the lobby groups in different countries coordinate their efforts. However, policymakers can use different trade agreement designs to dilute the private sector's influence.

Chapter 2 focuses on the international governance of solar geoengineering. If deployed in a globally coordinated regime, interventions with geoengineering could reduce adverse climate change impacts. In the absence of such a regime, however, strategic incentives of individual actors might result in detrimental outcomes. This chapter investigates different international governance structures using a game of farsighted coalition formation. In this setting, a country pondering whether to leave or join a coalition anticipates that its decision could spark another (dis)integration process among the other players. In contrast to the static models of international environmental agreements, this dynamic structure enables a more realistic picture of what coalitions are likely to form and remain stable. The model also provides a unified framework for comparing different institutional arrangements for geoengineering deployment, such as consensus-based and majority-vote coalitions.

Chapter 3 develops a modelling framework for estimating the long-run economic impacts of tropical cyclones. The framework combines a numerical general equilibrium model of the economy with a probabilistic disaster modelling platform. Both components are global in their regional scope, which allows a consistent comparison across countries with different disaster risk profiles. The results highlight how the recovery after a single cyclone shock can take several decades, with the negative disaster impacts accumulating quickly for the frequently exposed regions. Assumptions regarding the drivers of economic growth and climate change's impact on future cyclone damages affect the numerical results. However, they do not change the overall qualitative conclusions.

Chapter 4 explores the applicability of learning-based model predictive control for solving stochastic economic decision problems. The high computational cost of sequential decision-making under uncertainty and the long planning horizons encountered in many economic applications severely restrict the usability of conventional dynamic programming techniques in high-dimensional settings. This chapter demonstrates that learning-based model predictive control provides a promising alternative approach as it can often deliver high-quality approximate solutions with a significantly smaller computational budget. In addition, stochastic model predictive control enables an intuitive formulation of *cautious* behaviours that complement the conventional risk-averse decision rules. An integrated climate-economy assessment model provides a challenging environment to evaluate the proposed method.

Zusammenfassung

In Kapitel 1 werden die Zusammenhänge zwischen internationalen Handelsabkommen, Umweltvorschriften und Wohlfahrt bei Vorhandensein von politisch einflussreichen Unternehmen untersucht. Die Motivation ergibt sich aus der Beobachtung, dass nach Jahrzehnten der schrittweisen Handelsliberalisierung die Einfuhrzölle in den meisten Sektoren bereits niedrig sind und wenig Spielraum für weitere Senkungen besteht. Infolgedessen konzentrieren sich moderne Handelsabkommen zunehmend darauf, andere handelsbeschränkende regulatorische Unterschiede zwischen den Ländern auszugleichen. Eine solch tiefgreifende Handelsintegration ist oft sehr umstritten, da ein regulatorischer Wettlauf nach unten befürchtet wird. In diesem Kapitel wird ein theoretisches Zwei-Länder-Modell entwickelt, um die strategischen Interaktionen zwischen Herstellern, politischen Entscheidungsträgern und Industrie-lobbys zu analysieren. Die Ergebnisse verdeutlichen, wie die Einmischung des Privatsektors zu wohlfahrtsmindernden Handelsabkommen führen kann, insbesondere wenn die Lobbygruppen in verschiedenen Ländern ihre Bemühungen koordinieren. Politische Entscheidungsträger können jedoch verschiedene Formen von Handelsabkommen nutzen, um den Einfluss des Privatsektors zu schwächen.

Kapitel 2 befasst sich mit der internationalen Koordination des solaren Geoengineering. Im Rahmen eines global koordinierten Systems könnte Geoengineering die negativen Auswirkungen des Klimawandels verringern. Ohne diese Art von Koordination könnten die strategischen Anreize einzelner Akteure jedoch zu nachteiligen Ergebnissen führen. In diesem Kapitel werden verschiedene internationale Governance-Strukturen anhand eines spieltheoretischen Modells mit vorausschauender Koalitionsbildung untersucht. Ein Land, das darüber nachdenkt, eine Koalition zu verlassen oder ihr beizutreten, antizipiert dabei, dass seine Entscheidung einen weiteren (Des-)Integrationsprozess bei den anderen Akteuren auslösen könnte. Im

Gegensatz zu statischen Modellen internationaler Umweltabkommen ermöglicht diese dynamische Struktur ein realistischeres Bild davon, welche Koalitionen sich wahrscheinlich bilden und stabil bleiben werden. Darüber hinaus bietet das Modell einen einheitlichen Rahmen für den Vergleich verschiedener institutioneller Rahmenbedingungen für den Einsatz von Geoengineering, wie z.B. konsensbasierte Koalitionen oder jene die auf Mehrheitsbeschlüssen beruhen.

In Kapitel 3 wird ein Modellierungsrahmen für die Abschätzung der langfristigen wirtschaftlichen Auswirkungen tropischer Wirbelstürme entwickelt. Die Methode kombiniert ein numerisches allgemeines Gleichgewichtsmodell einer Volkswirtschaft mit einer probabilistischen Katastrophenmodellierungsplattform. Beide Komponenten sind in ihrem regionalen Geltungsbereich global, was einen konsistenten Vergleich zwischen Ländern mit unterschiedlichen Katastrophenrisikoprofilen ermöglicht. Die Ergebnisse zeigen, dass die Erholung nach einem einzigen Wirbelsturm mehrere Jahrzehnte dauern kann, wobei sich die negativen Auswirkungen der Katastrophe in häufiger betroffenen Gebieten schnell akkumulieren. Die Annahmen zu den Triebkräften des Wirtschaftswachstums und zu den Auswirkungen des Klimawandels auf künftige Katastrophen beeinflussen die numerischen Ergebnisse, ändern jedoch nichts an den qualitativen Schlussfolgerungen.

In Kapitel 4 wird die Anwendbarkeit der lern-gestützten prädiktiven Modellsteuerung zur Lösung stochastischer wirtschaftlicher Entscheidungsprobleme untersucht. Der hohe Rechenaufwand der sequentiellen Entscheidungsfindung unter Unsicherheit und die langen Planungshorizonte, die in vielen wirtschaftlichen Anwendungen anzutreffen sind, schränken die Anwendbarkeit konventioneller dynamischer Programmierungsverfahren in hochdimensionalen Anwendungen stark ein. In diesem Kapitel wird gezeigt, dass auf maschinellem Lernen basierende Steuerungstechniken einen vielversprechenden alternativen Ansatz darstellen, da sie häufig qualitativ hochwertige Näherungslösungen mit einem deutlich geringeren Rechenaufwand liefern können. Darüber hinaus ermöglicht die stochastische prädiktive Modellsteuerung eine intuitive Formulierung vorsichtiger Verhaltensweisen, die die konventionellen Entscheidungsregeln des Risikoverhaltens ergänzen. Ein integriertes Klima-Wirtschafts-Bewertungsmodell bietet eine anspruchsvolle Umgebung für die Evaluation der vorgeschlagenen Methode.

Acknowledgments

Finishing a doctoral thesis is a team effort. Above all, I am grateful to my supervisor, Prof. Antoine Bommier, for his tireless guidance and the trust and freedom his students can enjoy exploring uncharted research topics. I would also like to thank Prof. Lucas Bretschger and Prof. Lars Grüne for evaluating my research and providing excellent feedback along the way. In addition, I am grateful to Prof. Hans Gersbach for chairing the examination committee.

During my thesis, I had the great privilege to work with Jean-Philippe Nicolai, Daniel Heyen, Clément Renoir, and Aleksei Minabutdinov on various research projects. These collaborations brought inspiration, improved my work, and deeply enriched my time at ETH. Thank you all for the invaluable help. Also, a special thanks to Adrien, Amélie, Arnaud, Aurore, Cheng, Claudio, Elias, Ewelina, Giovanni, Hélène, Irina, Julia, Markus, Moritz, Onur, Philippe, and Yen for being such fantastic colleagues and for your infinite intellectual and emotional support.

I am indebted to Juha Honkatukia and Prof. Markku Ollikainen for their initial encouragement in taking up this challenge. Without them, perhaps I would have never even started. Finally, I would like to thank Tyyra for her loving support when the task felt overwhelming. Without her, I know I would have never finished.

Jere Lehtomaa
Zürich, November 2022

Contents

Thesis Summary	I
Zusammenfassung	III
Introduction	1
1 Strategic Lobbying Under Deep Trade Agreements	6
1.1 Introduction	7
1.2 Theoretical framework	12
1.3 Comparison of trade agreements	20
1.4 Endogenous lobbying efforts	26
1.5 Conclusion	30
Appendices	33
1.A Proofs	33
1.B Parameter values	36
1.C Numerical results	37
2 Geoengineering Governance: A Coalition Formation Approach	38
2.1 Introduction	39
2.2 Theoretical framework	40
2.3 Numerical examples	42
2.4 Conclusion	48
Appendices	50
2.A Theoretical framework	50
2.B Geoengineering deployment equilibria	57
2.C Equilibrium strategy profiles	60
2.D Additional examples	61

3	Natural Disasters and Economic Dynamics	67
3.1	Introduction	68
3.2	Methodology	72
3.3	Results	79
3.4	Discussion	88
3.5	Conclusion	89
	Appendices	91
3.A	General equilibrium model	91
3.B	Additional results	107
3.C	Disaster impact distributions	110
3.D	Cyclone damages under climate change	110
4	Learning-based Control for Climate Policy Analysis	114
4.1	Introduction	115
4.2	Preliminaries	118
4.3	Environment	127
4.4	Experiments	133
4.5	Conclusion	146
	Appendices	148
4.A	Sampling-based MPC validation	148
4.B	Environment calibration	151
4.C	Implementation details	154
4.D	Soft constraints	155
	References	156
	Curriculum Vitae	176

Introduction

Human activity has shaped the atmosphere's chemical composition since the early years of 18th-century industrialization. At first unaware of the effects carbon emissions had on the Earth's temperature, the link has been established in the scientific literature for more than a century [1–3]. Today, with the mean global surface temperature already a degree above its preindustrial level, human-induced climate change is causing adverse effects such as more frequent and intense weather extremes, ocean acidification, ecosystem deterioration, food and water insecurity, and a higher incidence of vector-borne diseases [4]. Limiting the temperature increase to 1.5 °C, a target set in the 2015 Paris Agreement, requires rapid reductions in global emissions, reaching net zero around 2050 [5].

Despite the mounting evidence of climate change's adverse effects and the unequivocal role human activity has in causing it, the global policy response remains feeble. According to recent estimates, a complete and timely implementation of the updated Paris Agreement pledges might limit global warming to just below two degrees Celsius [6]. However, these trajectories contain substantial uncertainties, as the plausible peak temperature range extends from 1.4 °C to 2.8 °C. Other studies note that the remaining carbon budget for meeting the 1.5°C target will, at current rates, deplete in less than a decade. Even the unprecedented drop in economic activity due to the COVID-19 pandemic only brought a temporary reduction in global annual CO₂ emissions of less than 6%, highlighting the insufficiency of the currently agreed measures [7].

Basic economic theory suggests an elegant solution for regulating negative externalities such as carbon emissions: a tax equal to the externality's marginal social cost [8]. A carbon tax is desirable for its cost-effectiveness in reducing emissions and its simplicity when used to replace other overlapping climate policy instruments.

As with other Pigouvian taxes, pricing carbon would also generate government revenues for reducing distortionary taxes elsewhere, such as those on labour supply. Ideally, the carbon price should be uniform across economic sectors and world regions to minimize the overall cost of climate action.

In reality, however, implementing a global carbon price is notoriously difficult. First, there is no consensus on the adequate carbon price level. Given the uncertainties regarding climate system dynamics and climate change's future impact on human and ecological systems, finding the societal cost of each emitted CO₂ tonne is impossible. For instance, the monetary estimates are sensitive to the functional representation of climate damages in integrated climate-economy models [9]. There are also fundamental ethical considerations that affect the choice of the carbon price, such as how to discount the welfare of future generations vis-à-vis the immediate costs of reducing emissions [10]. As a result, the global climate policy landscape remains fragmented. There are currently 68 active carbon pricing instruments worldwide, covering approximately a quarter of global CO₂ emissions. The prices range from less than one US dollar up to USD 137 per tonne of CO₂ [11].

On top of the uncertainty over future climate impacts comes the difficulty of policy coordination in a world with power imbalances and conflicting interests. Designing meaningful climate policies while taking into account the vulnerabilities, capabilities, and responsibilities of sovereign states requires robust *governance* structures both on the national and international level [12]. The difficulty of international climate governance also implies that the relevant climate policy alternatives are usually not the most cost-efficient but those that are good enough and politically feasible. Unsuccessful governance efforts can lead to undesirable outcomes, for instance, to individual countries watering down consensus-based climate agreements [13], or in the future, to countries deploying geoengineering technologies unilaterally at the detriment of others [14]. The governance of non-state actors is also important, as private interest groups often exert substantial lobbying pressure on policymakers to steer environmental regulation [15, 16].

This thesis consists of two distinct thematical sections addressing the above discussion on climate policy design. The first two chapters focus on international environmental governance based on game-theoretic analysis. Chapter 1 applies the methodology to modern trade agreements, whereas Chapter 2 studies the gover-

nance of solar geoengineering deployment. The remaining two chapters concentrate on numerical climate impact assessment under uncertainty. Chapter 3 develops a framework for estimating long-term natural disaster impacts, and Chapter 4 explores learning-based control techniques for solving computationally intensive climate-economy integrated assessment models.

Chapter 1 studies the linkage of national environmental regulations and international trade agreements.¹ The motivation comes from the observation that after decades of gradual trade liberalization, import tariffs are already low in most sectors, with little room for further cuts. As a result, modern trade agreements increasingly focus on *deep trade integration*, that is, smoothing out other trade-restricting regulatory differences across countries. Deep trade integration typically covers areas such as intellectual property rights, environmental standards, and labour market rules [17]. Among the various regulations that can be subject to international coordination, those related to the environment often cause the most public controversy due to the concern of a regulatory race to the bottom.

Trade negotiations provide an interesting forum for environmental governance. First, the countries negotiating the agreement can be highly asymmetric in bargaining power and trade policy objectives, allowing for rich strategic interactions between the governments. Moreover, large firms often dominate the value of international trade flows [18], meaning that firms with market power might also respond strategically to trade policy outcomes. Finally, trade representatives do not negotiate the agreements in isolation but under constant pressure from various special interest groups [19]. The special interests can be as diverse as producer lobbies, environmental NGOs, or consumer protection agencies. Efficient governance ensures, for instance, that the presence of powerful private-sector lobbies does not lead to excessive environmental deregulation as a part of a new trade agreement.

The first chapter contributes to the growing literature on deep trade integration [20, 21] by considering a theoretical two-country trade model under imperfect competition and politically influential producer lobbies. The results provide insight into how the motives of policymakers and firms interact under various trade agree-

¹Chapter 1 is joint work with Jean-Philippe Nicolai. J.N and J.L contributed equally to conceptualizing the study. J.N formulated the modelling framework. J.L conducted the theoretical analysis and implemented the numerical simulations. Both authors contributed equally to writing and editing the manuscript.

ment designs. A key determinant of the results is whether lobby groups in different countries collude and coordinate their efforts in a joint push for weaker regulation.

Chapter 2 focuses on the international governance of solar geoengineering.² If deployed in a globally coordinated regime, interventions with solar geoengineering could reduce adverse climate change impacts. In the absence of such a regime, however, strategic incentives of individual actors might result in detrimental outcomes [14, 22, 23]. Here, *governance* refers to bridging the gap between potentially harmful and beneficial geoengineering deployment [24]. We model different international governance structures as a game of farsighted coalition formation [25, 26]. In this setting, a country pondering whether to leave or join a coalition anticipates that its decision could spark another (dis)integration process among the other players.

The main contribution of the second chapter is to introduce the coalition formation dynamics into the solar geoengineering context. In contrast to the static models of international environmental agreements, this dynamic structure enables a more realistic picture of what coalitions are likely to form and remain stable. The model also provides a unified framework for comparing different institutional arrangements for geoengineering deployment, such as consensus-based and majority-vote coalitions. Our results highlight that even small changes in the institutional setting – interpreted as the outcomes of active international governance efforts – can lead to a diverse set of equilibrium geoengineering deployment strategies.

Chapters 3 and 4 turn the thesis focus to integrated climate-economy modelling and numerical climate impact assessment. In particular, the chapters focus on economic and climate *risks*. Due to the partial understanding of several climate system processes and the unpredictable developments in the future world economy, incorporating uncertainties into the analysis is crucial for producing realistic insights for climate policymaking.

Chapter 3 studies the long-term economic impacts of natural disasters, focusing specifically on tropical cyclones. Over the past two decades, natural disasters have caused more than 50'000 deaths and almost USD 160 billion in economic damages per average year [27]. In addition to the direct effects through loss of lives

²Chapter 2 is joint work with Daniel Heyen. D.H. conceptualized the research. D.H. and J.L. contributed equally to the investigation and the formal analysis. J.L. wrote the code and conducted the computational analysis. D.H. and J.L. contributed equally to writing, reviewing and editing the work.

and destruction of economic assets, natural disasters can have far-reaching regional consequences in the globally interlinked economy [28]. Better understanding the long-term consequences of tropical cyclones is essential as they cause almost 40% of the total reported climate-related disaster damages and deaths [29].

The main contribution of the third chapter is to develop a framework that joins together a numerical general equilibrium model of the economy with a probabilistic disaster modelling platform.³ The results provide insight into how disaster strikes affect the economy directly through damaged capital stocks and how significant are the indirect effects as the disaster shock makes its way through multiple interlinked sectors and world regions. We quantify the disaster impacts on consumption and investment patterns over several decades and show that the adverse effects accumulate quickly, especially in the frequently affected areas. The framework also allows consistently analyzing how climate change might affect future disasters' frequency, intensity, and impact.

Finally, Chapter 4 focuses on computational methods and explores the applicability of learning-based model predictive control for solving stochastic climate-economy integrated assessment models. Sequential decision-making problems under uncertainty are computationally expensive. As a result, even relatively simple economic decision models with adequate uncertainty representation have often required supercomputer infrastructure [30] or careful pruning of the original problem [31]. This chapter demonstrates that learning-based predictive control [32–34] provides a promising alternative approach as it can often deliver high-quality approximate solutions with a significantly smaller computational budget. In addition, stochastic model predictive control enables an intuitive formulation of *cautious* behaviours that complement the previous works on risk-averse climate policy decision rules.

³Chapter 3 is joint work with Clément Renoir. Both authors contributed equally to the chapter.

Chapter 1

Strategic Lobbying Under Deep Trade Agreements

Abstract

After decades of gradual tariff cuts, few conventional trade barriers remain. As a result, the focus of new trade agreements has shifted from further tariff concessions to harmonizing other trade-affecting regulations between countries. We build a theoretical model to analyze how such deepening trade integration affects environmental regulation, lobbying, and welfare in a strategic two-country setting. We first derive analytical expressions for the optimal regulatory levels under various trade scenarios. Our results show strong incentives for international regulatory cooperation, but the effects can be welfare-reducing in the presence of politically influential firms. Then, we consider a model with endogenous lobbying efforts. Deep integration tends to increase lobbying relative to a conventional shallow agreement and can provide incentives for lobbies in different countries to coordinate, which is always detrimental to welfare. Perfect regulatory harmonization can improve welfare compared to partial convergence by weakening the influence of lobbies. However, this is always in conflict with the policymaker's objective.

This chapter is joint work with Jean-Philippe Nicolai.

1.1 Introduction

The multilateral trade rules signed under the GATT/WTO framework have been a significant driver of economic integration during the past decades. A central outcome of this process has been the continuous decline in global trade-weighted average tariff rates [35]. As tariffs' role has gradually diminished, coordinating other national regulations has become an increasingly important part of the international trade agenda. This deepening trade integration is apparent from the growing complexity of preferential trade agreements. In addition to conventional trade topics, they can also cover a wide range of other policy areas related to the *"free (or freer) movement of goods, services, capital, people, and ideas"* [17, p.8]. The EU has been actively promoting such deep trade agreements (DTAs), recently signing the Comprehensive Economic and Trade Agreement (CETA) with Canada and a similar deal with the South-American Mercosur trade bloc.

However, several concerns surround deep trade integration. One worry is that harmonizing regulations between countries could dilute environmental standards, consumer rights, and labour codes as governments struggle to keep local firms competitive against foreign rivals. Another fear is that DTAs might increase the influence of lobbies and other private sector interests. As Ref. [36] points out, the perception of lobbying is fundamentally different between shallow and deep agreements. Whereas lobbying for lower tariffs can be beneficial as it keeps protectionists in check, efforts to boost trade via regulatory adjustments might provide the lobby groups with unwanted clout. Although DTAs can well increase welfare by promoting trade, they are *"as likely to produce welfare-reducing, or purely redistributive outcomes under the guise of free trade"* [36, p.76]. The overall welfare impact of DTAs, therefore, remains unclear.

We explore the impacts of deep trade integration on environmental regulation, lobbying, and welfare. Although the regulatory and welfare implications remain largely ambiguous, our main findings suggest that DTAs tend to increase lobbying efforts relative to a shallow trade scenario. There are two main factors at play. First, the lobbying incentives depend on the magnitude of trade, so an increase in the overall trade volumes directly scales up the lobbying efforts. Second, regulatory cooperation across markets can incentivize firms in different countries to coordinate

and commit to high lobbying levels, which is always detrimental to welfare. The net welfare impact of DTAs then depends on whether the gains from trade expansion can outweigh the redistribution of surplus to politically powerful firms and the potential erosion of regulatory standards.

To analyze the role of DTAs, we build a theoretical model of international trade under imperfect competition and local consumption externalities. Imperfect competition is a central feature of the model for two reasons. First, it gives rise to intra-industry trade even between two perfectly identical economies, focusing our analysis on trade policy's strategic uses. Second, it describes the roles of market power, between-firm interactions, and positive profits in a simple setting. We consider them essential determinants of the firms' capability to behave strategically against policymakers.¹

In our model, two countries trade a single homogeneous good under Cournot competition. Both countries have access to one policy instrument: a minimum product standard for all goods consumed within their borders. The home government's policy then determines minimum requirements for the domestic market supply of local firms and the export supply of foreign firms. We first solve the model for a benchmark shallow agreement scenario where both countries set their local regulations non-cooperatively. We then analyze two versions of a DTA: one where countries cooperatively choose their own standards and another where they agree on one perfectly harmonized regulatory level. As the focus here is on optimal regulation, the first part of our analysis assumes all lobbying to be exogenous.

We show that the DTA's impacts on regulation and welfare depend on the integration's depth and the asymmetries between the two trading partners. In the special case where firms face no additional costs from supplying two markets with different varieties of their good, international regulatory cooperation always leads to more lenient policies in both countries. With positive compliance costs, however, the result no longer holds, and a high cost of complying with foreign standards can generate stricter policies to narrow the regulatory gap. A commitment to perfectly harmonize standards across markets constrains governments as they must choose over a single shared policy instrument. However, it can still increase aggregate

¹For instance, data from the US suggest that relatively few large companies do most of the lobbying related to trade agreements [19].

welfare by weakening the influence of lobbies and better aligning the policymaker's biased objective with the society's welfare function.

Finally, we modify our basic setup by endogenizing firms' lobbying decisions. Two patterns emerge from our results. First, as discussed above, deep trade integration consistently increases lobbying efforts compared to a shallow benchmark agreement. In addition, under all trade scenarios, firms in both countries can either coordinate their lobbying efforts or set the levels non-cooperatively. If the firms cannot sustain a cooperative outcome, the total lobbying expenditure remains relatively low. If, on the other hand, coordination persists, it leads to a sharp increase in lobbying efforts as the firms have shared interests in weakening the regulation in all destination markets. DTAs can increase the gains from coordinated lobbying, making trade integration more susceptible to private-sector manipulation.

Our results have a direct connection to the ongoing debate over the costs and benefits of deep trade integration. For instance, an independent French commission covering the impacts of CETA on the environment, climate, and health highlights the potential intervention of private interests as a risk of the trade agreement [37]. In particular, mechanisms included in the agreement, such as the Regulatory Cooperation Forum (RCF), that promote further voluntary coordination between interested parties, might work as vehicles for increased political interference. Our results formalize this risk by showing that DTAs can boost private-sector influence through higher lobbying expenditures and by incentivizing firms to coordinate lobbying efforts. The risks are not only limited to CETA, as the EU might use it as a template for other future trade agreements [38].

Related literature

Our paper builds on three main strands of earlier work. The first considers the substitutability between trade policies and domestic environmental regulation. The key idea is that countries with internationally competing firms might distort their local environmental rules to gain a competitive edge over their foreign trade partners [39]. This incentive becomes even more evident in the context of trade agreements where countries jointly commit to lowering their conventional trade barriers but might then resort to replacing them with other protectionist policies [40].

A substantial part of this literature focuses on trade *liberalization* and the impact of tariff cuts on optimal domestic regulation [41–44]. Our model with bilateral trade and imperfect competition is similar to this approach. However, with tariffs already sufficiently low, we depart from this earlier work by assuming that the most significant trade barriers arise from differences in national regulations. Instead of trade liberalization, we then focus on trade *integration*, where trade talks revolve around smoothing out these regulatory gaps.

An important exception in this literature is Ref. [45], which also considers a bilateral trade model with minimum product standards as the only available policy instrument. The authors define protectionist standards as ones that “*exceed[s] what a planner would impose if all producers were local*” (p.395). They show that the social planner always opts for protectionist minimum standards. Depending on the setting, firms prefer either completely removing domestic regulations or a regulation stringent enough to force the foreign rival out of the local market. We diverge from their approach by allowing firms to compete in both markets simultaneously, endowing all firms and policymakers with a symmetric action set. Our approach, therefore, generates different strategic interactions compared to their work.

The second relevant strand of literature focuses on the new generation of trade agreements and international regulatory harmonization. The closest to our work is Ref. [20], assessing the welfare impacts of deep integration. The authors consider both product standards and process regulations. They demonstrate that if producer lobbies are sufficiently powerful, agreements on product standards result in excessive deregulation and reduced welfare, while agreements on process standards strengthen regulations and improve welfare. Our imperfectly competitive two-country partial equilibrium model is complementary to their competitive multi-country and multi-sector setting. However, we only consider the implementation of a product standard. Both approaches reveal relatively similar findings, such as the aligned lobbying incentives and overly stringent non-cooperative standards in the absence of international compliance costs. However, our approach highlights the additional strategic elements from market power and fixed compliance costs, producing an alternative rationale for deep integration.

Ref. [46] builds a Ricardian trade model where countries differ in regulatory preferences. Trade only occurs if the gains compensate for the disparate views over

optimal regulation. Our approach is similar in that the sole purpose of trade agreements is to coordinate product standards under a local consumption externality. But whereas their work focuses on how regulatory preferences affect trade flows and the trade blocs that emerge, we emphasize the strategic interactions, market imperfections, and political economy forces that drive the stringency of regulation. Ref. [21] also considers a setting where consumers in different countries have dissimilar tastes, driving a gap between local regulations. Firms can supply several markets simultaneously but incur a fixed cost that depends on the width of the regulatory gap. The new trade agreement's central goal is to restore the global efficiency lost in the patchwork of national policies.

Third, our work is closely related to the existing literature on the political economy of international trade. Ref. [47] sets up a model where industry lobbies make political campaign contributions to the incumbent government in exchange for political favours. The policymaker optimizes over total political contributions and average voter welfare, both affecting the policymaker's re-election chances. A reduced-form version of their lobbying game, which we also follow in this paper, has been used extensively to model the biased objectives of a policymaker subject to lobbying pressure [48–51].

Earlier work suggests that the broader impacts of political influence depend not only on the interaction between policymakers and private sector groups but whether lobbies in other sectors [52] or countries [20] have similar or divergent interests. The interplay between lobbies is also a central determinant of the results in our model.

To summarize, our main contribution comes from the treatment of lobbying and political economy forces under different forms of deep trade integration. We show that the lobbying efforts and welfare impacts depend on design issues such as whether the trade agreement involves perfect harmonization of standards or only partial regulatory convergence. Although extensive literature exists on the political economy of trade agreements, the issue has received relatively little attention in models featuring deep integration, especially in an imperfectly competitive setting.

The remainder of this paper is structured as follows. Section 2 sets up the modelling framework for the shallow benchmark agreement and the DTA extensions. Section 3 presents our main findings under exogenous lobbying, and Section 4 presents a model with endogenous lobbying decisions. Section 5 concludes.

1.2 Theoretical framework

This section describes the bilateral trade model under imperfect competition, consumption externalities, and producer lobbying. The drivers of international trade come from the cross-hauling motives as in Ref. [53] and Ref. [54].

1.2.1 Main assumptions

Production and international trade. We consider two countries, A and B , with one firm in each country. Both firms produce a single homogeneous good X and supply it to both markets simultaneously. Firms compete à la Cournot. We assume all national policies to be non-prohibitive so that trade flows between the two countries and all domestic output quantities are strictly positive. X_j^i denotes the output level of the firm in country i destined to market j and Q_i the aggregate supply in market i . That is, $Q_A = X_A^A + X_A^B$ and $Q_B = X_B^B + X_B^A$. A linear inverse demand function characterizes the prices in each market with $p_i(Q_i) = \alpha_i - Q_i$. Consumer surplus is then $CS_i = Q_i^2/2$.

Minimum product standards. Consumption generates local damages. The maximum per-unit externality in country i is $\delta_i > 0$. To limit the externality, each government can impose a product standard μ_i . Consumers cannot distinguish between goods produced under different standards, so changes in μ_i have no effect on demand.² For producers, a stricter standard implies higher production costs. Therefore, absent the policymaker, firms would have no incentives to produce at higher standards. To regulate the firms, policymakers in both countries implement a minimum standard $\bar{\mu}_i$ that all products within their market area must satisfy. Each country then controls the products consumed but not those manufactured within its national borders. We assume the regulation to be non-discriminatory: countries are free to treat imported goods with any regulation necessary for human or environmental well-being as long as domestic producers face the same requirements.

²A weaker interpretation of this assumption is that a regulatory standard, for instance, on vehicle exhaust particle limits, is not an important determinant of the purchase decision for most consumers.

Technology and production costs. The unit production costs are constant in output quantity but increase linearly with the minimum standard. The cost for the firm in country i of meeting the standards imposed in market j is $c_i(\mu_j) = \phi_i \mu_j$, where $\phi_i > 0$ captures the firm's technological efficiency. For the two firms to produce positive amounts, we bound the cost differences to satisfy $2\phi_i > \phi_j$. The marginal abatement efficiency $\omega_i > 0$ is constant, so that the effective per-unit externality in country i is $\delta_i - \omega_i \mu_i$. To keep the effective per-unit externality positive, we consider $\mu_i \in [0, \frac{\delta_i}{\omega_i}]$.

Fixed trade costs. In addition to the production costs, we consider a trade cost that is constant in terms of output quantity but increases in the distance between domestic and export market regulation. The term broadly captures the costs incurred by firms from complying with two different standards, whether through the cost of maintaining separate production lines or from additional red tape at border crossings. Specifically, we assume that each firm carries a fixed cost of $\psi_i(\mu_i, \mu_j) = (F_i/2)(\mu_i - \mu_j)^2$ for some $F_i \geq 0$. Lowering these fixed costs is a central political motivation for deep trade integration.

Damages. The total externality in country i is a function of aggregate consumption and the effective per-unit damages, $D_i = (\delta_i - \omega_i \bar{\mu}_i) Q_i$. Our implicit assumption is that externalities such as small particle emissions from vehicles or pesticide traces in food products are similarly harmful in both regions. However, we might also extend the interpretation of δ_i to wider gaps in *perceived* damages. For instance, the issue of chlorine-washed chicken, banned in the EU but common practice in the US, marks a symbolic difference between the US and EU legislations and is an oft-cited cause for public uproar against DTAs in Europe. We also assume that the unit externality is sufficiently low relative to the demand function intercept, $\alpha_i > \delta_i$. This rules out situations where the marginal damage of consumption would exceed the marginal benefits.

Political influence and welfare. In addition to market power, firms in our model possess political influence over the regulators. We follow an approach similar to the *Politically Realistic Objective Functions* [55], which is commonly used in

two-country trade models [20, 48–51, 56].³ Under this setting, the lobbying game between private interests and policymakers reduces to placing additional weights on profits π in the standard government objective function:

$$\Psi_i = CS_i + \gamma_i^i \pi_i^i + \gamma_j^i \pi_j^i - D_i. \quad (1.1)$$

The policymaker’s bias towards firm i ’s market j profits in Eq. (1.1) is captured in the terms $\gamma_j^i \geq 1$. Throughout the paper, we make a distinction between the (biased) objective Ψ_i and the (real) welfare $W_i = CS_i + \pi_i^i + \pi_j^i - D_i$. We start by treating all influence terms γ as exogenous parameters. Section 1.4 extends our analysis to firm-specific endogenous lobbying activity. In that case, firms can choose their (costly) lobbying efforts to influence the weights associated with their profits.

Timing. We solve the benchmark model as a two-stage game. In the first stage, both governments set their minimum standards. In the second stage, firms in both countries simultaneously choose their profit-maximizing output quantities for all destination markets.

Segmented market profits. The markets in our partial equilibrium economy are perfectly segmented, and firms can independently make output decisions for each destination market. We can then divide each firm’s profits into *i*) the local component where the products meet the domestic regulation and *ii*) the export component where the products meet the standard imposed by the foreign government. One helpful interpretation of this breakdown of profits is that there is one firm in each country only focusing on domestic operations and another firm focusing on exports. The profits for firms A and B are, respectively:

³Ref. [47] provides explicit microfoundations for this approach.

$$\begin{aligned}
\pi^A(X_A^A, X_B^A) &= \left[\alpha_A - (X_A^A + X_B^A) - c_A(\mu_A) \right] X_A^A \\
&\quad + \left[\alpha_B - (X_B^B + X_B^A) - c_A(\mu_B) \right] X_B^A - \psi_A(\mu_A, \mu_B), \\
\pi^B(X_B^B, X_A^B) &= \left[\alpha_B - (X_B^B + X_A^B) - c_B(\mu_B) \right] X_B^B \\
&\quad + \left[\alpha_A - (X_A^A + X_A^B) - c_B(\mu_A) \right] X_A^B - \psi_B(\mu_B, \mu_A).
\end{aligned} \tag{1.2}$$

Equilibrium trade volumes. Each firm maximizes profits by choosing production levels for the domestic and export markets. We determine the firms' best-response functions from the first-order optimality conditions to the system in Eq. (1.2).⁴ The Cournot-Nash equilibrium quantities are:⁵

$$\begin{aligned}
X_A^A &= \frac{1}{3} \left[\alpha_A - 2c_A(\mu_A) + c_B(\mu_A) \right], \\
X_B^A &= \frac{1}{3} \left[\alpha_B - 2c_A(\mu_B) + c_B(\mu_B) \right], \\
X_A^B &= \frac{1}{3} \left[\alpha_A - 2c_B(\mu_A) + c_A(\mu_A) \right], \\
X_B^B &= \frac{1}{3} \left[\alpha_B - 2c_B(\mu_B) + c_A(\mu_B) \right].
\end{aligned} \tag{1.3}$$

A higher destination market demand α_i increases the equilibrium output quantities both for the domestic supplier and the foreign exporter. In each country i , higher production costs for the foreign exporter, c_j , will reduce imports X_i^j and simultaneously boost domestic production X_i^i . Therefore, the outcome of a foreign cost increase is the domestic producer controlling a larger share of the local market. However, the total market supply Q_i will remain below the original level, as the marginal decrease in foreign exports will outweigh the boost to domestic production. That will also increase the local prices p_i .

On the other hand, an increase in the stringency of national minimum standards μ_i will cut both domestic production and foreign export supply. When the marginal costs ϕ_i are identical between the firms, this reduction will be equivalent in both

⁴The second derivatives are everywhere constant and strictly negative.

⁵The uniqueness of the Nash equilibrium follows from the linear demand schedules.

countries. However, when the foreign marginal costs exceed those of the domestic producer, the impact will fall disproportionately on the foreign producer, again securing the local firm a larger share of the domestic market.

Trade policy scenarios. To study the regulatory outcomes of DTAs as well as their different possible designs, we compare the three following trade scenarios:

1. **Business as usual:** Each government maximizes its national objective and sets the minimum standard non-cooperatively. In other words, countries are open to trade, but there is no coordination over the regulatory levels.
2. **DTA with regulatory cooperation:** Governments coordinate over their (separate) regulations by maximizing a joint objective.
3. **DTA with full regulatory harmonization:** Governments maximize a joint objective but commit to imposing equally stringent minimum standards.

1.2.2 Business as usual

The first scenario corresponds to an existing shallow trade agreement where countries trade in the absence of tariff policies but the lack of regulatory cooperation restricts the total traded quantity. Solving for $\frac{\partial \Psi_i}{\partial \mu_i} = 0$ gives for each country i the following optimal standard:

$$\bar{\mu}_i^{BAU} = \frac{\Lambda_j^{BAU} F_i \gamma_j^i - \Lambda_i^{BAU} \Xi_j^{BAU}}{9 [\Xi_i^{BAU} \Xi_j^{BAU} - F_i F_j \gamma_j^i \gamma_i^j]}, \quad (1.4)$$

where we denote:

$$\begin{aligned} \Lambda_i^{BAU} &= 3\delta_i \Phi - 2\alpha_i (\Phi - 3\omega_i + \varphi_i \gamma_i^i) > 0, \\ \Xi_i^{BAU} &= \frac{1}{9} [\Phi (\Phi - 6\omega_i) + 2\varphi_i^2 \gamma_i^i - 9F_i \gamma_j^i] < 0. \end{aligned}$$

The symmetric terms Λ_i correspond to the country i 's home market considerations and are positive under our parameter assumptions. The term trades off local

damages δ_i , local demand α_i , and the local firm's domestic profits through the influence-weighted cost (dis)advantage $\varphi_i = (2\phi_i - \phi_j)$. The joint production cost term is $\Phi = \phi_A + \phi_B$. The term Ξ_i corresponds to government i 's second-order optimality condition and is strictly negative by assumption.

A high value for Λ_i roughly implies stricter domestic regulation either due to high externalities, low demand, or high demand under relatively efficient abatement technology ω_i . If there are costs for complying with foreign standards, higher foreign Λ_j also drags up the market i regulation to narrow the regulatory gap. Intuitively, the higher the policymaker's valuation of export profits, as captured in γ_j^i , the stronger the foreign market effect.⁶

The reaction functions for both governments are linear in the foreign regulatory stringency, with the slopes determined by the fixed compliance costs and the influence of the domestic export lobby. In the absence of compliance costs, the connection between $\bar{\mu}_i$ and $\bar{\mu}_j$ in the non-cooperative scenario vanishes, and foreign regulation no longer has an impact on domestic standards.

It is worth emphasizing some particular features of our model economy with two market failures – consumption externalities and imperfect competition – and only a single policy variable at the disposal of each government. In such a setting, the regulatory level might be efficient relative to the incomplete policy toolbox but remain suboptimal compared to other efficiency benchmarks. When adjusting its regulation, each government then has to balance between *i*) increasing standards to reduce local damages, *ii*) decreasing standards to boost the (possibly) insufficient duopoly supply, *iii*) the strategic use of regulation to shift economic rents from the foreign producers to the domestic firms, and *iv*) taking into account the political clout of all the relevant special interest groups.

We next present some comparative statics of the optimal regulation in Eq. (1.4) to analyze these different considerations in more detail. First, as expected, the optimal domestic standard is always strictly increasing with local per-unit damages, $\partial\bar{\mu}_i/\partial\delta_i > 0$.⁷ On the other hand, higher domestic demand α_i can have either a

⁶The effect is similar to the *California effect* coined by Ref. [57] for situations where exporters voluntarily adopt the stricter standards enforced in an important destination market.

⁷We rule out situations where a government would regulate the domestic market *only* to manipulate the compliance cost term in its favour. In other words, the home government will always impose a strictly positive minimum standard even in the absence of compliance costs and foreign

positive or negative effect on standards. If the policymaker attaches a high weight to domestic profits or the home firm has a cost advantage and high abatement efficiency, the rent-shifting effect dominates, and the government can impose stricter local standards. Otherwise, higher domestic demand implies that boosting supply and upholding consumer surplus generates a more significant improvement in the government objective than cutting consumption externalities, resulting in more lenient regulation.

The same reasoning applies to the abatement cost terms ϕ and compliance costs F . If the domestic demand and cost advantages are sufficiently high, an uptick in foreign costs can make local standards stricter, cutting market access from the foreign producer.

Finally, consider the marginal impact that lobbying has on optimal regulation. Even when there are fixed compliance costs, the import-competing sector will always push for relaxing the domestic regulation as it will always result in strictly lower production costs.⁸ As a result, a firm's growing political influence always translates into higher profits and a jump in aggregate supply. However, the marginal returns from lobbying are diminishing, as lobbying for weaker regulation at home will also increase the foreign rival's export supply. The export lobbies' influence on domestic policies depends on foreign regulation. A higher weight on export profits motivates the home government to narrow the international regulatory gap, potentially leading to stricter home standards.

regulation. In effect, this constrains the allowed distance between the production costs ϕ_i and the abatement efficiency ω_i .

⁸Since we assume all trade flows to be positive, the home firm cannot lobby for a regulatory gap wide enough to push the foreign rival out of the market.

1.2.3 Deep trade agreements with regulatory cooperation

Consider next international regulatory cooperation under the joint government objective $\Psi = \Psi_A + \Psi_B$. The timing of this game is as follows:

First stage: Governments set their optimal national standards maximizing the joint objective.

Second stage: Firms choose their profit-maximizing output quantities for all markets simultaneously.

We solve the game using backward induction. Since the consumption externality is purely local, and governments can only regulate the goods consumed within their borders, one country's minimum standard will not directly affect the damages occurring in the other. Therefore, the only change relative to the non-cooperative benchmark is that policymakers now consider how their domestic rules affect foreign export profits. Solving for $\frac{\partial \Psi}{\partial \mu_i} = 0$ yields for each country i :

$$\bar{\mu}_i^{DTA} = \frac{\Lambda_j^{DTA}(F_i\gamma_j^i + F_j\gamma_i^j) - \Lambda_i^{DTA}\Xi_j^{DTA}}{9[\Xi_i^{DTA}\Xi_j^{DTA} - (F_i\gamma_j^i + F_j\gamma_i^j)^2]}, \quad (1.5)$$

where we again denote:

$$\begin{aligned} \Lambda_i^{DTA} &= 3\delta_i\Phi - 2\alpha_i(\Phi - 3\omega_i + \varphi_i\gamma_i^i + \varphi_j\gamma_i^j), \\ \Xi_i^{DTA} &= \frac{1}{9}[\Phi(\Phi - 6\omega_i) + 2\varphi_i^2\gamma_i^i - 9F_i\gamma_j^i - (9F_j - 2\varphi_j^2)\gamma_i^j]. \end{aligned}$$

As expected, Eq. (1.4) and Eq. (1.5) only differ in how the foreign firm's compliance costs F_j , cost (dis)advantage φ_j , and export sector influence γ_i^j affect optimal regulation. Trade integration, in effect, removes each government's incentives to manipulate standards for rent-shifting purposes and results in a reciprocal exchange of market access.

1.2.4 Deep trade agreements with perfect harmonization

In our final trade policy scenario, governments agree on one perfectly harmonized regulatory level across markets.⁹ We set up a game with the following timing:

First stage: Governments set a minimum standard $\bar{\mu}_A^{PER} = \bar{\mu}_B^{PER} = \bar{\mu}^{PER}$ maximizing the joint objective $\Psi = \Psi_A + \Psi_B$.

Second stage: Firms in both countries set their profit-maximizing output quantities for all markets.

Again, we use backward induction to solve the game. Setting $\frac{\partial \Psi}{\partial \mu} = 0$ gives:

$$\bar{\mu}^{PER} = \frac{\Lambda_i^{DTA} + \Lambda_j^{DTA}}{-9\Xi^{PER}}, \quad (1.6)$$

where, using $\Omega = \omega_A + \omega_B$ for an aggregate abatement efficiency term, we denote:

$$\Xi^{PER} = \frac{2}{9} [\Phi(\Phi - 3\Omega) + \varphi_i^2(\gamma_i^i + \gamma_j^i) + \varphi_j^2(\gamma_j^j + \gamma_i^j)].$$

The denominator in Eq. (1.6) has a very intuitive form where the first term trades off the aggregate production costs and abatement efficiencies, and the other terms weigh each firm's relative production costs φ according to its total political influence. Under the special case of two perfectly identical economies, Eq. (1.5) and Eq. (1.6) exactly coincide. Section 1.3 analyzes the regulatory differences and the resulting welfare impacts in more detail.

1.3 Comparison of trade agreements

This section compares the regulatory, lobbying, and welfare outcomes under the above trade scenarios. Appendix 1.A contains all proofs.

⁹The previous scenario with country-specific cooperative standards can also produce a perfect regulatory convergence across markets. However, as this only happens in special cases, a separate scenario for perfect harmonization is warranted.

1.3.1 Regulation

The regulatory impacts of deep trade integration hinge on the asymmetries between the trade partners and the depth of their integration. In the special case without export compliance costs, we find the following ranking of regulatory levels:

Proposition 1 *In the absence of export costs, the following rankings always hold:*

1. For any individual country k , $\bar{\mu}_k^{DTA} \leq \bar{\mu}_k^{BAU}$.
2. For any trading partners i and j , $\bar{\mu}_i^{DTA} \leq \bar{\mu}^{PER} \leq \bar{\mu}_j^{DTA}$.
3. The ranking between $\bar{\mu}_k^{BAU}$ and $\bar{\mu}^{PER}$ is ambiguous.

We can summarize Proposition 1 as follows. The regulation is always more lenient for both countries under the regulatory cooperation scenario (DTA) than under the shallow benchmark case (BAU). In addition, the regulation under perfect harmonization (PER) is always a weighted average of the standards imposed under the DTA scenario with partial convergence. For the movement from BAU to PER, two outcomes are possible. First, if sufficiently similar, the model produces a regulatory race to the bottom where both countries loosen their standards. On the other hand, it is possible for one of the countries first to relax its domestic standards when moving from BAU to DTA but then tighten the regulation in PER even beyond the original benchmark level.

Compare first Eq. (1.4) and Eq. (1.5) that give $\bar{\mu}_i^{DTA} \leq \bar{\mu}_i^{BAU}$ for all countries. As the policymakers under deep integration switch to maximizing a joint objective, both ignore their original incentive of shifting surplus across markets by manipulating regulation.¹⁰ Therefore, the outcome is a reciprocal exchange in market access through a reduction in local standards. The result holds even in the absence of special interests γ , and the effect becomes even stronger when lobbies are present.

Consider next the two cooperative outcomes in Eq. (1.5) and Eq. (1.6). If the countries are perfectly identical, the optimal standards under DTA and PER coincide, and $\bar{\mu}^{PER} = \bar{\mu}_i^{DTA}$ for $i = A, B$. The finding is intuitive: if two cooperating

¹⁰The implicit assumption behind a joint government objective is the existence of a perfect transfer mechanism between countries.

countries move from distinct policies to harmonizing their regulations, the new policy will be some weighted average between the pre-existing regulatory levels. Such averaging will not affect identical economies, and the regulation remains unchanged.

For a concrete example, consider two economies that only differ in their per-unit damages δ_i . As noted earlier, higher externalities will, other things equal, always tighten the optimal standards. Therefore, harmonizing regulations means that the country with higher externalities will loosen its regulation. More formally, we have that $\mu_i^{DTA} < \mu^{PER} < \mu_j^{DTA}$ when $\delta_j > \delta_i$. The intuition is clear: for the country with greater damages, harmonizing standards with a less sensitive trade partner corresponds to a setting where one large economy with lower *average* damages determines its regulation alone. Since the optimal policy is monotonically increasing in damages, the resulting harmonized minimum standard is below the DTA level for country j .¹¹

Introducing the fixed compliance costs drastically changes the results in Proposition 1. Especially the strong prediction regarding the regulatory drop between BAU and DTA no longer holds. Instead, it might well be that one of the countries has considerable externalities, strict regulation, and high export compliance costs. In such cases, it can be more profitable for the foreign policymaker to shift its minimum standard upward for a mutually beneficial compromise that maximizes the joint objective.

Proposition 1 also suggests that some lobbying can even be desirable under the shallow BAU scenario. The non-cooperative governments might set excessively stringent standards at the expense of the foreign country. The producer lobbies pushing for weaker regulation can then act as a counterweight to the strategic government behaviour.

The ambiguity over the regulatory outcomes highlights how the agreement's design details and country heterogeneity are crucial determinants of how deep trade integration affects domestic policies. Asymmetric countries are also more likely to see asymmetric effects. A new trade agreement might cause significant regulatory changes for one of the trading partners but only minor adjustments for the other.

¹¹For country i , the tighter regulation also means a cut in total supply Q_i . Therefore, if the ultimate goal of a DTA is to increase trade flows, the perfect harmonization of standards can even be counter-productive unless there are other cost reductions for firms from the unified regulation.

1.3.2 Lobbying

The lobbying motives vary across trade scenarios and the import-competing and exporting sectors. The influence that the import-competing sector has on domestic regulation is clear:

Proposition 2 *Import-competing firms always lobby for weaker local regulation.*

That the above proposition always holds is somewhat surprising. One might imagine a scenario where the local firm with a sufficient cost advantage over its foreign rival would benefit from stricter domestic regulation to push the competitor out. However, as we restrict our analysis only to scenarios with strictly positive trade flows, the local lobby can never achieve a regulatory gap wide enough to dominate the market entirely. Since this strategic channel is never available, the home firm is always better off with the cost reduction from weaker regulation.

The situation looks different for the exporting sectors. As the production costs depend on the foreign standard, the exporters only care about domestic regulation indirectly through the export compliance costs. Therefore, a higher weight for export profits in the benchmark scenario always drives convergence in international standards. With a tightly regulated export market, the firms might even lobby for stricter domestic regulation. However, the link between export sector influence and domestic regulation disappears in the absence of export costs.

International regulatory cooperation introduces new incentives for the export lobbies. As long as *any* of the firms face strictly positive export compliance costs, the export sector influence will contribute to the local standards. If only the foreign firm faces export costs, the marginal impact of greater local export influence on domestic regulation is negative through the following chain of events. The local exporter does not mind the regulatory gap, and a higher weight on its profits in the joint government objective will result in a drop in foreign standards. However, this might bring additional costs to the foreign firm due to a wider regulatory gap. To compensate, the home government will make a reciprocal downward adjustment to local regulation to avoid any changes in foreign costs. If both firms face strictly positive compliance costs, the marginal impact is unclear and depends on whether the compliance or production cost terms dominate the effects.

The perfect harmonization of standards removes all considerations regarding fixed trade compliance costs. As a result, all firms prefer lower standards in any market they supply. Therefore, perfect regulatory harmonization effectively turns the trade area into a single regulated world market. We collect these results to the following proposition:

Proposition 3 *The export sector's political influence has the following effect on domestic regulation:*

1. *In BAU, export lobbies push domestic regulation towards closing the international regulatory gap whenever they have positive trade costs. More formally:*

$$\text{sign}\left(\frac{\partial \bar{\mu}_i^{BAU}}{\partial \gamma_j^i}\right) = \begin{cases} \text{sign}\left(\bar{\mu}_j^{BAU} - \bar{\mu}_i^{BAU}\right), & \text{if } F_i > 0 \\ 0, & \text{otherwise} \end{cases}$$

2. *In DTA, the direction of export lobbying depends on the local and foreign firm's trade costs:*

$$\text{sign}\left(\frac{\partial \bar{\mu}_i^{DTA}}{\partial \gamma_j^i}\right) = \begin{cases} 0, & \text{if } F_i = F_j = 0 \\ -1, & \text{if } F_j > F_i = 0 \\ \in \{-1, 0, 1\}, & \text{otherwise} \end{cases}$$

3. *Under perfect harmonization, more powerful export lobbies push for weaker regulation. That is:*

$$\text{sign}\left(\frac{\partial \bar{\mu}_i^{PER}}{\partial \gamma_j^i}\right) = -1.$$

Proposition 3 suggests that depending on the trade scenario, firms might have conflicting preferences regarding the minimum standards, as one firm might prefer cutting the compliance costs rather than the unit production costs. However, the firms will always benefit from the two standards moving down in unison, as it does not affect the compliance costs but strictly reduces the production costs. Therefore, if the firms are able to coordinate their lobbying efforts across markets, they can always benefit from weaker regulation.

1.3.3 Welfare

Some ambiguity remains over the welfare effects of deepening trade integration. The result is partly due to the limited policy instruments (minimum standards) for tackling multiple market failures (environmental externalities and imperfect competition). Nevertheless, our results highlight some interesting patterns. The first one is related to the welfare changes between a non-cooperative shallow benchmark and a DTA with regulatory cooperation:

Proposition 4 *Regulatory cooperation can reduce welfare compared to a shallow trade agreement if lobbies are sufficiently strong.*

In the absence of lobbies, the government's objective and the society's welfare function are equivalent. In this case, cooperating on national regulations is guaranteed not to decrease aggregate welfare, as the governments can always replicate the non-cooperative outcome if they so choose. As explained before, non-cooperative governments can impose inefficiently strict standards, but the incentive vanishes with regulatory cooperation. However, if the political influence terms are high such that the government objective poorly represents societal welfare, the DTA can lead to an over-representation of private profits. Even if the agreement is welfare-improving, the existence of lobbies might dilute the gains from cooperation. It is, therefore, possible for policymakers to strike an agreement that appears profitable given their biased objective but makes the two economies worse off in aggregate real terms.

From a DTA with regulatory cooperation, a further step to perfect harmonization of standards is never an appealing choice in terms of the policymaker's biased objective function. After all, since harmonization limits the number of available policy instruments while nothing else changes in the structure of cooperation, the value of the planner's problem can never increase. Interestingly, an increase in joint welfare is still possible. Since both policymakers are more constrained by the negotiations, their ability to favour politically influential firms diminishes, thus better aligning the policymaker's biased objective with the true societal welfare.

Proposition 5 *Perfect regulatory harmonization can increase welfare relative to a partial regulatory convergence if lobbies are sufficiently strong.*

For instance, perfect harmonization can produce higher welfare than (partial) regulatory cooperation when lobbies in different markets are highly asymmetric. Intuitively, suppose a country with rampant lobbying signs a trade agreement with a large economy where private interests are absent (or less powerful). In that case, the bargaining power of the existing lobbies must play a smaller role in determining the standards than before. Therefore, although the regulatory level might deviate further from the optimal level in terms of tackling the negative environmental externality, perfect harmonization might, in some instances, still be welfare-increasing for the two economies as a whole.

1.4 Endogenous lobbying efforts

So far, we have treated the political influence terms γ as exogenous parameters to the policymaker. Next, we propose a model that endogenises the lobbying weights in the regulator's objective function. Our goal is to construct an ad-hoc model for investigating lobbying efforts in a framework that complements the previous analysis. We assume that each weight γ becomes an increasing function of the firm's (costly) lobbying effort and that the marginal efficiency of lobbying is decreasing. As we consider a simple partial equilibrium model with two countries and only a single homogenous good, we assume that an equilibrium also exists for this extended setup. Formally, we analyse the following game:

Stage 0: Governments choose the form of cooperation (BAU, DTA, or PER).

Stage 1: Firms choose their lobbying efforts.

Stage 2: Governments set the product standards.

Stage 3: Firms choose their profit-maximizing output quantities.

Notably, we assume that firms determine their lobbying efforts *after* the governments have agreed on the type of the trade agreement.¹² In addition, we focus on

¹²Clearly, this assumption restricts the types of lobbying dynamics our model can address, as firms cannot lobby for the *type* of the agreement. However, we use the EU-Canada CETA agreement as a guiding example. The trade partners announced their political commitment to

a special case where firms can only actively lobby to increase the weight of their domestic market profits, which is sufficient to highlight the main mechanisms at play. Lobbying foreign standards directly might be too costly or impossible due to other political constraints. The function that captures the political influence of firms satisfies the following shape requirements:

$$\begin{aligned}\gamma_i^{i'}(e_i) &> 0, \\ \gamma_i^{i''}(e_i) &< 0, \\ \gamma_i^i(0) &= 1,\end{aligned}\tag{1.7}$$

where e_i denotes the lobbying effort of firm i to boost its domestic market influence. Firms have perfect information on how the weights γ affect the implemented regulations throughout Eq. (1.4) to Eq. (1.6). They can also perfectly observe how their political influence depends on the lobbying effort. Using this information, firm i will then maximize its net profits, $\tilde{\pi}^i$, defined as:¹³

$$\tilde{\pi}^i = \pi_i^i + \pi_j^i - \beta_i e_i.\tag{1.8}$$

Above, β_i denotes the unit lobbying cost. The firms set their lobbying efforts according to the following first-order optimality condition:

$$\begin{aligned}\frac{\partial \tilde{\pi}^i}{\partial e_i} &= -X_i^i \left[\frac{dc_i^i}{de_i} + \frac{dX_i^j}{de_i} \right] - X_j^i \left[\frac{dc_j^i}{de_i} + \frac{dX_j^j}{de_i} \right] \\ &\quad - F_i(\mu_i - \mu_j) \left[\frac{d\mu_i}{de_i} - \frac{d\mu_j}{de_i} \right] - \beta_i = 0.\end{aligned}\tag{1.9}$$

For brevity, we denote $c_j^i \equiv c_i(\mu_j)$ for the firm i 's cost to satisfy market j 's regulatory requirements. Applying the envelope condition significantly simplifies the above

strengthen bilateral trade ties at the EU-Canada Summit in Prague in 2009 and concluded the CETA negotiations in 2014. Therefore, it is reasonable to consider these six years as the window of active lobbying efforts initiated by the announcement of the agreement.

¹³From Eq. (1.2) and Eq. (1.3), we get $\pi_j^i = (X_j^i)^2$ (ignoring the export compliance term).

expression, as firms have already optimized their profits with respect to market-specific output quantities in earlier stages of the game.

Eq. (1.9) highlights the different considerations each firm has when deciding its lobbying levels. First, higher lobbying reduces local regulation, thereby cutting production costs and boosting output and profits. However, what constrains the total expenditure is that more lenient standards at home will also benefit the foreign exporter and increase the competitor's output. We also see that the firm's equilibrium output quantity X always scales the corresponding lobbying expenditure in a given market. Moreover, the only potential motive for lobbying for stricter standards comes from the existence of fixed trade costs F_i .¹⁴

Next, we provide a numerical example to illustrate the profit-maximizing lobbying efforts under different trade scenarios. We follow the work in Ref. [50] and assume that the political clout evolves according to the following equation:

$$\gamma_i^i(e_i) = 1 + \rho_i(e_i)^{\rho_i}, \quad i = A, B, \quad (1.10)$$

with a curvature parameter $0 < \rho_i < 1$ to satisfy the shape requirements in Eq. (1.7). For all experiments we use the IPOPT interior-point optimizer [58] and the JuMP mathematical optimization language [59] in Julia [60]. Appendix 1.B contains the default model parameters.

Figure 1.1 illustrates the reaction functions of the two firms under the DTA scenario. The best responses are increasing and concave, such that the more any individual firm spends on lobbying, the lower the regulation and the more profits both firms will make, thus affording even higher levels of lobbying expenditure. Due to this cycle, the total lobbying efforts are substantially higher when the firms can credibly commit to levels that maximize their joint net profits.

As a final experiment, we simulate the economy 1,000 times, drawing parameters randomly such that they satisfy all parameter assumptions. For each parameter set, we solve the model under all possible trade scenarios.

The numerical example has two main goals. First, the goal is to verify the

¹⁴Eq. (1.9) refers to a case of individual profit maximization. If firms collude and maximize joint profits, all foreign terms enter the equation symmetrically.

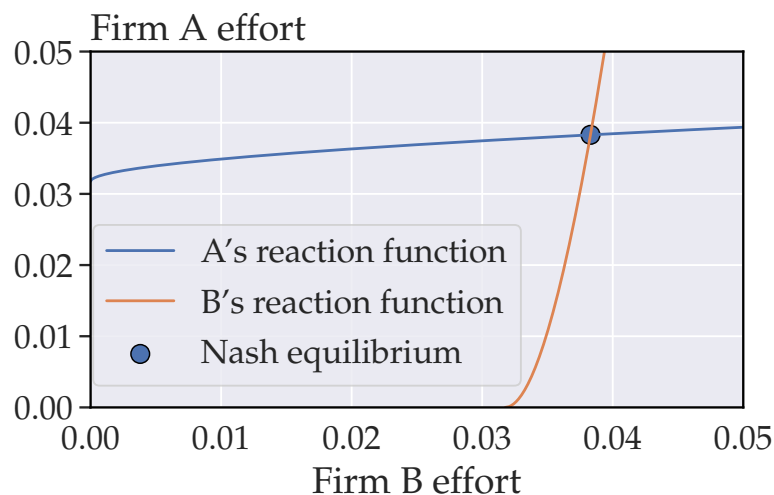


Figure 1.1: Reaction functions for domestic lobbying efforts under a deep trade agreement scenario. The circle highlights the Nash equilibrium lobbying levels. The best responses are increasing, such that more lobbying by one of the firms also increases the lobbying efforts of the other.

predictions from the first-order conditions in Eq. (1.9) and those in Proposition 1. Together they suggest that deeper trade integration will likely boost trade flows, incentivizing firms to spend more on lobbying. Moreover, we anticipate seeing a difference in lobbying activity under the two deep trade scenarios, as the perfect harmonization of standards can dilute the power of lobby groups if the countries are sufficiently asymmetric. Our second goal is to compare the lobbying levels in all trade scenarios when firms can coordinate their efforts over different markets, as deriving analytical expressions is infeasible. From the results in Figure 1.1, we expect the lobbying efforts to be significantly higher when firms can commit to coordinating, highlighting their aligned motives.

Figure 1.2 illustrates the simulation results.¹⁵ It also verifies many of the mechanisms discussed earlier. First, we note that the lobbying levels are generally higher in the DTA and PER scenarios where governments set regulations cooperatively. In contrast, the excessively stringent standards in the BAU scenario restrict trade and suppress lobbying incentives. If the firms can coordinate across markets, the lobby-

¹⁵See Appendix 1.C for a summary of the data.




Figure 1.2: Lobbying statistics under all trade scenarios when firms can either play non-cooperatively against each other (NASH) or cooperate by maximizing joint profits (COOP). Cooperating consistently yields higher lobbying levels than non-cooperative play, highlighting the aligned motives of the firms. The regulation is generally *too* strict in the BAU scenario. That implies lower standards, more trade, higher profits, and higher lobbying in the DTA and PER scenarios where governments maximize a joint objective.

ing levels are notably higher than under a non-cooperative Nash equilibrium. Even in scenarios where one of the firms would (in isolation) prefer a more stringent local standard to narrow the costly regulation gap, colluding firms can always lobby for a simultaneous downward shift in all standards. That keeps trade compliance costs unchanged but reduces overall production costs. Finally, we see a slight reduction in mean lobbying levels when comparing the DTA and PER scenarios.

The numerical values also show that the net profits tend to be highest under the DTA scenario, making it the favourite option for the firms. However, compared to the shallow benchmark scenario, any form of deep integration only improves aggregate welfare when the firms cannot coordinate their lobbying efforts.

1.5 Conclusion

A key feature of modern trade agreements is that countries not only negotiate over conventional trade policy objectives such as tariff cuts. In addition, they try to coordinate away any differences in their national regulations that restrict the total quantity of trade. These deep trade agreements (DTAs) can include various de-

degrees of regulatory cooperation, from the simple alignment of product certification rules to a perfect harmonization of national standards. Among the various existing regulations that can be subject to international coordination, those related to the environment often cause the most public controversy due to the concern of a regulatory race to the bottom.

This chapter investigates the role of special interest groups in shaping the DTAs. We focus on the interplay between firms, policymakers, and environmental regulation, assuming that firms possess market power and political influence. We set up a two-country trade model with imperfect competition, local consumption externalities, and producer lobbies. Although the underlying model is relatively simple, it highlights many complexities surrounding deep trade integration.

DTA's impacts on environmental regulation, lobbying efforts, and welfare hinge on the type of the agreement and the heterogeneity between the trading partners. We first show that in the absence of trade compliance costs, regulatory cooperation always relaxes environmental regulation in both countries compared to a benchmark shallow trade scenario. However, if supplying two markets with different standards is costly for the firms, tightening the local rules can be beneficial for narrowing the international regulatory gap.

Our results show that deep trade integration directly affects lobbying motives and that lobbying efforts tend to be higher under a DTA than in a conventional shallow trade agreement. Firms coordinating their efforts across markets is always detrimental to welfare as it results in rampant levels of lobbying. Whereas previous work suggests that foreign lobbies can affect domestic tariff policies [61], our findings extend the results to DTAs under minimum product standards.

The welfare impacts of DTAs remain ambiguous. Aggregate welfare can increase following an expansion of trade but diminish if the erosion of standards and the redistribution of surplus to the politically powerful firms dominate the effects. Perfect regulatory harmonization can reduce welfare as countries commit to fewer policy instruments. However, in some cases, harmonization can lead to a welfare gain if it limits the government's ability to favour politically powerful firms, thus better aligning the policymaker's biased objective with the society's actual welfare function.

Our results highlight the risk of DTAs strengthening the political influence of

private interests. For instance, under perfect harmonization of national standards, the higher incentives for firms to collude and the resulting increase in lobbying efforts might pose a far greater threat to welfare than the inefficiency from fewer available regulatory instruments. Moreover, if the firms collude in lobbying efforts, it is reasonable to assume they can coordinate output decisions as well, implying an even more significant welfare drop. In the real world, such risks could materialize through cooperation bodies that promote discussions over voluntary harmonization across stakeholders [37]. Therefore, limiting private sector influence will be an essential design issue for any successful deep trade agreement.

Appendix

1.A Proofs

Proof to Proposition 1.

We need to establish the following conditions:

1. $\bar{\mu}_i^{DEEP} \leq \bar{\mu}_i^{BAU}$
2. $\bar{\mu}_i^{DEEP} \leq \bar{\mu}^{FULL} \leq \bar{\mu}_j^{DEEP}$
3. The ranking between $\bar{\mu}_i^{BAU}$ and $\bar{\mu}^{FULL}$ is ambiguous.

To show the first part, we can write:

$$\mu_i^{DTA} - \mu_i^{BAU} = \frac{2\varphi_j \gamma_i^j (9\alpha_i \Xi_i^{BAU} + \varphi_j \Lambda_i^{BAU})}{81 \Xi_A^{BAU} \Xi_A^{DTA}},$$

whose sign only depends on the term inside the parenthesis. Under the special case of $F_A = F_B = 0$ we have that $\bar{\mu}_i^{BAU} = -\Lambda_i^{BAU} / (9\Xi_i^{BAU})$. Since $\frac{1}{2} < \frac{\phi_i}{\phi_j} < 2$ by assumption, we also know that $\varphi_j = 2\phi_j - \phi_i < \phi_i$. Finally, using the condition of strictly positive trade flows, we can write:

$$\begin{aligned} \alpha_i - c_i(\bar{\mu}_i^{BAU}) &> 0 \\ \alpha_i - \varphi_j \bar{\mu}_i^{BAU} &> 0 \\ 9\alpha_i \Xi_i^{BAU} + \varphi_j \Lambda_i^{BAU} &< 0. \end{aligned}$$

That $\bar{\mu}_i^{DEEP} \leq \bar{\mu}_i^{BAU}$ follows directly.

Next, to show that $\bar{\mu}_i^{DTA} \leq \bar{\mu}^{PER} \leq \bar{\mu}_j^{DTA}$, we can rewrite Eq. (1.5) as $\bar{\mu}_A^{DTA} = a/b$, $\bar{\mu}_B^{DTA} = c/d$, and Eq. (1.6) as $\bar{\mu}^{PER} = (a+c)/(b+d)$, where $a, b, c, d \in \mathbb{R}_{>0}$. This means that:

$$\frac{a}{b} = \frac{c}{d} \implies \frac{a}{b} = \frac{a+c}{b+d},$$

that is, $\bar{\mu}_A^{DEEP} = \bar{\mu}_B^{DEEP} \implies \bar{\mu}_A^{DEEP} = \bar{\mu}^{FULL}$. We also have that:

$$\frac{a}{b} > \frac{c}{d} \implies \frac{c}{d} < \frac{a+c}{b+d} < \frac{a}{b},$$

that is, $\bar{\mu}_A^{DEEP} > \bar{\mu}_B^{DEEP} \implies \bar{\mu}_B^{DEEP} < \bar{\mu}^{FULL} < \bar{\mu}_A^{DEEP}$, and reversely:

$$\frac{c}{d} > \frac{a}{b} \implies \frac{a}{b} < \frac{a+c}{b+d} < \frac{c}{d},$$

which gives $\bar{\mu}_B^{DEEP} > \bar{\mu}_A^{DEEP} \implies \bar{\mu}_A^{DEEP} < \bar{\mu}^{FULL} < \bar{\mu}_B^{DEEP}$. Therefore, $\bar{\mu}_i^{DEEP} \leq \bar{\mu}^{FULL} \leq \bar{\mu}_j^{DEEP}$ always holds.

The final remaining comparison is that between $\bar{\mu}_i^{BAU}$ and $\bar{\mu}^{PER}$. We simply show that there exist parameterizations of our model that satisfy all parameter restrictions and result in different rankings between the regulatory levels. Consider first two identical economies with parameter values specified in Table 1.B.1. In this case, we have that $\bar{\mu}_A^{BAU} = \bar{\mu}_B^{BAU} > \bar{\mu}^{FULL}$. Alternatively, consider the default parameters but let $\alpha_A = 3, \alpha_B = 3, \delta_A = 2.1, \delta_B = 2.71, \gamma_A^A = 1.1$. This gives $\bar{\mu}_A^{BAU} < \bar{\mu}^{FULL} < \bar{\mu}_B^{BAU}$, making the rankings ambiguous. ■

Proof to Proposition 2.

For the BAU and DTA policies in Eq. (1.4) and Eq. (1.5), we can separate the optimal standard in each scenario into $\bar{\mu}_i = a/b$, we can write:

$$\frac{\partial \bar{\mu}_i}{\partial \gamma_i^i} = \frac{2\Xi_j \varphi_i (b\alpha_i - a\varphi_i)}{b^2}.$$

Similarly for the perfect harmonization in Eq. (1.6):

$$\frac{\partial \bar{\mu}_i}{\partial \gamma_i^i} = -\frac{2\varphi_i (b\alpha_i - a\varphi_i)}{b^2}.$$

For all expressions, the sign depends inversely on $b\alpha_i - a\varphi_i = \alpha_i - \varphi_i \bar{\mu}_i$. Repeating the steps outlined in Proposition 1, we know the term to be strictly positive. Therefore, a higher weight γ_i^i on domestic profits always results in less stringent standards. ■

Proof to Proposition 3.

For the BAU scenario, we notice that:

$$\frac{\partial \bar{\mu}_i}{\partial \gamma_j^i} = \frac{(\bar{\mu}_i - \bar{\mu}_j) F_i \Xi_j}{\Xi_i \Xi_j - F_i F_j \gamma_j^i \gamma_i^j}.$$

Under our parameter assumptions, the sign of the above expression only depends on the difference between the equilibrium standards $\bar{\mu}$. With $F_i = 0$, the connection between regulation and export profit influence disappears.

Under a DTA scenario, denoting $\tilde{F} = (F_A \gamma_B^A + F_B \gamma_A^B)$, we can write:

$$\frac{\partial \bar{\mu}_i}{\partial \gamma_j^i} = \frac{F_i (\Lambda_j + 9\Xi_j \bar{\mu}_i) + \tilde{F} [9F_i (2\bar{\mu}_i - \bar{\mu}_j) - 2\alpha_j \varphi_i + 2\bar{\mu}_j \varphi_i^2]}{9(\Xi_i \Xi_j - \tilde{F}^2)}.$$

When both $F_A > 0$ and $F_B > 0$, the expression can take any sign. When $F_j > F_i = 0$, the sign of the above expression only depends on $-2F_j \varphi_i (\alpha_j - \varphi_i \bar{\mu}_j) \gamma_i^j$, which, repeating the analysis in Proposition 1, we know to be always negative. With both $F_A = F_B = 0$, the relationship again vanishes.

Finally, under perfect regulatory harmonization, we have that:

$$\frac{\partial \bar{\mu}}{\partial \gamma_j^i} = \frac{2\varphi_i (\alpha_j - \varphi_i \bar{\mu})}{9\Xi},$$

which is again always strictly negative under our parameter assumptions. ■

Proof to proposition 4. To prove the welfare-related propositions, we simply evaluate our model under different sets of parameters that satisfy all the assumptions and constraints. To see that regulatory cooperation can increase welfare when lobbies are sufficiently weak, use the default parameters in Table 1.B.1 with all influence terms γ set to the value of 1.1. This yields $W^{DTA} > W^{BAU}$. Next, repeating the analysis with all values of γ set to 1.3, we get that $W^{DTA} < W^{BAU}$. ■

Proof to proposition 5. We prove the proposition numerically as before. To see that perfect harmonization can increase welfare, simply use the default parameters in Table 1.B.1 with one import-competing lobby γ_i^i endowed with a high value of 1.4. This yields $W^{PER} > W^{DTA}$. Next, repeating the analysis with the default parameters and setting $\alpha_i = 4.2$ to introduce an asymmetry between national regulations, we get $W^{PER} < W^{DTA}$. ■

1.B Parameter values

Parameter	Description	Value
α_i	Market demand intercept	4.50
β_i^i	Unit lobbying cost	1.00
δ_i	Unit externality	3.00
ω_i	Abatement efficiency	1.00
ϕ_i	Production cost	1.05
ρ_i^i	Political influence curvature	0.20
F_i	Export compliance cost	0.50
γ_i^i	Political influence	1.00

Table 1.B.1: Default parameters for numerical simulations.

1.C Numerical results

	Agreement Scenario	BAU		DTA		PER	
		COOP	NASH	COOP	NASH	COOP	NASH
Q_A	max	2.4017	2.3577	3.7917	3.4547	3.7866	3.3823
	mean	0.8921	0.8830	1.4272	1.3564	1.3988	1.3288
Q_B	max	2.7112	2.6996	4.5600	4.1770	4.5731	4.1687
	mean	0.8564	0.8477	1.4107	1.3398	1.4269	1.3569
W	max	7.8678	7.9209	7.2770	8.0671	7.2713	8.0313
	mean	1.6182	1.6267	1.4873	1.7117	1.4819	1.7027
e_A	max	0.8239	0.2613	5.0000	2.6889	5.0000	2.6864
	mean	0.0233	0.0086	0.2090	0.0500	0.2035	0.0481
e_B	max	1.1108	0.3663	5.0000	3.0073	5.0000	2.7224
	mean	0.0166	0.0058	0.2423	0.0572	0.2410	0.0562
μ_A	max	19.4216	19.4478	16.1741	16.4975	16.1719	16.4952
	mean	2.7798	2.7967	1.8756	1.9971	1.9043	2.0251
μ_B	max	19.4218	19.4482	16.1673	16.4902	16.1719	16.4952
	mean	2.8465	2.8628	1.9200	2.0418	1.9043	2.0251
$\tilde{\pi}_A$	max	3.5795	3.5737	8.0950	7.2001	8.0927	7.0676
	mean	0.5337	0.5319	1.1629	1.1266	1.1541	1.1175
$\tilde{\pi}_B$	max	2.9314	2.9153	6.6517	5.9085	6.6846	5.9082
	mean	0.4701	0.4671	1.0736	1.0476	1.0684	1.0432

Table 1.C.1: Main results for the numerical example.

Chapter 2

Solar Geoengineering Governance: A Coalition Formation Approach

Abstract

Climate interventions with solar geoengineering could reduce climate damages if deployed in a globally coordinated regime. In the absence of such a regime, however, strategic incentives of single actors might result in detrimental outcomes. Governance structures, i.e. more or less formal institutional arrangements between countries, could steer the decentralized geoengineering deployment towards the preferable global outcome. We show that the coalition formation literature can make a valuable contribution to assessing the relative merit of different governance schemes. A country pondering whether to leave or join a coalition anticipates that its decision could spark another (dis)integration process among the other players. This dynamic structure enables a more realistic picture of what coalitions are likely to form and remain stable. Moreover, the model provides a unified framework for comparing different institutional settings for geoengineering deployment, such as consensus-based and majority-vote coalitions.

This chapter is joint work with Daniel Heyen. Reprinted with minor edits from Ref. [62].

2.1 Introduction

Solar geoengineering (SG) denotes technologies for reflecting solar radiation away from Earth to counteract climate change [63, 64]. A prominent proposal is the injection of aerosols into the stratosphere. Such global cooling can significantly reduce climate damages *if* appropriately deployed but could severely worsen the situation if not [64, 65]. The main concerns in this context are unilateral SG deployment and the lack of international cooperation [14, 22, 23, 66, 67].

Bridging the gap between potentially disadvantageous and beneficial SG deployment can be achieved by *governance of SG* [24, 68–72]. Assessing the merit of different governance schemes requires predictions about how they would alter countries’ incentives to cooperate on SG deployment. Our basic assumption in this article is that countries rationally deploy SG to further their interests. As the international sphere is anarchic in the sense that no supranational body can enforce a certain use of SG, cooperation among countries is possible but needs to be in the interest of all cooperating partners [71].

We present a framework of SG coalition formation, drawing on a rich literature in economics [25, 26, 73]. This framework has two main advantages. The first advantage is the ability to represent complex dynamics of coalition formation. Existing papers on SG cooperation [74–76], rooted in a long tradition of analysing cooperation incentives for greenhouse gas abatement [77–80], model SG coalition formation as a two-stage game: in the first stage, countries decide whether to enter a treaty, i.e. commit to deploying SG in cooperation with their fellow treaty members; in the second stage, countries decide on SG deployment. According to that modelling approach, a coalition is stable if it is internally stable (no member wants to leave) and externally stable (no outsider wants to join). While insightful, this approach is static and therefore unable to capture – as we will demonstrate – important dynamics of coalition formation. More realistic contributions, in particular those with farsighted players anticipating ripple effects from changes to a coalition, have appeared in the general climate context [81–85]. Except for a two-player repeated game in Ref. [86], to our knowledge, the SG literature has yet to adopt that dynamic approach for analysing cooperation incentives.

The second advantage of the coalition formation framework is its flexibility.

We interpret the governance of SG deployment as setting the international rules of who can use SG under what circumstances. For instance, whether treaties are irreversible or renegotiable and whether coalition membership requires consent by existing members (exclusive membership game) or not (open membership game) are easy to capture in our framework. This flexibility of the framework allows modelling a wide range of formal and informal international arrangements, thus facilitating the interdisciplinary debate about SG governance.

2.2 Theoretical framework

This section gives a concise description of the SG coalition formation framework. More technical details are available in Appendix 2.A. The framework rests on the vast general literature on coalition formation; see Ref. [73] for an overview. There are n countries and $t = 0, 1, \dots$ periods. Of central importance is the **state space** X : state $x_t \in X$ describes the composition of SG coalitions at time t . Countries in a coalition cooperatively decide on the level of SG deployment and behave non-cooperatively towards other countries.¹ We allow countries to differ in terms of **power**. In a given coalition, more powerful countries have more say in how the coalition deploys SG; more powerful countries are also more likely to initiate changes to the existing state of coalitions (see the "protocol" below). The **period-payoff function** $u_i : X \rightarrow \mathbb{R}$ captures the payoffs of country i for all possible states: payoff $u_i(x_t)$ fully summarizes the entire impact to country i at time t when at that time countries are organized into coalitions as described by state x_t . The payoff function clearly rests on several assumptions about the "global thermostat game" of SG deployment, such as country-specific climate damages, how SG affects the climate, the direct and indirect costs associated with SG deployment, and what level of SG coalitions choose. In Section 2.3, we will develop an example with specific assumptions.

The rules of coalition formation depend on the specification of a protocol and approval committees, and the flexibility allows to model very different governance arrangements on SG deployment. In every period, the **protocol** randomly selects

¹We assume all coalitions to be disjoint, such that any country can only belong to a single coalition at a time.

one country to be a proposer (more powerful countries have a higher chance of being selected) and describes which countries are open to receiving proposals from the selected country. *Reversible agreements* are those where the set of potential partners is always the entire set of countries. *Irreversible agreements*, in contrast, can be modelled by allowing to approach only countries as potential partners that have not formed any previous coalitions during the game so far. Moreover, irreversible agreements can be *temporary* instead of *permanently binding*. In such a case, the protocol simply ceases to select any committed countries as new proposers or potential partners until their going agreements have expired.

The second key building block of the framework, **approval committees**, specifies which state transitions require the consent of others. The proposer must send its preferred state transition to an approval committee. The new state only realizes if the committee approves the transition (some transitions might be infeasible in the first place; in such a case, no approval committee can validate the move). The framework allows for different approval rules such as *unanimity* (all members of the approval committee need to consent) or a *majority* vote. Other interesting features, such as the renegotiation of treaties and the distinction between open membership (countries entering existing treaties at will) and exclusive membership games (existing treaty members having to approve joining countries) [79], can also be considered. The set of respondents in the approval committee can also be larger than the set of countries invited to move to a new coalition, reflecting a situation in which all countries whose treaties are affected by the proposed move need to be consulted.

For a given protocol and approval committees, what coalitions might form? Countries formulate **strategies**, i.e. full plans of the form "if the state is x , what new coalition will I propose?" and "if the state is x and country j proposes moving to state y , am I going to accept?". Mixed strategies where countries can propose and accept transitions with a probability between 0 and 1 are also allowed. The strategies of all players, together with the initial state x_0 and the protocol, imply a stochastic process P for the evolution of states $(x_t)_t$ over time. Accordingly, all countries can calculate the expected value at the initial state as:

$$V_i(x_0) = \mathbb{E}_P \sum_{t=0}^{\infty} \delta^t u_i(x_t). \quad (2.1)$$

The parameter δ – for simplicity assumed to be the same for all countries – controls the level of **farsightedness**. The higher the parameter δ , the more important are long-term permanent relative to short-term temporary gains. For a value of δ close to 1, countries might be willing to push their payoffs far into the future or even endure temporary losses to better position themselves for future negotiations. As usual in game theory, an **equilibrium** of a coalition formation game is a strategy profile such that no country wants to change their strategy given the strategies of others. An equilibrium may consist of a never-ending coalition formation process, continuously cycling, or settling on some specific *absorbing* state. An equilibrium exists under mild conditions (see Appendix 2.A.4) but is not unique in general.

2.3 Numerical examples

This section presents a simple setting with three countries and stylized assumptions about climate change, damages, and SG to illustrate the potential of the coalition formation framework from Section 2.2.

2.3.1 Climate change, damages, and solar geoengineering

Consider three countries with different baseline temperatures: W (warm), T (temperate), and C (cold). At the moment countries decide about SG, climate change has warmed all of them uniformly by 3°C . For the clarity of exposition, we assume SG is costless, has no side effects, and cooling is uniform. We can then consider the global SG level G as the cooling experienced in every country. For this stylized example, we assume that climate damages only stem from (local) temperature levels. In line with Ref. [87] and as used by Ref. [76], we consider quadratic damages that are minimal at an ideal temperature, here 13°C . As Appendix 2.B.1 demonstrates, we can write country i 's period payoffs as a function of the global SG level G with:

$$u_i(G) = G \cdot (2\alpha_i - G). \quad (2.2)$$

The payoffs are highest at country i 's ideal SG level α_i . The warmer the country, the higher the ideal SG level; see Table 2.3.1. Eq. (2.2) is normalized such that in the absence of SG, $G = 0$, the period payoff is zero for all countries. We can then interpret a positive period payoff as the country benefitting from SG deployment.

	Country		
	Warm W	Temperate T	Cold C
Baseline temperature	21.5	14.0	11.5
Temperature with climate change	24.5	17.0	14.5
Ideal local temperature	13.0	13.0	13.0
Ideal SG level α_i	11.5	4.0	1.5

Table 2.3.1: A stylized example featuring three countries with different baseline temperatures, uniform temperature increase under climate change, and uniform cooling under SG. The assumption of the same ideal local temperature, consistent with Ref. [87], results in heterogeneous ideal SG levels α_i . All numbers in [$^{\circ}$ C].

2.3.2 Assumptions on the coalition formation framework

We assume that all countries have the same level of farsightedness, $\delta = 0.99$. We also assume they have the same power, which has two consequences. First, the protocol is uniform, i.e. in every period, each country gets to propose with a probability of $1/3$. Second, countries in a coalition aim to implement their average ideal SG level. With only three countries, $N = 3$, at most one non-singleton coalition can form. Therefore, we denote the possible states of the coalition formation process as the absence of SG coalitions (\emptyset), the two-country coalitions (TC), (WC) and (WT), and the grand coalition (WTC).

In the following, we make two assumptions about the international sphere in which SG deployment occurs. The first assumption is that every country can unilaterally leave any treaty.² In the language of our framework, this means that, for example, country C can suggest the transition from (WTC) to (WT), and no approval committee needs to consent. If a country leaves, the remaining members

²Many international treaties explicitly allow countries to withdraw. For instance, the U.S. was able to withdraw from the Paris Agreement.

are assumed to stay at least temporarily in the treaty. In later periods, however, those remaining countries can break apart themselves, and the anticipation of the ensuing process critically shapes the incentives to leave the treaty in the first place.

The second assumption is to rule out side payments (transfers) from one country to another to lure them into treaties. This implies that Eq. (2.2) fully determines country i 's period payoff. While side payments play a beneficial role in some international treaties, the absence of full cooperation suggests that their role is limited [88]. Therefore, ruling out side payments is a useful starting point for illustrating the coalition formation framework.

Weak governance

We will now contrast two scenarios of SG governance. In our first scenario, "weak governance", countries are free to deploy SG as they please. Panel A in Figure 2.3.1 summarizes the relevant aspects of the equilibrium strategies in that scenario. Both T and C would prefer the grand coalition (WTC). However, country W can always do better by walking out of any existing treaty and unilaterally deploying its ideal SG level $\alpha_W = 11.5$. Any other option would imply compromising with countries of lower ideal SG levels and dragging down W's payoffs. The stochastic process reflects the same reasoning; see panel A in Figure 2.3.2. With probability $1/3$, W is the proposer and will leave any previously agreed coalition.

In contrast, both T and C prefer any coalition with W to () or (TC) and therefore do not change the state whenever they are to propose. Eventually, the coalition formation process converges to the absorbing state ().³ Therefore, the prediction is clear: under weak international SG governance, W will deploy a high SG level, causing substantial external damage to others. That is a stylized representation of the free-driver concern [14]. Note that the prediction would be the same also for the more conventional, static approach: any coalition with W as a member would violate the condition of internal stability as W always prefers to leave.

³Note that in the weak governance scenario, the state (TC) is equivalent to (). Irrespective of whether T and C form a coalition, W is free to deploy SG, and by even higher cooling, T and C would harm themselves. Therefore, all payoffs coincide in (TC) and (). Accordingly, the equilibrium strategy in Figure 2.3.1 is not unique. Since all countries are indifferent between staying in (TC) and moving to (), there exist equilibrium strategies under which either () or (TC) can be an absorbing state.

A. Weak Governance				B. Power Threshold			
State	Country			State	Country		
	W	T	C		W	T	C
() G=11.50	*132.25 132.25	-40.25 -40.25	-97.75 -97.75	() G=0.00	0.00 55.13	0.00 14.29	0.00 0.68
(TC) G=11.50	*132.25 132.25	-40.25 -40.25	-97.75 -97.75	(TC) G=2.75	55.69 55.69	*14.44 14.44	*0.69 0.69
(WC) G=6.50	*107.25 131.51	9.75 -38.78	*-22.75 -95.54	(WC) G=6.50	*107.25 56.66	9.75 14.16	*-22.75 -0.01
(WT) G=7.75	*118.19 131.84	*1.94 -39.01	-36.81 -95.96	(WT) G=7.75	*118.19 56.99	*1.94 13.93	-36.81 -0.42
(WTC) G=5.67	*98.22 131.25	*13.22 -38.68	*-15.11 -95.32	(WTC) G=5.67	*98.22 57.44	*13.22 14.03	*-15.11 -0.44





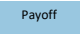

Legend:	 (Coalition) SG level	 high period payoff  medium period payoff  low period payoff	 Payoff *Payoff	country is not deploying SG country is deploying SG	 Period payoff Value function Proposition	W's consent needed, but W rejects C's consent needed, and C accepts
----------------	---	---	---	--	--	--

Figure 2.3.1: Summary of period payoffs and equilibrium strategy profiles under weak governance (panel A) and power threshold (panel B). Period-payoffs u_i , cf. Eq. (2.2), and long-term expected values given strategies V_i , cf. Eq. (2.1), are shown in red (country W), green (country T), and blue (country C) boxes. Dark and light fills represent more and less attractive constellations, respectively. An asterisk in front of the payoff indicates that the respective country is deploying SG. Every row is dedicated to a certain going state. The dark grey box on the left indicates the coalition structure and the total global SG level in that state. For every state, the light grey boxes show the proposal of the respective country in equilibrium. If that proposition needs approval by others, the countries needed for approval are indicated by small coloured boxes to the right of the proposal: solid filled if the respective country approves the proposal in equilibrium and filled with a diagonal line if the respective country rejects the proposal.

Power threshold

In the second governance scenario, we explore the case in which SG deployment requires the deploying coalition to possess at least half of the global power, a scenario considered in Ref. [74]. The underlying story could be of strict sanctions by non-deploying countries. Whether such sanctions are incentive-compatible is not the question we ask here; see the discussion in Section 2.4. In our three-country setting

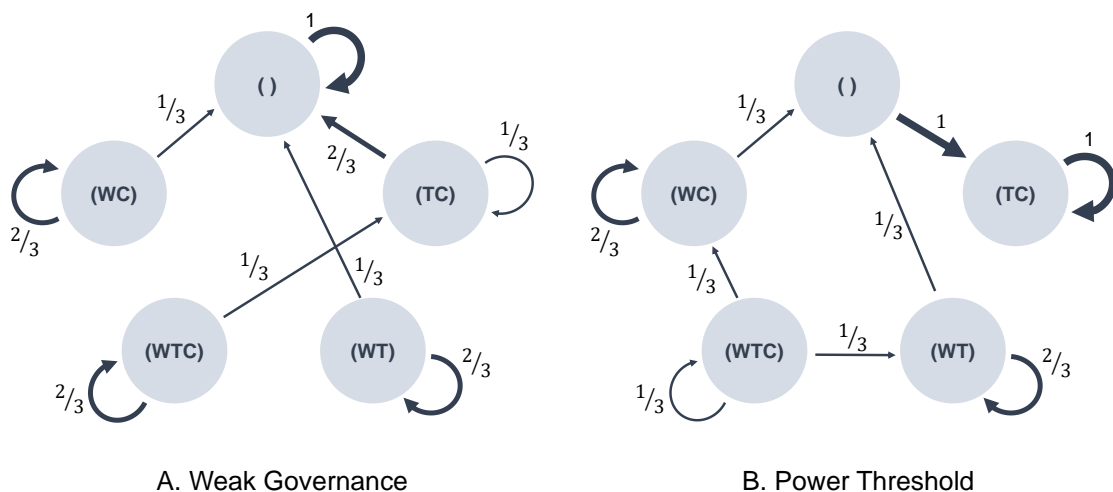


Figure 2.3.2: State transition probabilities under weak governance (panel A) and power threshold (panel B). The probabilities directly follow from the uniform protocol (in every period, each country gets to propose with equal probability) and the equilibrium strategies in Figure 2.3.1. For weak governance, the states () and (TC) are equivalent as country W is free to implement its preferred SG level irrespective of the coalitional configuration of the two other countries. Accordingly, both () and (TC) are potential absorbing states of the equilibrium coalition formation process. For the power threshold, the state (TC) is the unique absorbing state, and the corresponding SG level is low (the midpoint of α_T and α_C).

with equal powers, the power threshold implies that SG deployment can only occur in a coalition of two or three countries. Other than the power threshold, the game remains exactly as above.

The resulting coalition formation game is substantially different from the weak governance scenario. As panel B in Figure 2.3.1 demonstrates, country W is now willing to cooperate to access the global thermostat; the best partner is T, the country with the most similar ideal SG level. On the other hand, T and C prefer a coalition among themselves. Both benefit from SG, and teaming up is the least costly compromise; recall that T and C are closer in ideal SG levels than T and W.

The new institutional setting also changes the dynamics of the state transitions; see panel B in Figure 2.3.2. Due to the power threshold, W tries to maintain any of the high-deployment states when it is the proposer. However, T and C now have a more active role to play. Consider first why, given that countries are farsighted,

the grand coalition (WTC) cannot be an absorbing state, even though none of the players have profitable single-step deviations available. By moving the game from (WTC) to (WT), country C suffers high temporary losses by giving up its access to the global thermostat but knows that it will trigger a further deviation by country T from (WT) to (). This is because from there, the game moves directly to (TC) with relatively high payoffs for both players T and C and no further incentives to deviate.⁴ Country T has the same motive, with the transition going through (WC) instead of (WT). This cycle of disintegration and regrouping occurs because W can block a direct move from the grand coalition to (TC).

Such transitions would be lost under a static approach. In fact, all of the states (WTC), (WT), and (TC) would be equilibrium predictions if the countries were not farsighted. In the dynamic framework, instead of the per-period payoffs, countries rank different states in terms of their expected long-run value. They only accept transitions to a new state if it yields a higher value and only propose states where the expected value, given the probability of that state being approved, is maximized. It is also the reason why country T, if sufficiently farsighted, would in () reject a proposal from W to implement (WT), although there would be a short-term improvement for both countries. In our setting, the power threshold governance effectively moves the deployment power from the free driver to any coalition of natural partners with a high enough share of world power. Indeed, if the ideal SG level of T were higher and closer to that of W than C, the absorbing state would have been (WT) instead of (TC).

That (TC) is an absorbing state when it yields the highest possible payoff for both T and C is easy to see. However, the equilibrium strategies in Figure 2.3.1 do not rest on this assumption. In Appendix 2.D, we provide an additional example where all countries receive their highest period payoffs in different states: W in (WT), T in (WTC), and C in (TC). The strategies in panel B in Figure 2.3.1 still constitute a valid equilibrium. Interestingly, country T would still choose to break out of the grand coalition, only in anticipation of C doing so. The reason is that by acting first, T can receive a higher cumulative payoff as the game transitions to the

⁴Note that in equilibrium, the warm country W would also propose a move from () to (TC), despite not having any influence on the SG deployment. This is because it knows that T and C would reject any other proposals, only waiting to implement (TC) themselves. Since there is no SG in (), the only way for W to get to a state with positive SG deployment is by proposing (TC).

absorbing state (TC) via (WC) instead of (WT). The whole reason for the grand coalition (WTC) not being stable is that C would deviate in anticipation of (TC) ultimately materializing, and T would deviate in anticipation of C's strategy.

In a second supplementary example, we consider a case where a majority approval committee can enforce state transitions instead of a full unanimity requirement. Such an institutional setting can support even the grand coalition as an absorbing state. Consider the above examples where country C wants to deviate from the grand coalition, anticipating (TC) to emerge as the absorbing state. In (TC), however, T could invite (and approve) W to join the grand coalition without consulting C. Since deviating from (WTC) implies heavy temporary losses for C, it no longer wishes to initiate the disbanding cycle, and the grand coalition emerges as an absorbing state. Appendix 2.D contains further details on the supplementary examples.

The examples highlight that small changes in the geoengineering game or institutional setting can result in notably different equilibrium strategies, even for a simple game with just three countries. This brings to the fore the subtleties of farsighted behaviour and dynamic strategies that would be lost in a static framework.

2.4 Conclusion

The illustrative example in Section 2.3 is designed to demonstrate the power and subtleties of the coalition formation framework. Clearly, we have made several simplifying assumptions. First, the underlying geoengineering deployment game is deliberately simple. In the example, we only consider three stylized countries, temperature changes both from climate change and SG are uniform, climate damages are only temperature related, and SG is costless. All those simplifying assumptions can and should be replaced by more realistic modelling choices [89]. Also, note that the equilibria of coalition formation are, in general, not unique. This leads to an important equilibrium selection problem in which players might try to convince others to play a certain equilibrium instead of a different one.

The second simplification was to study SG deployment in isolation, leaving aside other climate policy instruments. Accordingly, future research can connect the coalition formation literature to the interplay of SG with mitigation [e.g. 75, 90, 91] and

counter-geoengineering [76, 86, 92, 93]. Research and development of geoengineering capabilities has not been the focus here either [94–97], but could be tackled with the same approach, thus contributing to debates about research governance [63, Chapter 5]. Transfers between countries could also be easily included in the dynamic coalition formation game, with interesting impacts on efficiency and transition dynamics [98].

For our simple example, we have assumed existing governance in the background (such as weak governance and a power threshold) and looked at coalition formation within those exogenous governance boundaries. Obviously, countries negotiate around those governance rules as well. More complex models of coalition formation can put those negotiations over governance arrangements centre-stage. In the spirit of the backward induction method, the result from the coalition formation game on SG deployment under specific governance arrangements, such as the two in our example in Section 2.3 or more sophisticated versions, will then constitute the *states* in the coalition formation game on SG governance. This game will allow predictions about which governance arrangements are likely to emerge.

Appendix

2.A Theoretical framework

This section presents the main elements of a general coalition formation framework, making extensive use of Ref. [25], Ref. [26], and Ref. [73].⁵ We will illustrate the components with a simple example.

Consider a set N of countries with cardinality n , an infinite set $t = 0, 1, \dots$ of periods, and a compact set X of states. A coalition is a non-empty subset of N . A state captures (at the very least, see extensions below) the *coalition structure*, i.e. how countries are arranged into coalitions. Such a coalition structure is described by a *partition* π of the set of all countries $\{1, \dots, n\}$. We assume all coalitions to be disjoint, such that any country can only belong to one coalition at a time.⁶ For example, in the case of three countries, $N = \{1, 2, 3\}$, the following coalition structures are possible:

$$\begin{aligned}x^0 &= \{\{1\}, \{2\}, \{3\}\}, \quad x^1 = \{\{23\}, \{1\}\}, \\x^2 &= \{\{13\}, \{2\}\}, \quad x^3 = \{\{12\}, \{3\}\}, \quad x^{GC} = \{\{123\}\},\end{aligned}$$

where, for instance, $\{\{12\}, \{3\}\}$ means that country 1 and country 2 form a coalition and country 3 remains outside as a singleton.⁷ State x^{GC} denotes the *grand coalition* comprising all countries, and state x^0 consists of only singleton coalitions.

⁵We adopt the non-cooperative bargaining approach instead of cooperative blocking as the default form of interaction between coalitions. The reason is that we think it fits better with international negotiations.

⁶Ref. [84] considers a nested structure of coalitions, and in Ref. [99] the coalitions can overlap.

⁷By definition, singletons are also considered as coalitions.

Countries differ in their political **power** $\gamma_i > 0$, as used in Ref. [74] and Ref. [84]. Political power might correspond to (a combination of) economic power (measured by GDP), military power, or international reputation. We normalize the terms γ_i with $\sum_{i \in N} \gamma_i = 1$, so that each individual term corresponds to the country's share of total world power.

Power plays a dual role in our framework. First, we assume that a coalition determines its level of geoengineering deployment based on its members' relative power. A more powerful country thus has more say on how a coalition sets the global thermostat. Second, we consider a governance scheme with a minimum power threshold $\underline{\gamma}$ required to implement geoengineering. A coalition C can only deploy geoengineering if its total power exceeds the power threshold, $\sum_{i \in C} \gamma_i \geq \underline{\gamma}$. This power threshold is the first (exogenously defined) way to capture geoengineering governance rules. Ref. [74], for instance, sets $\underline{\gamma} = 0.5$. The threshold only reflects a *necessary* lower bound on "big enough even to start it" or "big enough to push it through against others". For instance, military opposition or economic sanctions to prevent geoengineering might only work if the coalitions that deploy geoengineering are sufficiently weak. The power threshold framework is also fairly flexible. It spans cases as diverse as $\underline{\gamma} = 1$, where only the grand coalition can deploy geoengineering, and $\underline{\gamma} = 0$, which does not impose any power restrictions at all.

For each country there is a continuous one-period payoff function $u_i : X \rightarrow \mathbb{R}$ and a (common) discount factor $\delta \in (0, 1)$. The instantaneous payoff for country i , $u_i(x)$, is then fully described by the going state x . This rests on the assumptions on the equilibrium of the underlying geoengineering deployment game and on whether and how coalitions can transfer utility among their members. We describe these assumptions in detail in Appendix 2.B.

2.A.1 Institutional rules

A dynamic coalition formation game requires clear institutional rules to govern which changes in the coalition structure are possible at a given period. Some changes might be infeasible due to existing agreements' irreversibility or if the renegotiation requires the approval of some minimum share of affected players. On the other hand, a game where coalitions can invite new members, dissolve, and merge with

other groups over time can be a better description for some settings. To consider a general framework that can handle a range of possible institutional settings, we model the evolution of coalition structures as a sequential game of non-cooperative bargaining. Under this approach, one country in each period can propose a new coalition to form, while other countries can either accept or reject the proposition. We next introduce two complementary concepts that determine the rules of such a game. A protocol specifies an active proposer, a set of potential partners, and the order of responders for each period. The proposer sends the proposition to an approval committee – a collection of countries with the power to validate the transition.

The **protocol** determines how reversible the existing agreements are. Each period t starts with the protocol ρ choosing one country S as the active proposer. For instance, the protocol can be a random process, selecting countries with uniform probabilities or weighted according to their share of global political power. The protocol can also depend on the history of the game in arbitrarily complex ways. Along with the active (singleton) proposer S , the protocol specifies the collection of potential partners $P \subseteq N$ to be all countries that can receive proposals at time t . If the agreements are fully reversible, the set of potential partners is always the entire set of countries N . However, for irreversible treaties, the set only consists of countries that have not formed any previous coalitions during the game so far. Finally, for some subset of all potential partners $Q \subseteq P$, S proposes a new coalition R entering into force.⁸ Agreements can also be temporary instead of permanently binding. In such a case, the protocol stops selecting any committed countries as new proposers or potential partners until their going agreements have expired.

Approval committees specify the rules for renegotiation and how much power a particular group has over a new state becoming effective. More formally, for any current state x , let $F_A(x) \subseteq X$ be the set of feasible next states that a set of countries A has the power to approve. Any proposition to move from the current state x to some new state y needs approval from A such that $y \in F_A(x)$. Notably, the set of responders in the approval committee A can potentially be much larger than the set of countries invited to move to a new coalition. Consider the example coalition

⁸The proposer can always choose to maintain the status quo. That is, proposing an existing coalition that contains S is always trivially accepted and requires no approval committee.

structures above. If the going state is x^3 and country 3 is the active proposer, it might propose moving the system to x^1 where it forms a new coalition with country 2. Whether country 1 first needs to free country 2 from any previous obligations depends on the game at hand through restrictions on the approval committee. A unanimous acceptance by the approval committee moves the game into the new state.⁹ Any rejections mean that the state remains unchanged, and the game moves to the next period with a new proposer selected by the protocol. This game continues indefinitely.

Assumptions on the protocol and approval committees reflect the sphere of international agreements that we are interested in and the bargaining tools that countries have at their disposal in the international arena. First, we place no restrictions on what transitions a country can propose. Any player might then suggest other countries to join forces or existing coalitions to disintegrate, even if it is not self a member of those coalitions. Another form of leverage that we allow the countries to have is the possibility to walk out of their current treaties unilaterally without consulting the other members. For more strictly binding agreements, one could assume, for instance, that a country must be on the approval committee for any transition that affects the composition of its current coalition [26].

The **history** of the game, H_t , keeps track of previous coalition structures, propositions, and responses. It can also be included in the state representation x together with the coalition structure π . In some instances, the history is merely a bookkeeping item with little impact on the actual bargaining process. However, it also enables the implementation of very detailed protocol-approval committee combinations. One potential use case is the introduction of rejector power, where the choice of new proposers depends on the history of recent rejectors [100]. In an extreme case, the first country to reject a proposition automatically becomes the new proposer, making the rejection of a proposal a much more valuable tool than under a uniformly random proposer selection.

⁹Extensions are also possible where some non-unanimous vote is enough to pass a proposed transition.

2.A.2 Partition functions and cross-coalition externalities

Externalities across coalitions play an essential role in our geoengineering game. For any individual country, the instantaneous payoff depends not only on the composition of its own coalition but also on what other coalitions exist simultaneously within the entire coalition structure. Partition functions enable a flexible representation of these cross-coalition externalities.

As before, let N be the set of all countries and the coalition structure π a partition of N . Define $M(C, \pi)$ to be a mapping that specifies for each coalition $C \in \pi$ a set of payoff vectors in $\mathbb{R}^{|C|}$.¹⁰ That is, if we assume some transferability of utility, the elements in this set reflect the possible ways any coalition can divide its aggregate payoffs between members, given the ambient coalition structure. A partition function game (N, M) then defines the value of each coalition in every ambient coalition structure that it can be an element of.

A consequence of coalitional payoffs depending on the entire coalition structure is that countries' propositions must depend on what happens elsewhere in the system. In our context, this might mean that a geoengineering coalition conditions its value-sharing rule on whether the outside countries form a counter-geoengineering coalition or remain as singletons. Formally, then, any proposition is a tuple $(C, \mathbf{u}(\pi))$ where the vector $\mathbf{u}(\pi)$ contains a proposed allocation of value across members in a new coalition C for every conceivable partition π that contains C .

The extent to which utility is transferable across countries is context-specific. For international environmental agreements, transfers might be technically feasible but politically challenging to implement. The additional complexity of partition functions is that any coalition structure π may allow for multiple equilibrium payoff vectors. To simplify the analysis, it is also possible to impose a fixed sharing rule within a coalition as in Ref. [101]. In our game, such fixed rules could be, for instance, proportional to the political power shares, an equal split across members, or even forbidding any within-coalition utility transfers altogether.

¹⁰In a characteristic function game without externalities, the payoffs are independent of the ambient coalition structure.

2.A.3 Farsightedness and bargaining friction

The discounting parameter δ serves two roles in our approach. It introduces bargaining frictions if rejected proposals cause delays in payoffs and controls the countries' level of farsightedness concerning possible future transitions. Countries receive an infinite stream of real-time payoffs based on the going state and their static payoff function $u_i(x)$. Their goal is to maximize the expected sum of discounted payoffs:

$$\mathbb{E}_P \sum_{t=0}^{\infty} \delta^t u_i(x_t), \quad (2.3)$$

where P is a stochastic process governing the state transitions (defined in more detail below). An unsuccessful round of bargaining means that any future gains from coalition formation will be delayed and discounted. A value of δ close to 1 implies that countries might be willing to push their payoffs far into the future or even endure temporary losses to better position themselves for future negotiations.

When choosing their propositions and responses, countries consider not just the one-step consequences of their actions but an entire sequence of new propositions, responses, and transitions that their current actions might trigger. A high level of δ implies that countries consider a long chain of possible future events when choosing their strategies. For simplicity, we assume that all countries share the same level of farsightedness.

2.A.4 Strategy profiles

A strategy profile defines the country-specific decision rules. We assume all strategies to be stationary and Markovian so that all proposals and responses only depend on the going state with no dependence on history or time. All countries have complete information regarding the game's details and other players' strategies.

Let $p_i(x, y, A)$ be the probability that country i proposes moving to state y to an approval committee A when chosen as the active proposer under a going state x . Similarly, let $r_i(x, y, A)$ be the probability that country i , if a member of the approval committee, accepts the proposed transition. The strategy profile is a collection of these proposer and responder distributions across all countries, $\sigma = \{p_i, r_i\}_{i=1}^n$. We can then condense all stochastic components of the model into

a single process P_ρ^σ that governs the evolution of our geoengineering game. In other words, the probability of any state transition depends on the protocol picking a particular country k to choose a proposition from p_k and all responding countries $j \in A$ replying according to r_j . The value function of any state x for country i under strategy profiles σ and protocol ρ is then:

$$V_i(x) = (1 - \delta)u_i(x) + \delta \sum_{y \in X} P_\rho^\sigma(x, y)V_i(y). \quad (2.4)$$

An equilibrium strategy profile is a collection of strategies such that no country wants to change their own strategy given the strategies of others. The existence of an equilibrium is guaranteed under the conditions of continuous payoff functions $u_i(x)$ and compact state space X [25, 102], but in general the equilibria are not unique¹¹. The following two consistency conditions must apply to all players for a strategy profile to constitute an equilibrium [26].

Condition 1. Proposer consistency.

$p_i(x, y, A) > 0$ only if $y \in F_A(x)$ and

$$(y, A) \in \arg \max_{(k, Q): k \in F_Q(x)} \{\Psi(x, k, Q)V_i(k) + (1 - \Psi(x, k, Q))V_i(x)\}$$

Condition 2. Responder consistency.

$$r_i(x, y, A) \begin{cases} = 1, & \text{if } V_i(y) > V_i(x) \\ = 0, & \text{if } V_i(y) < V_i(x) \\ \in [0, 1] & \text{otherwise.} \end{cases}$$

Above, $\Psi(x, k, Q)$ denotes the probability that the approval committee Q approves the proposed transition from state x to k . In words, the first condition states that for a country to propose a transition, there must be no other possible proposals to make that would yield a higher expected long-run payoff, given the strategies of

¹¹A stationary Markov equilibrium exists for finite X [26, Supplementary notes].

others. The second condition requires that countries never accept proposals that would take them to a state with a lower expected long-term payoff.

2.B Geoengineering deployment equilibria

The equilibrium of the underlying geoengineering deployment game depends on assumptions regarding the preferences of countries and the global effects and costs of geoengineering. This section outlines these assumptions.

2.B.1 Period payoffs are avoided temperature damages

Each country has a baseline temperature T_i^{Base} . Climate change is assumed to cause a uniform temperature increase by $\Delta > 0$,

$$T_i^{CC} = T_i^{Base} + \Delta. \quad (2.5)$$

Geoengineering reduces global temperatures, and we assume for clarity of the exposition that it reduces temperatures uniformly,

$$T_i = T_i^{CC} - G, \quad (2.6)$$

where we measure the geoengineering level G in terms of the uniform temperature effect. Following Ref. [87], we assume there is a uniform ideal temperature, T^{Ideal} , with climate damages increasing quadratically in the deviation from that ideal temperature,

$$u_i(T_i) = -d_i(T_i - T^{Ideal})^2 + K_i, \quad (2.7)$$

where K_i is a constant. We can rewrite Eq. (2.7) as a function of the global geoengineering level,

$$u_i(G) = d_i \cdot G \cdot (2\alpha_i - G), \quad (2.8)$$

where $\alpha_i = T_i^{Base} + \Delta - T^{Ideal}$ is the ideal SG level of country i and $K_i = d_i(T_i^{Base} + \Delta - T^{Ideal})^2$ are country i 's climate damages at the local temperature caused by

climate change. This choice of K_i ensures we can interpret u_i as country i 's benefit from SG. It is Eq. (2.8) that we use in the main text to calculate period payoffs for countries. In the main text we set $d_i = 1$ and $T^{Ideal} = 13^\circ\text{C}$. The constant term K_i could also capture any within-coalition transfer schemes, which we ignore in this paper.

2.B.2 Global geoengineering deployment

There are n countries in total. Apart from decisions on coalition formation, which is our main focus, countries choose geoengineering levels g_i . We assume that there are no direct costs of geoengineering. If only solar geoengineering is allowed, then $g_i \geq 0$. With counter geoengineering, $g_i \in \mathbb{R}$. Geoengineering efforts by different countries simply add up to the total global geoengineering level,

$$G = \sum_{i=1}^n g_i. \quad (2.9)$$

2.B.3 Best response functions

Here, we determine the best response function of a coalition. Let $C \subseteq N$ be a coalition that exceeds the minimal power threshold, $\sum_{i \in C} \gamma_i \geq \underline{\gamma}$, and let $G_{-C} \geq 0$ be the amount of cooling provided by all other countries. If the coalition falls short of the minimal power threshold, its best response trivially is zero. The coalition's choice variable is $G_C \geq 0$, and the final total geoengineering level will be $G = G_C + G_{-C}$. The coalition's objective is to maximize the power-adjusted sum of utilities of its members,

$$\max_{G_C \geq 0} \sum_{i \in C} \gamma_i \cdot (\kappa_i - d_i(G_C + G_{-C} - \alpha_i))^2. \quad (2.10)$$

The constant κ_i is irrelevant for the optimization, so Eq. (2.10) is equivalent to maximizing:

$$- \sum_{i \in C} \eta_i (G_C + G_{-C} - \alpha_i)^2, \quad (2.11)$$

where $\eta_i = \gamma_i d_i$ combines a country's political power and climatic exposure, measured in the marginal temperature damage. Eq. (2.11) is strictly concave in G_C with a maximum at

$$G_C(G_{-C}) = \max(\bar{\alpha}_C - G_{-C}, 0), \quad (2.12)$$

where

$$\bar{\alpha}_C = \left(\sum_{i \in C} \eta_i \right)^{-1} \sum_{i \in C} \eta_i \alpha_i \quad (2.13)$$

is the coalition's *weight-adjusted average ideal level*. Note that Eq. (2.13) simplifies to the arithmetic mean if all η_i are the same; also, a singleton's weight-adjusted ideal level is the original ideal level.

2.B.4 Nash equilibrium

Take any coalition structure $\pi = C_1 \cup \dots \cup C_K$. Assume the coalitions are ordered according to their weight-adjusted ideal level, i.e. $\bar{\alpha}_{C_1} > \dots > \bar{\alpha}_{C_K}$. For simplicity, we assume that the $\bar{\alpha}_{C_k}$ are pairwise-different. We can ensure pairwise different weight-adjusted average ideal levels by very small changes to either powers γ_i , marginal damages d_i and/or ideal deployment levels α_i .

We have determined the best response functions in Eq. (2.12). All best response functions have the same slope. They do not overlap because the $\bar{\alpha}_{C_k}$ are pairwise different. Taken together, this shows that the unique Nash equilibrium is:

$$G_{C_1}^* = \bar{\alpha}_{C_1}, \quad G_{C_k}^* = 0 \text{ for } k = 2, \dots, K. \quad (2.14)$$

The result is in line with the free-driver narrative: the coalition of countries with the highest ideal geoengineering level, here captured by the weight-adjusted average ideal level, implements that level and imposes it on everyone else. The total geoengineering level of coalition structure π can be written as:

$$G^*(\pi) = \max_k \{\bar{\alpha}_{C_k} | C_k \in \pi\}. \quad (2.15)$$

2.C Equilibrium strategy profiles

Tables 2.C.1 and 2.C.2 describe the equilibrium strategy profiles in full detail, corresponding to the two examples in Section 2.3. As noted in the main text, a strategy implies countries having contingency plans for all possible states where they might be chosen as the proposer or asked to respond to a proposal as a member of an approval committee. We construct the strategy profiles manually from an initial guess according to the pre-specified approval committee structures, and verify their correctness using Condition 1 and Condition 2 in section 2.A.4.

		Proposer W					Proposer T					Proposer C						
		()	(TC)	(WC)	(WT)	(WTC)	()	(TC)	(WC)	(WT)	(WTC)	()	(TC)	(WC)	(WT)	(WTC)		
()	prop.	1					1					1						
()	acce. W	1	Grey				0 0 0					0 0 0						
()	acce. T	0	1		1	1	1	Grey				1 1						
()	acce. C	0	1	Grey		1	0	1	Grey		1	1	Grey					
(TC)	prop.	1					1					1						
(TC)	acce. W	1	Grey				0 0 0					0 0 0						
(TC)	acce. T	0	1		1	1	1	1	Grey				1 1 1					
(TC)	acce. C	0	1		1	1	Grey		1	1	1	1	1	Grey				
(WC)	prop.	1					1					1						
(WC)	acce. W	1	1		Grey			1	1	Grey		1	0	Grey			1	0
(WC)	acce. T	0	Grey		0	1	Grey		1	Grey			0		0		1	
(WC)	acce. C	0	Grey		0	1	0	0	Grey		0	1	0	1		Grey		
(WT)	prop.	1					1					1						
(WT)	acce. W	1	Grey		1	Grey			1	0	Grey		0	Grey			1	0
(WT)	acce. T	0	1	Grey		1	0	Grey		1	Grey			0		0		1
(WT)	acce. C	0	1	Grey		1	0	1	Grey		1	Grey			1		Grey	
(WTC)	prop.	1					1					1						
(WTC)	acce. W	1		Grey			1	1	Grey		1	Grey			Grey			
(WTC)	acce. T	0	0		0	Grey			0		1		0			0		1
(WTC)	acce. C	0	0		0	Grey			0	0	Grey		0			0		1

Table 2.C.1: Strategy profile under the weak governance scenario. Light blue rows denote state transition proposals. Dark orange tiles indicate transitions the proposer can initiate without consulting others (maintaining the status quo or unilaterally leaving a coalition). Grey tiles indicate that a country is not a member of the corresponding approval committee. Light green tiles are transitions for which a country is in the approval committee and must provide an acceptance probability. Light orange tiles on the left correspond to the going state of the game. Numerical values indicate proposal and acceptance probabilities between 0 and 1.

		Proposer W					Proposer T					Proposer C					
		()	(TC)	(WC)	(WT)	(WTC)	()	(TC)	(WC)	(WT)	(WTC)	()	(TC)	(WC)	(WT)	(WTC)	
()	prop.	1					1					1					
()	acce.	W	1						1	1	1			1	1	1	
()	acce.	T		1		0	0	1					1		0	0	
()	acce.	C		1	0			0		0		1					
(TC)	prop.	1					1					1					
(TC)	acce.	W		1					1	1	1			1	1	1	
(TC)	acce.	T	0		0	0	0	0	1					0	0	0	
(TC)	acce.	C	0		0	0	0		0	0	0	0	1				
(WC)	prop.	1					1					1					
(WC)	acce.	W	0		1			0	0		1	1		0		1	1
(WC)	acce.	T		1		0	0		1				1		0	0	
(WC)	acce.	C		1		0	0	1	1		0	0	1		1		
(WT)	prop.	1					1					1					
(WT)	acce.	W	0			1			0	0		1		0	0	0	1
(WT)	acce.	T		1	1		1			1			1	1	1		1
(WT)	acce.	C		1	1		0		1	1		0			1		
(WTC)	prop.	1					1					1					
(WTC)	acce.	W		0			1	0	0		0		0	0	0		
(WTC)	acce.	T	1		1	0			1		1		1	1	1		
(WTC)	acce.	C	1		1	1		1	1		1				1	1	

Table 2.C.2: Strategy profile under the power threshold scenario. Light blue rows denote state transition proposals. Dark orange tiles indicate transitions the proposer can initiate without consulting others (maintaining the status quo or unilaterally leaving a coalition). Grey tiles indicate that a country is not a member of the corresponding approval committee. Light green tiles are transitions for which a country is in the approval committee and must provide an acceptance probability. Light orange tiles on the left correspond to the going state of the game. Numerical values indicate proposal and acceptance probabilities between 0 and 1.

2.D Additional examples

For the clarity of exposition, the examples presented in Section 2.3 assumed that the countries differ only in terms of their ideal temperature levels. We also contrasted the weak governance scenario with free-driver behaviour to only one possible governance scheme with a minimum power threshold. In this supplementary section, we provide two additional examples. First, we consider the same minimum power threshold rule as in the main text but now allow countries to differ in their marginal damages d_i . Then, building on the example with differentiated marginal damages, we consider an alternative set of institutional rules where a unanimous approval

committee is no longer needed, but a majority approval is enough to verify a state transition. As we will show, both extensions will directly affect the transition dynamics of the model and, in the latter example, also the predictions of the absorbing state.

Power threshold with country-specific marginal damages

Consider a setting similar to Section 2.3 with a minimum power threshold for SG deployment. All countries still have the same ideal temperature levels shown in Table 2.3.1. But whereas all marginal damages d_i are fixed to unity in the main text, we now consider $d_W = 0.75$, $d_T = 1.25$, and $d_C = 1.0$. Alternatively, we could introduce variation in the damage functions by considering additional terms (such as precipitation effects) or by considering different marginal impacts from excessive and insufficient levels of geoengineering.

The main consequence of the new marginal damages is that all countries now receive their highest possible period payoffs in different states: country W in (WT), T in the grand coalition (WTC), and C in (TC) (see panel A in Figure 2.D.1). The exact same strategy profile shown in Figure 2.C.2 still constitutes an equilibrium in this new example. The key thing to note is that country T still wants to break out of the grand coalition, even though it maximizes the period payoffs for T. This means that country C has a high degree of veto power on the system dynamics through being able to break the grand coalition and by T anticipating its strategy.

Consider the game in (WTC). Country T knows that W remains passive, as the warm country tries to maintain any high-deployment state for as long as possible. Therefore, the bargaining occurs essentially between the two natural partners, T and C, who have relatively similar ideal temperature levels and together enough political clout to deploy SG.

Since the countries are identical in power and the proposer protocol is a uniform distribution, there is a 1/3 change of C becoming the proposer in (WTC) and breaking the system into (WT), knowing that from there it is in the interest of T to further deviate into all singletons (), again with probability 1/3. Anticipating these moves and acting first, country T can first deviate from (WTC) to (WC), forcing C to perform the additional deviation to (). For T, such a sequence of states implies a

higher sum of cumulative payoffs due to the more favourable transition. Therefore, country T, anticipating the actions of C, also prefers to break out from the state where the SG deployment level would bring it closest to its ideal temperature, simply in anticipation of the strategies of others.

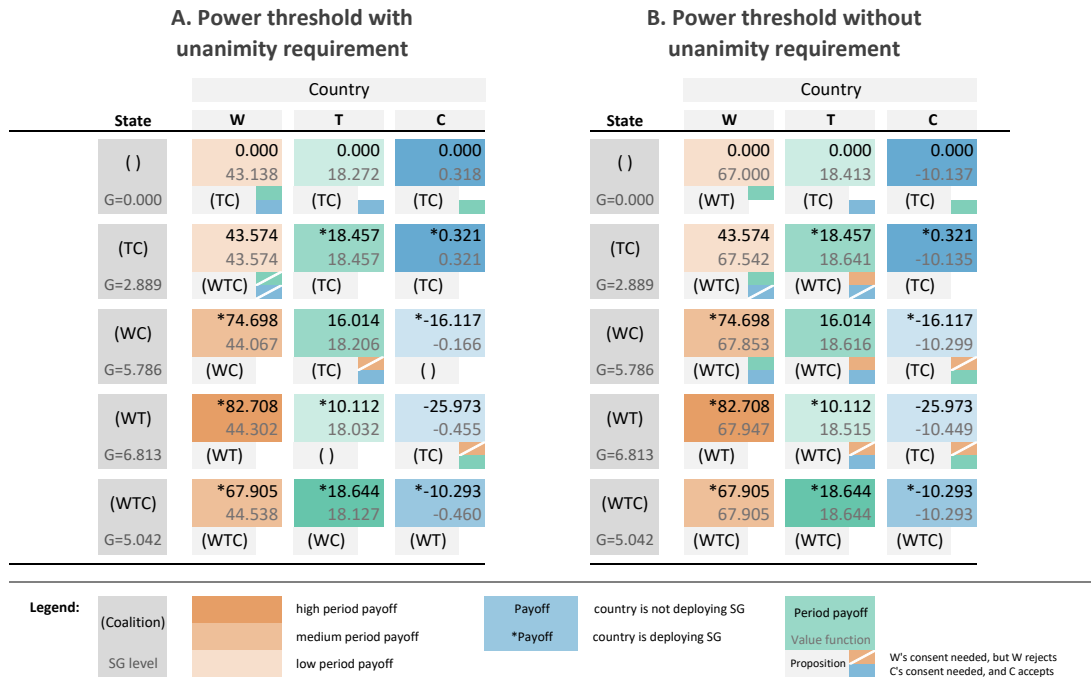


Figure 2.D.1: Summary of period payoffs and equilibrium strategy profiles under differentiated marginal damages with the unanimity requirement (panel A) and without it (panel B). Period-payoffs u_i , cf. Eq. (2.2), and long-term expected values given strategies V_i , cf. Eq. (2.1), are shown in red (country W), green (country T), and blue (country C) boxes. Dark and light fills represent more and less attractive constellations, respectively. An asterisk in front of the payoff indicates that the respective country is deploying SG. Every row is dedicated to a certain going state. The dark grey box on the left indicates the coalition structure and the total global SG level in that state. For every state, the light grey boxes show the proposal of the respective country in equilibrium. If that proposition needs approval by others, the countries needed for approval are indicated by small coloured boxes to the right of the proposal: solid filled if the respective country approves the proposal in equilibrium, filled with a diagonal line if the respective country rejects the proposal.

Power threshold with majority approval committees

Finally, consider the same setting as above with a minimum power threshold and differentiated marginal damages. However, this time, assume that a majority committee is enough to validate a state transition. More concretely, for our game, this implies that all new members must still approve a transition (e.g. W proposing the move from $()$ to (WTC) needs the consent of both T and C), but only half of the existing members (for instance, W can propose, and T can approve the transition from (TC) to (WTC) without consulting C).

Dropping the unanimity requirement can have a significant impact on the transition dynamics. Consider the strategies in panel B in Figure 2.D.1, under which the grand coalition (WTC) is an absorbing state. So far, country C always had an incentive to break the grand coalition, expecting T to eventually play along and move the system to the absorbing state (TC) . However, without the unanimity requirement, it would be in the interest of T to approve W entering (TC) , thus restoring the grand coalition. Since C needs to endure high transitory losses when aiming for (TC) to materialize, this strategy is no longer profitable without the unanimity requirement.

The resulting transition probabilities are presented in Figure 2.D.2, and the detailed strategy profile in Table 2.D.1. Note that all countries now want to avoid the singleton structure $()$. That is intuitive: since it is now easier for countries to trigger state transitions, there is less veto power from dropping out of any treaty. All countries anticipate that some SG deploying coalition will eventually emerge. It is, therefore, in their interest to join in and influence the thermostat setting instead of remaining outside. Note that in this scenario, the geoengineering deployment levels, and hence the period payoffs, remain unchanged from panel A in Figure 2.D.1 with the unanimity requirement still in place. On the other hand, the value functions for each state, as described in Eq. (2.4), change, due to the new transition probabilities.

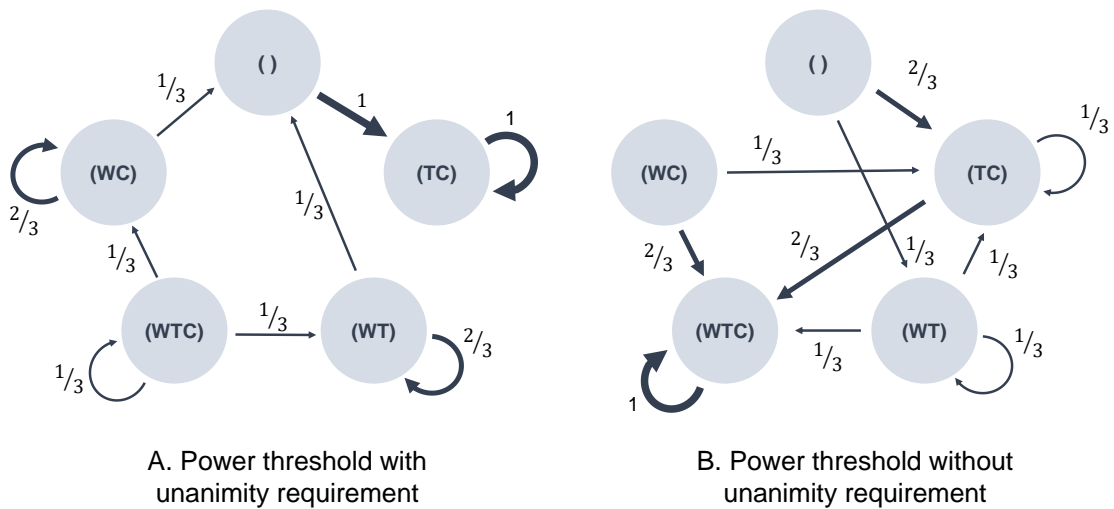


Figure 2.D.2: State transition probabilities under power threshold, country-specific marginal damages, unanimity requirement (panel A), and without unanimity requirement (panel B). The values in panel A correspond to the power threshold scenario in Section 2.3 and strategy profile in Table 2.C.2, repeated here for convenience. The transition probabilities in panel B follow from the equilibrium strategies shown in Figure 2.D.1. Without a unanimity requirement, the absorbing state is (WTC), as country C no longer has the incentive to deviate, knowing that countries W and T have the power to restore the grand coalition from the state (TC).

		Proposer W				Proposer T				Proposer C			
		() (TC) (WC) (WT) (WTC)	() (TC) (WC) (WT) (WTC)	() (TC) (WC) (WT) (WTC)	() (TC) (WC) (WT) (WTC)								
()	prop.	1				1				1			
()	acce. W	1					1	1	1		1	1	1
()	acce. T		1		1	1					1		1
()	acce. C		1	0		0			0	1			
(TC)	prop.	1				1				1			
(TC)	acce. W		1				1	1	1		1	1	1
(TC)	acce. T	0		0	0	1				0	0	0	1
(TC)	acce. C	0		0	0	0			0	0	1		
(WC)	prop.	1				1				1			
(WC)	acce. W	0		1		0	0		1	1		0	1
(WC)	acce. T		1		0	1					1		0
(WC)	acce. C		1		0	1	1		0	1	1		
(WT)	prop.	1				1				1			
(WT)	acce. W	0			1		0	0		0	0	0	
(WT)	acce. T		1	1		1			1		0	1	1
(WT)	acce. C		1	1		1	1		1			1	
(WTC)	prop.	1				1				1			
(WTC)	acce. W		0			1		0	1		0	0	0
(WTC)	acce. T	0		0	0			0		1	0	0	0
(WTC)	acce. C	1		0	0		1	1	0			0	1

Table 2.D.1: Strategy profile without unanimity requirement under a power threshold and differentiated marginal damages. Light blue rows denote state transition proposals. Dark orange tiles indicate transitions the proposer can initiate without consulting others (maintaining the status quo or unilaterally leaving a coalition). Grey tiles indicate that a country is not a member of the corresponding approval committee. Light green tiles are transitions for which a country is in the approval committee and must provide an acceptance probability. Light orange tiles on the left correspond to the going state of the game. Numerical values indicate proposal and acceptance probabilities between 0 and 1. For the strategy profile under unanimity requirement, see Table 2.C.2.

Chapter 3

Natural Disasters and Economic Dynamics: An Application to Tropical Cyclones

Abstract

This chapter presents a framework for estimating the long-run economic impacts of natural disasters. The approach combines a probabilistic disaster modelling platform with a dynamic general equilibrium model of the economy. We apply the methodology to study the effects of tropical cyclones in the United States, the Caribbean islands, Japan, China, and the Philippines. Our results show that the post-disaster recovery can take several decades, with notable cumulative adverse effects for the frequently exposed regions. For instance, cyclone activity reduces long-run aggregate consumption between 0.3 - 22 %, depending on the region. To evaluate the robustness of our results, we extend the model with two additional scenarios. First, we consider endogenous economic productivity gains from specialization. Second, we add a scenario where climate change alters the intensity and frequency of future disasters. The extensions affect the numerical results but do not change the qualitative conclusions.

This chapter is joint work with Clément Renoir. An earlier overview of the methodology and results has been published in Ref. [103].

3.1 Introduction

Tropical cyclones are among the costliest natural catastrophe events, causing approximately USD 50 billion in damages per average year over the past two decades [27]. In addition to their direct effects through loss of lives and destruction of economic assets, tropical cyclones can permanently shape regional growth dynamics by causing prolonged reconstruction periods [104], forced relocation [105], or even changing risk attitudes toward future disasters [106]. Future damages could be even higher as coastal development increases the value of exposed assets [107] and as climate change modifies cyclone intensity and frequency [108]. However, estimates regarding cyclones' long-term impacts on economic dynamics are sometimes inconclusive due to the complex interplay of direct and indirect effects over time.

We present a new framework for estimating the long-run economic impacts of tropical cyclones. First, we employ a probabilistic disaster impact model [109] to quantify the direct annual losses to regional capital stocks using historical and synthetic cyclone tracks. Then, we feed these estimations of capital destruction into a dynamic, multi-sectoral general equilibrium model of the economy [110]. With the economic model, we track the evolution of regional capital stocks and GDP components over time, analysing both the immediate disaster impacts and the long-run cumulative effects. The combination of models allows us to derive globally consistent cyclone damage estimates, providing new insights into the magnitude of tropical cyclones' role in long-term economic development.

We gather data from different sources. The disaster impact model uses satellite data from the International Best Track Archive for Climate Stewardship database [111] for historical cyclone tracks. We also use satellite imagery for estimating the spatial distribution of exposed assets based on nighttime light intensities [112]. The global grid resolution of the resulting damage estimates is approximately 10×10 km. We calibrate the economic model with the Global Trade Analysis Project dataset. It provides a sectoral decomposition of economic activities and bilateral trade flows for 129 world regions. We focus our analysis on the US, the Caribbean islands, Japan, China, and the Philippines. These regions are frequently hit by tropical cyclones and show considerable heterogeneity in size, economic structure, and overall cyclone exposure.

In our setting, the immediate economic response after a cyclone strike is a jump in aggregate investments and a consequent drop in consumption to replace the damaged capital stock. Ref. [113] refers to this tradeoff as "*forced investment*" since the reconstruction efforts can spur economic activity while there is still an overall reduction in welfare. Although the reconstruction in our model is relatively quick, it can take several decades to catch up with the benchmark growth path where no shocks occur. For some economic variables, such as the aggregate output and consumption levels, the post-disaster trajectory remains below the reference path for the entire simulation period.

Our results also highlight the dissimilarity between cyclone impacts on GDP and welfare. Whereas the long-run average drop in consumption ranges from 0.3% in the US to more than 20% in the Philippines, the respective GDP reductions are only 0.1% and 6%. Since consumption is our model's sole determinant of welfare, the GDP changes alone fail to capture the cyclones' full welfare effect. GDP, in this case, masks the opposite impacts cyclones have on individual GDP components, producing artificially small aggregate changes [114].

A long-run quantification of tropical cyclones' economic impacts is likely to be sensitive to assumptions regarding *i*) the underlying mechanisms of economic growth and *ii*) the effects of climate change on future disasters. The general equilibrium model allows us to consider growth dynamics based on either physical capital accumulation or knowledge creation with endogenous productivity gains from specialization. Compared to physical capital accumulation, the knowledge-based growth engine dampens the negative cyclone impacts due to higher productivity and additional incentives for investing in new capital varieties. Under this growth specification, in extreme cases, cyclones' long-run effect on GDP can even become positive due to very high investment levels and the within-sector positive spillovers that follow.

Finally, to study the role of climate change on future cyclone intensity and frequency, we recalibrate the regional cyclone damage distributions under two Representative Concentration Pathway (RCP) scenarios. Under the intermediate RCP4.5 scenario, cyclone intensity increases in the North Atlantic basin (the US Atlantic coast and Caribbean islands) with no changes in event frequency, driving up the total economic losses. On the other hand, cyclone frequency in the Northwestern

Pacific basin (Japan, China, and the Philippines) falls with only a slight increase in intensity, leading to lower mean damages at the end of the century. Under the high-emissions RCP8.5 scenario, cyclone damages increase in all considered regions. In the US, for instance, the aggregate capital stock in 2100 under RCP8.5 is approximately 1% lower compared to the same year under constant climate conditions.

Related literature

The empirical work linking natural disasters and economic growth is somewhat inconclusive. For instance, Ref. [115] finds a positive relationship between climatic disaster frequency and economic growth. On the other hand, Ref. [104], analyzing tropical cyclones, finds a systematic and substantial negative long-run impact on national incomes with no clear evidence of a rebound effect during the two decades following the catastrophe. Moreover, recurring shocks make it impossible to catch up with the initial GDP trend. Ref. [116] focuses on the Central American and Caribbean regions and estimates the average hurricane to reduce the output growth rate by roughly 0.84%.

Several works attempt to reconcile the diverging empirical evidence. One explanation is the difference between the *risk* of disasters occurring and the consequences after experiencing a disaster *strike* [117, 118]. Whereas disaster strikes can cause output losses due to capital destruction and business interruptions, disaster risk might induce higher precautionary savings, thereby inflating the economy's growth rate. However, Ref. [119] casts doubt on the possible role of precautionary savings in explaining the positive relationship between economic growth and disasters. A positive relationship would require unreasonably high values of intertemporal substitution elasticity and relative risk aversion when calibrating an endogenous growth model with data from the US.

Another explanation for the inconclusive empirical findings lies in the relative damages disasters might cause depending on the capital variety. As cyclones are particularly destructive to physical capital, frequent disasters might steer investments towards accumulating human capital instead, thereby enhancing productivity [115, 117, 120]. Whether natural disasters mainly affect productive capital stocks or durable consumption goods might also play a role [121]. Losing productive capital

harms economic performance, whereas only replacing damaged durable goods can boost output, potentially pushing GDP above the pre-disaster level.

A third possible explanation for positive tropical cyclone impacts is the process of creative destruction. As Ref. [117, p.90] writes, “*By destroying old factories and roads, disasters allow new and more efficient infrastructure to be built, providing an opportunity for the economy to transform itself into a more productive one in the long run.*” Older capital vintages might also be more susceptible to disaster damages than newer variants, amplifying the effect [122]. However, much depends on the affected region and the economic sector. For instance, Ref. [123] finds that the creative destruction effect only occurs in sufficiently developed economies. According to Ref. [124], on the other hand, storms can cause significant damage to agriculture, while capital stock upgrades only boost industrial growth. Other mechanisms might also dampen the productivity gains from creative destruction. For instance, small firms might not afford the business interruptions and worker re-training that are often necessary when replacing lost capital goods with new variants [113, 125].

Several studies also highlight the role of institutions as a determinant of disaster impacts. Ref. [126], for instance, finds that countries with higher-quality institutions suffer fewer disaster-related deaths. Education, trade openness, and financial system maturity also matter for disaster resilience [127–129]. The high institutional quality helps endure the initial catastrophe shock and enables faster deployment of resources for reconstruction, thus reducing negative disaster spillovers to the broader economy [130].

Specialized cyclone impact models provide another way to estimate the disasters’ long-run economic consequences [107, 131]. In particular, future losses might increase as coastal development inflates the value of exposed assets and as climate change alters the intensity and frequency of disasters. Although rich in spatial detail, these analyses frequently rely on predefined GDP projections to quantify long-term effects. However, as disasters become increasingly harmful, they are more likely to affect consumption, investment patterns, and the underlying growth trajectories. Models featuring fixed economic growth paths cannot – by design – capture these feedback mechanisms.

This chapter contributes to the literature by considering the impacts of cyclone strikes on long-run economic development. Empirical works such as Ref. [104] pro-

vide insight into the causal effect of cyclone shocks on GDP. However, they have to deal with several confounding factors and only analyze the effects in partial equilibrium. On the other hand, disaster impact models such as Ref. [107] represent cyclone damages in great detail but typically do not capture the economic adjustments over time. Our general equilibrium approach uses the spatial detail of a full disaster impact framework while capturing the endogenous recovery period dynamics. All model components rely on global datasets, allowing us to consistently evaluate country-level impacts, incorporating direct damages and the secondary effects through trade linkages. Finally, through changes in a single elasticity parameter, the economic model can also capture different assumptions regarding the underlying determinants of growth, from a standard capital accumulation setting to an endogenous representation where gains from specialization drive growth.

3.2 Methodology

This section describes the disaster impact framework, the general equilibrium model of the economy, and the integration of the two systems. We focus our study on five regions: the US, the Caribbean islands, Japan, China, and the Philippines. These regions are frequently exposed to tropical cyclones and vary drastically in the structure and size of their economies. All model components rely on globally consistent datasets, which makes extending the regional coverage of the analysis straightforward.

3.2.1 Disaster impact model

Quantifying the direct disaster impacts requires data describing the exposures (the spatial distribution of vulnerable physical assets) and the hazards (tracks and wind speeds of historical cyclones). We access both datasets and conduct the analysis using the open-source CLIMADA (CLIMate ADaptation) platform [109, 132].

Exposure

To estimate the annual disaster impacts, we first need to construct the spatial distribution of physical assets in all regions of our study. We use the *LitPop* model

[133], which combines nighttime light satellite imagery with gridded population accounts to obtain a globally consistent estimate of the asset distribution.

Satellite imagery is convenient for its public availability, global spatial coverage, and frequent update schedule. Our nighttime light intensity data come from NASA's Black Marble suite [112], and we use 2016 as the base year. However, some known caveats exist in using light intensities as a proxy for economic activity. These include, among others, high measurement errors in luminosity data, saturating pixel values, and bright pixels leaking light into their adjacent pixels, thus inflating their value [133, 134].

To overcome some of the above issues, the LitPop model supplements the nighttime light data with global population estimates from the Gridded Population of the World database [135]. The database provides globally disaggregated population counts with a resolution of up to 1×1 km. We give equal weight to the light intensity (Lit) and population data (Pop) and compute the share of the physical assets (A_i) in each pixel i out of N total pixels for a given country as:

$$A_i = \frac{Lit_i^n Pop_i^m}{\sum_i^N (Lit_i^n Pop_i^m)},$$

where $m = 1$ and $n = 1$ are tuneable share parameters. Since we model cyclone damages as destroyed capital stock, we use the value of the produced capital stock from the World Bank wealth accounts as the region-specific indicator for aggregate physical asset value [136]. Hence, the value of physical capital per land area is the product of the country's total capital stock value and the pixel-specific capital share. Figure 3.2.1 illustrates the resulting distribution of exposed capital stocks for the US state of Florida.

Hazard

Next, we estimate the cyclone damages based on historical cyclone tracks. We obtain the path and the maximum sustained wind speed of each recorded cyclone from 1950 to 2019 from the International Best Track Archive for Climate Stewardship (IBTrACS) database [111]. The dataset contains 6'907 tropical cyclone tracks, illustrated in Figure 3.2.2. Among them, 1'079 happened in the North Atlantic basin (containing the Caribbean islands and the US Atlantic coast), and




Figure 3.2.1: Estimated distribution of capital stocks in the US state of Florida. Each pixel is weighted according to its nighttime light intensity and population density. For each region in our sample, we distribute the aggregate produced capital stock value according to the pixel-specific shares.

2'040 in the Northwestern Pacific Ocean basin (containing Japan, China, and the Philippines). For each historical cyclone in the IBTrACS database, we construct 50 synthetic tracks as perturbed random walks under parameters controlling their distance from the original observations. The synthetic tracks inherit several features from their historical counterparts, such as changes in wind speeds on landfall, which are relevant for damage computation. These additional synthetic data improve the probabilistic description of the annual cyclone activity compared to using only historical storm tracks. The augmented dataset contains 55'029 and 104'040 tropical cyclone events for the North Atlantic and the Northwestern Pacific basins, respectively.

We use maximum wind speed to measure storm intensity, a common choice in the literature [137–140]. The storm intensity allows us to construct a proxy for capital destruction caused by each cyclone. We use the cubic damage function from Ref. [141] to translate cyclone wind speeds into capital destruction. The fraction of capital damaged by storm j at location i , and time t , $\delta_{i,j,t}$, varies with wind speeds V exceeding a threshold value:

$$\delta_{i,j,t} = \frac{v_{i,j,t}^3}{1 + v_{i,j,t}^3}, \quad (3.1)$$

with a normalized wind speed:

$$v_{i,j,t} \equiv \frac{\max\{V_{i,j,t} - V_{thresh}, 0\}}{V_{half} - V_{thresh}}. \quad (3.2)$$

Similar to Ref. [141], we set the wind speed below which there are no damages at $V_{thresh} = 25.7m/s$. The parameter $V_{half} = 74.7m/s$ determines the wind speed that destroys 50% of the capital stock [142]. We aggregate the damages $\delta_{i,j,t}$ by year to compute statistics such as the mean and the standard deviation of damages for all regions.




Figure 3.2.2: Global tropical cyclone activity for 1950-2019 based on the IBTrACS database [111]. The intensity levels from lower to higher wind speeds are tropical depression (TD), tropical storm (TS), and hurricanes of category 1 to 5 (Cat. 1-5) on the Saffir-Simpson scale.

Ultimately, we want to represent cyclone damages as annual economic shocks. Although the synthetic cyclone tracks extend our pool of disaster events, our dataset still contains estimates for only seventy years of cyclone activity (from 1950 to 2019). The small sample size may prohibit us from constructing a realistic damage distribution. To overcome this limitation, we enlarge our sample of annual damage estimates by creating a set of synthetic years. More specifically, we create 5'000 additional synthetic years of cyclone activity by assuming the disaster frequency

to follow a Poisson distribution [118, 143] and resampling a corresponding number of random events from the collection of synthetic and historical cyclone tracks. Appendix 3.C describes the resulting distributions in more detail.

Table 3.2.1 summarizes the damage estimates for the historical sample and the augmented sample containing synthetic tracks and all 5'070 synthetic years for each region of our study. We detect no systematic bias between the historical and augmented samples. The three statistics we use to compare the two samples are roughly similar, although the augmented sample has a much broader set of annual damages. Table 3.2.1 also highlights the considerable regional variation in relative cyclone-induced capital losses, ranging from an additional depreciation rate of 0.15% in the US to 2.67% in the Philippines.

	USA	CAR	JPN	CHN	PHL
Historical sample					
% of events causing damage	29.84	13.07	19.41	33.28	22.21
Mean damage, $\bar{\delta}_{TC}$ (%)	0.15	0.59	1.10	0.40	2.67
Std. of damages, σ_{TC}	0.21	1.51	1.82	0.51	6.43
Augmented sample					
% of events causing damage	26.82	13.50	18.68	28.74	21.64
Mean damage, $\bar{\delta}_{TC}$ (%)	0.09	0.65	0.84	0.40	2.73
Std. of damages, σ_{TC}	0.14	1.67	1.51	0.50	5.67

Table 3.2.1: Yearly cyclone damage statistics by region. The historical sample only contains the cyclone observations between 1950-2019. The augmented sample corresponds to the 5'070 synthetic years that we simulate from the historical observations and the perturbed synthetic tracks.

3.2.2 Economic model

This section describes the main features of the model economy, as well as the datasets used for calibration.

Numerical framework

We employ a dynamic, multi-regional and multi-sectoral numerical general equilibrium model based on Ref. [144] and Ref. [110]. The production structure of the economy consists of *i*) final good producers, *ii*) producers of intermediate goods, and *iii*) producers of intermediate composites. The separation between intermediate goods and intermediate composites is one of the framework’s key features. It enables switching on endogenous productivity gains from specialization with a simple change of model parameters. The time horizon of the theoretical model formulation is infinite with discrete increments but approximated using a finite number of periods in the numerical implementation. A detailed technical description of the model is available in Appendix 3.A.

Each regional economy consists of a forward-looking representative household, maximizing the discounted sum of utility from consumption over time. Households also own all firms and factors of production. Labour and capital are mobile across sectors (but not regions), and all countries are open to trade. We model international trade using the Armington assumption [145], which treats goods produced in different regions as imperfect substitutes. As opposed to a model with small open economies, our trade specification allows changes in regional production and demand patterns to affect world prices. Consequently, disaster impacts in one country can spill over to other regions via global supply chain links. The model consists of nested constant elasticity of substitution (CES) blocks that combine domestic and imported goods from various sectors into consumption aggregates and production input bundles.

Economic accounts and model calibration

We calibrate the economic model using the Global Trade Analysis Project (GTAP) database [146]. GTAP provides unified base-year economic accounts for 129 regions, 57 commodities, and five primary production factors. The dataset describes the flow of goods across sectors and regions and how the regional agents allocate them between final demand, intermediate production inputs, or trade. We use the GTAP data as a static snapshot of the economy and extrapolate—using a set of exogenous parameter assumptions—an initial balanced growth path on which all sectors grow

at the same rate. The dataset also includes sectoral greenhouse gas emissions, which allows constructing additional climate policy scenarios. Appendix 3.A.6 contains details on the sectoral and regional aggregation of the raw GTAP data.

In addition to the dollar-valued economic accounts from GTAP, the general equilibrium model requires various sector- and region-specific elasticity parameters. These include the elasticity estimates for consumer demand and the substitution elasticities between different production inputs. We use estimates mainly from the MIT Economic Projection & Policy Analysis model [147] and Ref. [146]. The numerical values are available in Appendix 3.A.3.

3.2.3 Model integration

In summary, we can describe our modelling framework as follows. We generate data on cyclone activity in all five regions of our study based on historical and synthetic cyclone tracks. By combining the cyclone tracks data with a wind-based damage function and the spatial distribution of economic assets, we compute the capital destruction caused by each cyclone in the sample. We aggregate this capital destruction estimate by year and obtain a distribution over the annual capital depreciation due to cyclone exposure for each economy. Finally, we construct a pool of synthetic cyclone activity years from which we draw the cyclone shocks. We consider the region-specific shocks as an unexpected annual increase in the natural depreciation level of capital. We first calibrate our economic growth model to a balanced growth path without cyclones. Finally, introducing the shocks, we can run counterfactual simulations and compare how the economic trajectories differ between the reference growth path and the one affected by cyclones.

Numerical general equilibrium models provide a flexible instrument for analysing the multi-sectoral adjustment of prices after an economic shock. Their deterministic structure, however, imposes some limitations on modelling the impacts of rare natural disasters such as tropical cyclones. Introducing disaster impacts in an arbitrary time step t , without further adjustments, would imply that for the periods preceding t , all agents in the model have perfect information over the timing and magnitude of the upcoming event. Agents would then react to disasters with optimal precautionary savings, producing an overly optimistic description of disasters'

consequences.

We choose a modified solution algorithm that maintains the forward-looking nature over the model's economic variables but treats the disaster realisations as unanticipated shocks. To model an unanticipated disaster occurring at time τ , we first solve for a reference equilibrium path without shocks from the initial period t_0 to the terminal time T , such that $t_0 < \tau < T$. We then fix all the variables from the reference equilibrium until τ and re-run the model with the shock. In other words, we only allow the agents to adjust their behaviour in the periods $t \geq \tau$. In the newly constructed sub-model, from τ to T , the shock occurs in the first period of the simulation, such that the agents have no chance of anticipating the shock. We combine the solution from the reference equilibrium and the one from the sub-model by using the reference equilibrium values for $t < \tau$ and the sub-model values for $t \geq \tau$. In the absence of shocks, this algorithm produces the same numerical results as only simulating the reference equilibrium path.

3.3 Results

We analyze the simulation results in four parts. First, we study the impulse response of the economy after a single year of cyclone activity. It illustrates the primary economic mechanisms and provides intuition for the recovery period dynamics under a single fixed-magnitude shock. Second, for the paper's main results, we run the model with recurring random shocks to study the cumulative long-run disaster effects. We then provide two sensitivity scenarios to scrutinize the central long-run modelling assumptions: the underlying drivers of economic growth and the role of climate change. We first recalibrate the general equilibrium model to introduce endogenous productivity gains from specialized capital varieties. Finally, we alter the regional cyclone damage distributions to consider how climate change might affect the intensity and frequency of future disasters. We present all numerical results as counterfactual simulations to the corresponding no-shock baseline economy.

3.3.1 Impulse response to a single cyclone shock

Consider first the effect of an individual cyclone shock. Figure 3.3.1 decomposes the general equilibrium response into expenditure-side GDP contributions. We simulate a shock at time $t = 5$ that increases capital depreciation compared to a year without disaster events. The magnitude of the shock for all countries is one standard deviation above the regional mean, as described in Table 3.2.1.¹ The overall picture is relatively similar for all regions. GDP falls on impact, followed by a catch-up period of faster growth and reinvestment. The higher depreciation increases the marginal productivity of capital, bringing greater returns on investment. Consequently, savings increase as a response to the reconstruction efforts. However, the increase in savings comes at the expense of lower consumption, reducing welfare. The magnitude of the investment jump ranges from 0.06% in the US to more than 4% in the Philippines.

International trade linkages are another determinant of the recovery period's shape and duration. Following the disaster, the trade balance deteriorates in all regions. Countries use more imports to facilitate reconstruction efforts, while exports suffer from the lost local production capacity and increased domestic investment demand. Thus, the trade channel highlights the additional flexibility that international openness can provide in the disaster aftermath. For most regions, the trade volumes converge relatively quickly to their original levels, closing the gap between the benchmark trajectory within a few years of the shock.

For the post-disaster periods, capital depreciation returns to its natural level, and the regional economies gradually return to their original steady-state path. Reconstruction is often relatively fast. For instance, reaching the pre-disaster capital stock level takes three years in the Philippines, whereas the US reinvests the lost capital amount already in the first post-disaster period. However, compared to the benchmark economic trajectory that evolves without interruptions, the catch-up recovery period can take up to several decades. Although the aggregate capital stock eventually reaches the reference trajectory, households spread the required additional investments over multiple years to avoid a drastic drop in consumption. As a result, the consumption (and GDP) levels remain permanently below the ref-

¹The shock distributions have a strong positive skew, producing relatively low mean damages.




Figure 3.3.1: Impulse response to a single cyclone shock. The magnitude of the shock in all regions is one standard deviation above the regional mean, as described in Table 3.2.1. The change in GDP is the sum of changes in aggregate consumption, investments, government expenditure, and net trade. All values are relative to a benchmark economy that grows on a steady-state path without shocks.

erence path in all regions. The long-run gap is 0.01% in the US, but up to 1.6% in the Philippines.

That the recovery back to the pre-disaster growth path can take decades might sound surprisingly slow. In reality, however, several factors can contribute to long recovery times. The first is the limited reconstruction capacity. With insufficient financial resources, a rapid reconstruction can only come with a sudden drop in consumption. Therefore, households' tendency to smooth consumption patterns over time directly leads to allocating the reinvestments over longer periods. Technical limitations, such as the lack of a sufficient reconstruction workforce, are also possible. Moreover, business cycles can further amplify the effects, particularly if a disaster strikes during a high cycle where available resources are already scarce [113]. There can also be significant production factor rigidity, especially between-sector

capital immobility, that complicates reconstruction efforts. The long recovery times are a common finding both in empirical works [104] and studies based on numerical general equilibrium simulations [148].

Governments usually assume an active role in the disaster aftermath. Yet changes in public demand are absent from the results shown in Figure 3.3.1. Since the cyclone impacts in our framework occur solely through losses in capital stock, the primary recovery mechanism is reinvesting. However, our numerical model makes no distinction between the private sector and government investments. Therefore, the numerical results we present on post-disaster aggregate investment levels include the increased public investment demand. Although our framework is flexible enough to consider additional transfer schemes from the government to households [148], we ignore them since the shock process does not automatically trigger any. Introducing these measures would require ad-hoc assumptions on government payouts and their effects. In reality, however, transfers such as medical payments and unemployment support can significantly increase in response to disaster events [149].

Finally, disasters' consequences unquestionably go beyond their impacts on physical assets and direct loss of lives. In addition, disasters cause traumatic injuries, stress, and diseases that can have long-lasting effects on welfare and productivity [150, 151]. However, these effects are likely to vary depending on the local institutions and the type of disaster. As this paper focuses on constructing a globally consistent modelling framework, we omit these effects from the numerical model but acknowledge that their unmeasured cost can be substantial.

3.3.2 Cumulative effect of recurring cyclone shocks

Whereas the previous section illustrated the model dynamics, we now turn to the paper's main results quantifying the long-run cumulative cyclone impacts. We randomly draw annual capital depreciation shocks from the augmented disaster event pool constructed in Section 3.2.1. We run 500 Monte Carlo simulations of the economic model for each region and provide aggregate results over a 30-year period.

When cyclone shocks are frequent and random, the economy is constantly adjusting to new conditions. Capital depreciation is, therefore, always above its natural level, hampering growth. Figure 3.3.2 shows the cumulative impact on aggregate




Figure 3.3.2: Change in aggregate consumption by region with recurring cyclone shocks. The solid lines represent the means of 500 Monte Carlo runs. The shaded areas denote the 5th to 95th percentile ranges. All values are relative to benchmark economies that grow on a steady-state path without shocks.

consumption. In the US, where direct cyclone damages typically only occur in specific regions, the cyclone-induced consumption drop is approximately 0.3% after 30 years of simulation. However, the reduction can be significantly greater for regions where a larger share of the capital stock is exposed. For instance, consumption in the Caribbean islands is more than 3% below the baseline level, whereas, in the Philippines, the long-run reduction exceeds 20%.

Compared to prior works, the magnitude of the results appears reasonable. For instance, Ref. [152] finds that typhoon exposure in the Philippines leads to an approximately 7% drop in the next year’s household expenditure. On the other hand, the US Congressional Budget Office estimates the country’s annual hurricane damages at 0.16% of GDP [153].

Appendix 3.B illustrates the results for additional economic variables. Investment levels are consistently above the no-disaster baseline trajectory for all regions, reflecting the dynamics explained in the stylized single-shock scenarios above. Sim-

ilarly, industry output and capital intensity remain consistently below the reference values for the entire simulation horizon.

Notably, compared to the changes in aggregate consumption and investments, the impacts on GDP appear relatively small. For instance, the long-run GDP in the US is only 0.1% below the reference growth trajectory, compared to a drop of 0.3% in consumption. The difference highlights that the long-term welfare implications of tropical cyclones are likely to be higher than what the GDP impacts alone might suggest. Ref. [114] also finds similar results, where collapsing the cyclone effects to a single GDP value masks the heterogeneous impacts cyclones have on macroeconomic activity through consumption, investment, and trade patterns. Nevertheless, the unambiguous finding here is that cyclone activity negatively affects GDP, consumption, and welfare in all regions. In terms of equivalent variation, cyclones reduce welfare by 0.35% in the US, 4.25% in the Caribbean islands, 6.2% in Japan, 3.4% in China, and 30.3% in the Philippines.

3.3.3 Extensions

Productivity gains from specialization

Our regional sample consists of heterogeneous countries in terms of size and economic structure. It is, therefore, important to consider alternative assumptions regarding the underlying drivers of growth as an explanation for varying disaster effects. Different mechanisms might allow, for instance, some countries to exhibit a post-disaster growth spurt due to "build back better" dynamics. In contrast, others may never recover to their original growth trend.² The capital structure is also likely to play a role. Whether the exposed regions and sectors rely primarily on physical, natural, or human capital can significantly affect the overall cyclone impacts and recovery dynamics.

We model the endogenous productivity gains with a simple model reparameterization. Instead of considering capital only as the physical stock, we make a broader interpretation of a capital composite that includes both the physical stock and the immaterial knowledge capital. Intermediate firms can invest in new sector-specific

²See Ref. [104] for a thorough discussion on different post-disaster recovery hypotheses.

capital varieties. The varieties are imperfectly substitutable, such that the intermediate firms make positive profits due to a monopoly markup. Similarly to the growth dynamics in Ref. [154], firms investing in new varieties receive a perpetual blueprint for their product. Compared to the previous section, where growth is solely due to physical capital accumulation, there is now an additional incentive for conducting R&D investments. The positive spillovers from specialization, on the other hand, enhance the overall economic growth rate.

We simulate the economies with the same shock realizations as in the previous section but now turn on the endogenous gains from specialization. We again use 500 Monte Carlo runs and report the results after 30 years of simulation. The resulting growth trajectories vary significantly compared to the previous section, as summarized in Table 3.3.1. Overall, the endogenous growth engine substantially dampens the negative impacts of repetitive cyclone events. Under gains from specialization, post-disaster investments increase more than under the standard case. As a result, the capital stock, although damaged by the same amount, gets rebuilt faster, leading to lower capital losses in the long run. The existing capital stock is also more productive, affording a faster reconstruction and attenuating the overall drop in consumption and welfare.

For most regions, the GDP reduction after 30 years of cyclone activity is more than 50% smaller than under the standard scenario. The consumption losses are between 17% to 26% smaller than before, depending on the region. For extreme shock realizations, GDP impacts under endogenous productivity gains can even become temporarily positive for the most affected regions, driven by the higher investment levels and increasing capital returns. However, even under these extreme realizations, the overall consumption impact remains negative. GDP only appears higher as the destroyed capital stocks are not measured in GDP, whereas the reconstruction efforts are. Even with gains from specialization, the long-run average GDP impacts remain negative for all regions.

The above results suggest that the substitutability between physical and knowledge capital can be an important determinant of long-run cyclone impacts, as discussed in prior studies [115, 120]. However, despite the considerable reduction in negative impacts, the effect does not appear strong enough to turn the cyclone impacts positive. Therefore, the overall qualitative findings from the previous section

	USA	CAR	JPN	CHN	PHL
Consumption	-23.18	-24.50	-26.30	-22.30	-17.21
Investments	0.60	7.72	24.60	14.80	51.73
Cap. intensity	-7.25	-4.22	-6.55	-5.73	-12.96
Capital Stock	-15.45	-24.88	-25.20	-24.67	-44.90
Industry output	-45.54	-66.04	-63.01	-59.73	-84.22
GDP	-45.03	-57.16	-56.73	-59.24	-87.95
Welfare	-16.28	-20.40	-20.45	-17.87	-14.77

Table 3.3.1: Change in mean cyclone impacts after 30 years under endogenous gains from specialization compared to a model without productivity gains. All values denote a percentage change difference relative to the cyclone impacts in Section 3.3.2. For instance, compared to the scenario without productivity gains, the aggregate long-term consumption drop in the US is 23.18% smaller, the reduction in capital stock 15.45% smaller, and the increase in investments is 0.60% larger. Welfare impacts are measured in terms of Hicksian equivalent variation.

remain unchanged even under the alternative growth engine.

Climate change impacts

Climate change can make weather-related extreme events more frequent and intense [4]. At the same time, economic growth and coastal development can increase future cyclone damages as there are more assets in harm's way [107]. For instance, the US Congressional Budget Office estimates the country's hurricane damages to rise from the current levels of USD 28 billion per year to USD 39 billion by 2075, attributing half of the increase to climate change and another half to further coastal development [153]. Therefore, in the final scenario, we study how our framework's main long-run economic variables react to the assumptions regarding future climate conditions.

In the following, we extend our model horizon from 30 to 80 years for the climate change impacts to take effect. We run the model in ten-year increments to compensate for the resulting increase in computational cost. We consider two possible greenhouse gas concentration pathways: the RCP4.5 with intermediate

emissions and the high-emission RCP8.5 scenario. We calibrate a new decadal damage distribution for each concentration scenario by tuning the cyclone intensity and frequency values based on Ref. [155]. Appendix 3.D documents the steps in more detail.

In the RCP4.5 scenario, we assume the cyclone intensity in the North Atlantic basin to increase by 4.5% by the late twenty-first century. In contrast, over the same period, there is no significant change in cyclone frequency [155]. As a result, in our framework, the mean cyclone damage δ_{TC} gradually increases, leading to a 30% higher value for the US and 24% higher for the Caribbean islands by 2100 compared to constant climate conditions. For the Northwestern Pacific basin, on the other hand, the cyclone frequency falls by 34.5%, but the intensity increases by 5.5%. The result is a slight increase in mean cyclone damages for the first decades, while in the longer run, the frequency effect dominates, decreasing the mean damages. The end-of-century reduction in mean damages varies from 9.7% in Japan to 13.5% in the Philippines. For the RCP8.5 scenario, the cyclone intensity and frequency increase for both basins, leading to a substantial increase in modelled damages.

We conduct 1000 Monte Carlo sample trajectories for each RCP scenario, randomly drawing shocks in each period as in the previous section. We run the analysis only for the US and Caribbean islands, focusing only on the North Atlantic cyclone basin regions.

	RCP4.5		RCP8.5	
	USA	CAR	USA	CAR
Mean damage, δ_{TC}	30.16	24.03	112.92	87.95
Std. of damages, σ_{TC}	25.94	18.72	93.11	66.39
Consumption	-0.10	-0.57	-0.25	-1.19
Capital stock	-0.41	-2.22	-1.09	-5.18
GDP	-0.05	-0.20	-0.12	-0.39

Table 3.3.2: Percentage change in tropical cyclone impacts in 2100 under different climate scenarios compared to estimates in the same year under constant climate conditions.

Table 3.3.2 reports the main simulated variables at 2100. The increasing cyclone

intensity directly leads to stronger capital depreciation shocks for both countries in both RCP scenarios. For instance, under RCP4.5, the average aggregate consumption levels are 0.1% (USA) and 0.69% (CAR) lower than at the same time under constant climate conditions. The impacts are even more pronounced in the RCP8.5 scenario. The mean cyclone damage more than doubles for the US and increases by almost 90% for the Caribbean islands. As a result, the end-of-century capital stocks are almost 1% and 5% smaller for the two regions, respectively, compared to an average simulation without climate impacts.

3.4 Discussion

We have constructed a modelling framework to represent probabilistic, region-specific cyclone damage functions in a dynamic economic growth model. That enables us to isolate how tropical cyclone strikes affect economic variables over time in general equilibrium. The chapter's main goal was to set up a globally consistent modelling framework. However, several possible extensions remain for studying additional cyclone impact channels and the role of public policies in disaster impact management.

First, we have excluded the role of adaptation. In our model, local adaptation measures could affect either the cyclone impact function (for instance, the construction of sea walls, mangrove restoration, or the implementation of new building codes) or the distribution of exposed assets (such as spatial planning in high-risk areas), and therefore have interesting broader economic effects. The optimal allocation of investments between productive capital stocks and unproductive adaptive capital can also produce interesting results for the long-term dynamics [156].

In our coupled system, tropical cyclones only enter the economy through damages to capital stock. In reality, cyclone impacts are much more complex. Disaster strikes might reduce the economy's total factor productivity, for instance, through electricity blackouts [118] or business interruptions [148]. Moreover, without additional assumptions, the general equilibrium response leaves out many secondary impacts. For instance, a drop in post-disaster quality of public services might make attracting workers more difficult, directly affecting the recovery period [157].

Throughout, we have considered countries as the smallest units of regional aggre-

gation. With appropriate economic accounts, it is possible to achieve even greater regional detail. For instance, Ref. [158] uses a sub-national model to study flood impacts in Italy, and Ref. [148] builds a numerical general equilibrium model calibrated to a single city. Greater regional detail might allow studying relevant local questions such as labour reallocations or comparing specific adaptation alternatives. We also assume that the cyclone shocks are uniform over all sectors of the economy. Combining the damage estimates with additional land use data enables deriving sector-specific shocks to represent, for instance, the increased vulnerability of agricultural sectors to natural disasters [124].

Finally, we acknowledge some caveats. First, by design, our model economy is always in equilibrium. That might be unrealistic, especially in the periods directly after a disaster where bottlenecks and resource misallocations are likely to happen in all markets. Immediately reaching an equilibrium might make the early stages of the recovery path overly optimistic, ignoring some real-world rigidities. We also assume that the regional distribution of exposed assets remains constant throughout the simulation. In reality, there might be considerable shifts if people and firms leave the most exposed areas or if urban expansion gradually covers new regions. We have also explicitly focused on tropical cyclone impacts based on wind intensity estimates. However, incorporating the simultaneous effects from cyclones, storm surges, and in the long term, even sea level rise, might give a more complete picture of the disaster impacts. The impact function parameterization in Eq. (3.2) also relies on data solely from the US and might not be directly applicable to other regions.

3.5 Conclusion

We develop a framework for quantifying the long-run economic impacts of tropical cyclones. Our approach combines a dynamic general equilibrium model of the economy with damage estimates from a probabilistic disaster impact platform. The coupled system allows us to consider region-specific damage functions and post-disaster recovery profiles, assuming that cyclone strikes enter the economy through unexpected capital depreciation shocks.

We apply our framework to five regions: the US, the Caribbean islands, Japan,

China, and the Philippines. The general findings are similar for all regions. By destroying capital stocks, cyclones lead to increased reconstruction investments, thus reducing consumption levels, GDP, and welfare. After 30 years of recurring cyclone shocks, the aggregate capital stock in the US is approximately 0.5% smaller compared to a no-shock baseline path. In the Philippines, which is the most affected region in our sample, the difference is almost 13%. Our results also highlight the need for disentangling the GDP and welfare effects. Changes in post-disaster GDP appear artificially small since GDP aggregates cyclones' broad macroeconomic impacts with opposite signs into a single value, masking the full welfare effect.

Assumptions regarding the economic growth engine and the role of future climate change affect the numerical results but do not change the overall qualitative findings. When endogenous productivity gains from specialization drive growth, mean cyclone impacts are smaller but still unambiguously negative. Under the RCP4.5 climate scenario, cyclone damages increase in the North Atlantic Ocean basin (the US and the Caribbean islands). However, they fall slightly in the Northwestern Pacific basin (China, Japan, and the Philippines) by the end of the century. Under a high-emission climate scenario (RCP8.5), cyclone damages increase in all regions compared to the current climate conditions.

Appendix

3.A General equilibrium model

We employ a dynamic, multi-regional, and multi-sectoral numerical general equilibrium model following Ref. [144] and Ref. [110].

3.A.1 Production

In each region, we model the economy's production structure as the interaction between three agents: *i*) final good producers, *ii*) producers of intermediate goods, and *iii*) producers of intermediate composites. The markets for the final goods and intermediate composites are perfectly competitive, whereas the market for intermediate goods can also be monopolistic. Figure 3.A.1 provides an overview of the nested production structure.

Final good production

The final good producers in sector i , region r , and time t produce output $Y_{i,r,t}$ according to the following constant elasticity of substitution (CES) function:

$$Y_{i,r,t} = \left[\alpha_{i,r} Q_{i,r,t}^{\frac{\sigma_{i,r}-1}{\sigma_{i,r}}} + (1 - \alpha_{i,r}) B_{i,r,t}^{\frac{\sigma_{i,r}-1}{\sigma_{i,r}}} \right]^{\frac{\sigma_{i,r}}{\sigma_{i,r}-1}}. \quad (3.3)$$

Above, $Q_{i,r,t}$ is the sector-specific composite intermediate good. $B_{i,r,t}$ denotes the composite final good from all other sectors needed for producing in sector i , capturing how different sectors (and regions) interact through a complex network of

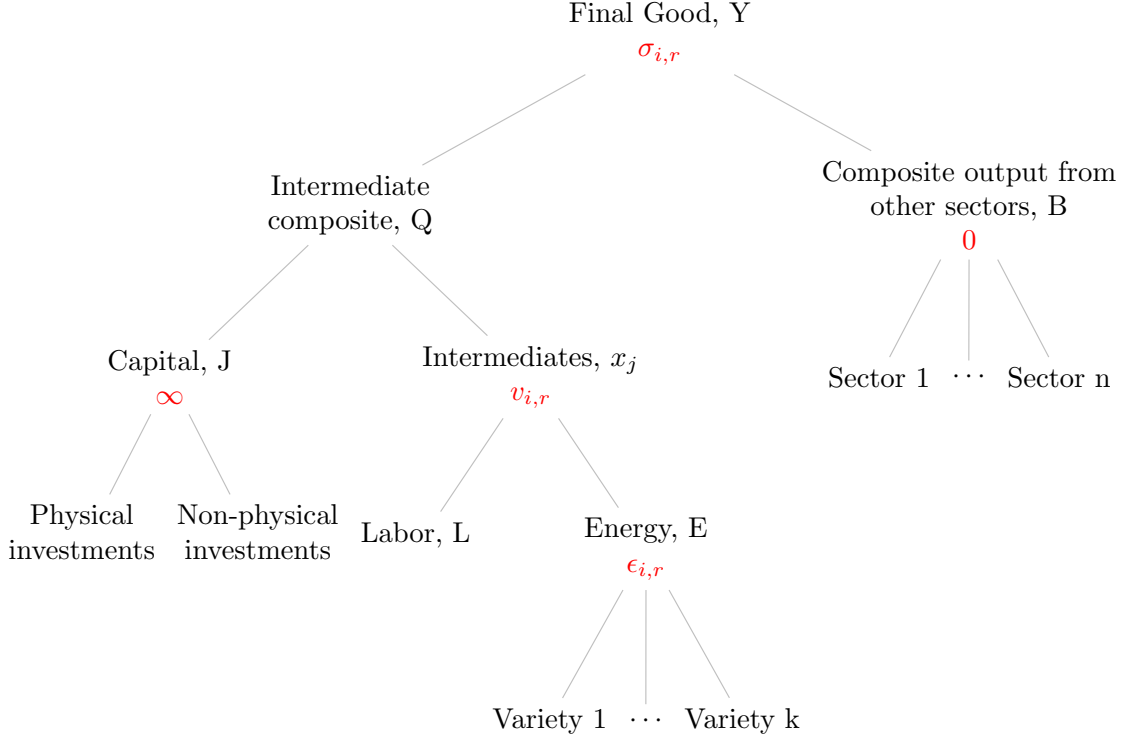


Figure 3.A.1: Production structure of the economy. Each production nest assumes a constant elasticity of substitution structure. The elasticity value within a sub-nest is highlighted in red.

value chains. Outputs from different sectors are assembled to $B_{i,r,t}$ according to a Leontief-type production function, that is, in fixed proportions. Parameters $\alpha_{i,r}$ determine the constant value shares between $Q_{i,r,t}$ and $B_{i,r,t}$ in the production function. The elasticity of substitution between the two types of inputs is $\sigma_{i,r}$. All elasticity and share parameters are sector- and region-specific.

In each sector, the final good producer maximizes profits in a perfectly competitive market by solving:

$$\max_{Q_{i,r,t}, B_{i,r,t}} p_{i,r,t}^Y Y_{i,r,t} - p_{i,r,t}^Q Q_{i,r,t} - p_{i,r,t}^B B_{i,r,t}, \text{ w.r.t 3.3}, \quad (3.4)$$

where $p_{i,r,t}^Y$, $p_{i,r,t}^Q$ and $p_{i,r,t}^B$ denote the prices of final goods, intermediate composite, and other inputs, respectively. Solving Eq. (3.4) and combining the resulting

demand functions for $Q_{i,r,t}$ and $B_{i,r,t}$ yields the following condition for the optimal input use:

$$\frac{Q_{i,r,t}}{B_{i,r,t}} = \left(\frac{\alpha_{i,r}}{1 - \alpha_{i,r}} \right)^{\sigma_{i,r}} \left(\frac{p_{i,r,t}^B}{p_{i,r,t}^Q} \right)^{\sigma_{i,r}}. \quad (3.5)$$

According to Eq. (3.5), an increase in the price of one input type will increase the share of the other input in the optimal bundle. Generally, we assume imperfect substitutability between different input types.

Intermediate composite production

In the second step of the production nest, producers of a sector-specific intermediate composite assemble their output $Q_{i,r,t}$ by combining different varieties of individual intermediate goods according to a standard Dixit-Stiglitz production function:

$$Q_{i,r,t} = \left[\int_{j=0}^{J_{i,r,t}} x_{j,i,r,t}^{\kappa} dj \right]^{\frac{1}{\kappa}}, \quad (3.6)$$

where $x_{j,i,r,t}$ denotes the j^{th} type of intermediate good variety that is available in sector i . $J_{i,r,t}$ denotes the sector-specific capital stock. We treat new innovations (that is, new varieties of $x_{j,i,r,t}$) as new varieties of capital, so new types of $x_{j,i,r,t}$ also imply an expansion in the capital stock. This specification gives us two channels through which the intermediate sector can induce growth in the overall economy. One is to produce a larger amount of any single variety $x_{j,i,r,t}$ by employing more labour and energy. The other is to expand the number of available varieties through investments in the sector-specific capital stock. The parameter κ measures the substitutability between different varieties, or equivalently, the gains from specialization. Setting $0 < \kappa < 1$ allows the increasing number of varieties to enhance final sector productivity in an endogenous manner [154], whereas $\kappa = 1$ switches off these productivity gains.

The producer of the intermediate good composite $Q_{i,r,t}$ maximizes profits on a competitive market, taking all prices as given, and solving:

$$\max_{x_{j,i,r,t}} p_{i,r,t}^Q Q_{i,r,t} - \int_{j=0}^{J_{i,r,t}} p_{j,i,r,t}^x x_{j,i,r,t} dj, \text{ w.r.t 3.6,} \quad (3.7)$$

where $p_{j,i,r,t}^x$ is the price of intermediate varieties. Solving the optimization problem in Eq. (3.7) determines the optimal demand for $x_{j,i,r,t}$ as:

$$x_{j,i,r,t} = \left(\frac{p_{i,r,t}^Q}{p_{j,i,r,t}^x} \right)^{\frac{1}{1-\kappa}} Q_{i,r,t}. \quad (3.8)$$

From here onwards, we assume that all varieties of the sector-specific intermediate good are perfectly symmetrical, i.e. $x_{j,i,r,t} = x_{i,r,t}$.

Intermediate good production

As described in Eq. (3.6), the amount, variety, and substitutability between different intermediate goods determine the expansion of each production sector i . We assume that a new firm first invents an intermediate variety $x_{i,r,t}$, and then produces it under a perpetual patent from the moment of invention. Therefore, the growth rate of the overall economy depends on the decisions of profit-seeking intermediate firms. To describe these intermediate firms in full, we need to describe both their optimal output decision for the already invented varieties, as well as their incentives to innovate new varieties.

i) Optimal output of existing varieties

To produce one unit of output, the intermediate good producer combines two types of inputs, labour $L_{i,r,t}$ and energy $E_{i,r,t}$, according to the following CES technology:

$$x_{i,r,t} = J_{i,r,t} \left[\lambda_{i,r} L_{i,r,t}^{\frac{v_{i,r}-1}{v_{i,r}}} + (1 - \lambda_{i,r}) E_{i,r,t}^{\frac{v_{i,r}-1}{v_{i,r}}} \right]^{\frac{v_{i,r}}{v_{i,r}-1}}, \quad (3.9)$$

where parameters $\lambda_{i,r}$ denote the value shares and $v_{i,r}$ the substitution elasticities. From Eq. (3.9), there are within-sector spillover effects from the expanding capital

stock $J_{i,r,t}$. We assume the supply of labour to be inelastic throughout the modelling horizon, mobile between sectors within a country, but immobile between countries. The energy aggregate $E_{i,r,t}$, on the other hand, is a combination of K available energy varieties:

$$E_{i,r,t} = \left[\sum_{k \in K} \phi_{k,i,r}(Z_{k,i,r,t})^{\frac{\epsilon_{i,r}-1}{\epsilon_{i,r}}} \right]^{\frac{\epsilon_{i,r}}{\epsilon_{i,r}-1}}. \quad (3.10)$$

We denote the amount of every energy input $k \in K$ by $Z_{k,i,r,t}$, and the respective price by $p_{k,r,t}^Z$. The output decision of the intermediate monopoly can be derived from two parts. First, it chooses an optimal bundle of labour and energy inputs under competitive profit-maximizing conditions:

$$\max_{L_{i,r,t}, Z_{k,i,r,t}} \psi_{i,r,t}^x x_{i,r,t} - w_{r,t} L_{i,r,t} - \sum_k p_{k,r,t}^Z Z_{k,i,r,t}, \quad (3.11)$$

where $\psi_{i,r,t}^x$ is the price that would prevail under a perfectly competitive market. The firm, however, exploits its monopoly power in the output market and sets the optimal output price solving:

$$\max_{p_{i,r,t}^x} p_{i,r,t}^x x_{i,r,t} - \psi_{i,r,t}^x x_{i,r,t}, \quad (3.12)$$

taking the demand for $x_{i,r,t}$ in Eq. (3.8) as given. Thus, it sets prices according to:

$$p_{i,r,t}^x = \frac{1}{\kappa} \psi_{i,r,t}^x, \quad (3.13)$$

making profits of:

$$\pi_{i,r,t} = (1 - \kappa) p_{i,r,t}^x x_{i,r,t}. \quad (3.14)$$

This brings us to an alternative definition of the substitutability term κ . As the individual intermediate goods $x_{i,r,t}$ are imperfect substitutes, and the intermediate good producers compete in a monopolistic market with an output price $p_{i,r,t}^x$ and markup $\frac{1}{\kappa}$, we can consider $(1 - \kappa)$ as the profit fraction of revenues from the intermediate composite sector going to the households that own the firms.

ii) Investments to new varieties

The capital stock consists of physical and non-physical capital, which together make up the sector-specific capital composite $J_{i,r,t}$. Firms conduct innovation by investing an amount $I_{i,r,t}$ to this composite capital good. Access to the investment market is unrestricted. New innovations then occur until the marginal cost of investments to the composite capital equals the firm's value so that no real profits remain. As in Ref. [154], we assume the knowledge capital from the innovation process to be non-rival but partially excludable with the use of patents. The equation of motion of the capital stock is:

$$J_{i,r,t+1} = I_{i,r,t} + (1 - \delta_{i,r,t})J_{i,r,t}, \quad (3.15)$$

with $\delta_{i,r,t}$ denoting the capital depreciation rate. The depreciation parameter has a particular role in our work as it ultimately depends both on the baseline depreciation rate and the exposure to cyclones that varies by year and region.

Finally, the capital accumulation process requires introducing a no-arbitrage condition. New firms (capital varieties) emerge as a result of household investment. In equilibrium, households must be indifferent between investing in a new firm and a riskless loan with a return $r_{i,r,t}$. As in standard endogenous growth models based on expanding input varieties, the value of the monopolist firm, that is, the value of owning a technology blueprint, is equal to the discounted value of all future profits. In our setting, this is equal to the cost of investing in a new firm. We can write the relationship between the new firm value $V_{i,r,t}$, instantaneous profits $\pi_{i,r,t}$, and the interest rate as $r_{i,r,t}$ with the following asset value equation: ³

³For details on deriving the relationship, see e.g. Ref. [159, Ch. 13]

$$\pi_{i,r,t} + \Delta V_{i,r,t} = r_{i,r,t} V_{i,r,t}, \quad (3.16)$$

where $\Delta V_{i,r,t}$ denotes the change in firm value. We can then extend Eq. (3.16) by writing:

$$\underbrace{\pi_{i,r,t}}_{\text{Direct return}} + \underbrace{\frac{p_{i,r,t+1}^J}{1+r_{i,r,t}} - p_{i,r,t}^J - \delta_{i,r,t} p_{i,r,t}^J}_{\text{Time } t \text{ present value of the capital gain}} = \underbrace{r_{i,r,t}}_{\text{Interest rate of a riskless loan}} \times \underbrace{V_{i,r,t}}_{\text{Firm value = investment cost}}, \quad (3.17)$$

where p_t^J is the price of capital. The intermediate good producer borrows from households to pay the innovation activities in advance. We can also re-write the sectoral profits from Eq. (3.14) as:

$$\pi_{i,r,t} = \underbrace{(1-\kappa)}_{\text{Monopoly profit share}} \underbrace{p_{i,r,t}^Q Q_{i,r,t}}_{\text{Sectoral revenue}} / \underbrace{J_{i,r,t}}_{\text{Number of varieties}}. \quad (3.18)$$

Inserting Eq. (3.18) into Eq. (3.17) then yields the expression for equilibrium interest rates and thus completes the no-arbitrage condition.

3.A.2 International trade

Our baseline dataset contains economic accounts of 129 regions, covering most of the global economy. Representing how different countries interact through international trade is, therefore, a central feature of our underlying general equilibrium model and an important determinant of how countries can adapt to economic shocks.

All final sectors of the economy are open to international trade. That is, all producers can employ both domestically produced and imported inputs, and consumers can purchase both domestic and imported consumption goods. To give more structure to the representation of international trade, we follow the Armington approach [145], which is a standard assumption in the numerical general equilibrium literature. With this approach, the suppliers of the final good use both domestically produced and imported goods as inputs in creating an Armington aggregate, which is the final good demanded in the economy. The domestic and imported inputs

are treated as imperfect substitutes. Intuitively, this means that consumers in any country can prefer domestically produced goods over imported varieties. More importantly, this allows for a realistic description of international trade, where any production sector in any region can simultaneously be an exporter and an importer of the same good, which is what we also observe in real economies.

More formally, denoting domestic sectoral production in region r by $D_{i,r,t}$ and imports from region s to r by $M_{i,s,r,t}$, the Armington aggregate is given by:

$$A_{i,r,t} = \left(\zeta_{i,r} D_{i,r,t}^{\frac{\eta_{i,r}-1}{\eta_{i,r}}} + (1 - \zeta_{i,r}) \left(\left[\sum_{s \neq r} m_{i,s,r} M_{i,s,r,t}^{\frac{\phi_{i,r}-1}{\phi_{i,r}}} \right]^{\frac{\phi_{i,r}}{\phi_{i,r}-1}} \right)^{\frac{\eta_{i,r}-1}{\eta_{i,r}}} \right)^{\frac{\eta_{i,r}}{\eta_{i,r}-1}}, \quad (3.19)$$

where we denote by $\zeta_{i,r}$ the share of domestic goods, and by $m_{i,s,r}$ the share parameters of different regions in the basket of imports. Parameters $\eta_{i,r}$ and $\phi_{i,r}$ are the respective substitution elasticities. With $p_{i,r,t}^A$ being the price of the Armington composite, and $p_{i,r,t}^Y$ the price of the domestic output, the profit maximization the final suppliers face is then:

$$\max_{D_{i,r,t}, M_{i,s,r,t}} p_{i,r,t}^A A_{i,r,t} - p_{i,r,t}^Y D_{i,r,t} - \sum_{s \neq r} p_{i,s,t}^A M_{i,s,r,t}. \quad (3.20)$$

We allow countries to run either trade surpluses or deficits, as also observed in the baseline dataset.

3.A.3 Preferences

For each region, we assume an infinitely-lived, forward-looking representative household. The household derives utility from consumption according to a standard constant intertemporal elasticity of substitution function:

$$U = \sum_{t=0}^{\infty} \left[\frac{1}{1 + \rho} \right]^t \frac{C_{r,t}^{1-\theta} - 1}{1 - \theta}, \quad (3.21)$$

where ρ denotes the time discounting parameter and θ the inverse of the intertem-

poral elasticity of substitution. As the economy consists of multiple production sectors, $C_{r,t}$ is a CES aggregate of the sector-specific consumption goods. Figure 3.A.2 illustrates the nested consumption structure.




Figure 3.A.2: Nested consumption structure. The substitution elasticity value within a sub-nest is highlighted in red.

The household also owns all firms in the economy, so its budget reads:

$$p_{r,t}^C C_{r,t} = w_{r,t} L_{r,t} - T_{r,t} - \sum_i p_{i,r,t+1}^J J_{i,r,t+1} + \sum_i (1 + r_{i,r,t}) p_{i,r,t}^J J_{i,r,t}, \quad (3.22)$$

where $w_{r,t}$ denotes the wage rate, $T_{r,t}$ a lump-sum tax which ensures the public budget to remain balanced, and $p_{r,t}^C$ is the price for the consumption aggregate.

3.A.4 Calibration details

Our model calibration follows the steps outlined in Ref. [160] and Ref. [161]. The key goal of the calibration process is to use the GTAP dataset as a static snapshot of the economy and extrapolate —using a set of exogenous parameter assumptions— a balanced growth path on which all sectors, and therefore all regional economies, grow at the same rate.

The household’s problem involves maximizing the stream of utility over time

in Eq. (3.21). The optimization is subject to the economy's production function $F(K_t, L_t)$, a resource constraint $F(K_t, L_t) = I_t + C_t$ dividing the output between consumption and investment, and the capital stock law of motion from Eq. (3.15). Assuming constant returns to scale and perfectly competitive markets, we can derive the following price relationships from the first-order optimality conditions [161]:

$$p_t = \left[\frac{1}{1 + \rho} \right]^t \frac{\partial U(C_t)}{\partial C_t}, \quad (3.23)$$

$$p_t^K = p_t \frac{\partial F(K_t, L_t)}{\partial K_t} + (1 - \delta)p_{t+1}^K, \quad (3.24)$$

$$p_t = p_{t+1}^K. \quad (3.25)$$

We can interpret these values as p_t being the price of output, $p_t \frac{\partial F(K_t, L_t)}{\partial K_t} =: R_t$ the rental rate of capital, that is, the value of the marginal product, and p_t^K the price of buying one new unit of capital. This distinction between capital stocks and capital services is central to the modelling approach: households own the stock, invest by buying new units of capital, and rent the capital to firms at the rate R_t .

Assuming a baseline interest rate \bar{r} , the calibration makes use of a declining reference price trajectory $p_t^{\text{ref}} = (1/(1 + \bar{r}))^t$. Then, for all prices at any arbitrary time instance τ , we have that:

$$p_{\tau+1} = \frac{p_\tau}{1 + \bar{r}}. \quad (3.26)$$

We can use the reference price path to further highlight the distinction between the capital rental and purchase prices. Combining Eq. (3.26) and Eq. (3.24) gives:

$$R_t = p_t^K \left(1 - \frac{1 - \delta}{1 + \bar{r}} \right), \quad (3.27)$$

which states that the capital rental price R_t is equal to the price of buying a new capital unit, subtracting the discounted value of the depreciated stock in the subsequent time period. Further normalizing $p_0 = 1$ allows us to write the first-period rental rate as $R_0 = \delta + \bar{r}$.

The benchmark GTAP data does not provide the capital stock values directly but only the base year capital earnings, denoted with V_0^K . Using $V_t^K := R_t K_t$

and the base year rental rate from above, we can derive the initial capital stock as $K_0 = V_0^K / (\delta + \bar{r})$. The next task is to calibrate the initial investments on a balanced growth path. Assuming a constant capital stock growth rate γ_K , we can write the next period capital stock either as in Eq. (3.15) or with $K_{t+1} = (1 + \gamma_K)K_t$. Combining the two equations gives $I_t = (\gamma_K + \delta)K_t$, such that on a balanced growth path, the annual investment level must cover both the capital growth rate and depreciation. Plugging in the definition of K_0 from before gives an initial investment level of:

$$I_0 = (\gamma_K + \delta)K_0. \quad (3.28)$$

The calibration process so far follows the standard conventions of numerical general equilibrium modelling. In our setting, however, the possibility of considering productivity gains from specialization requires some additional steps. When the gains from specialization are active, we assume that the size of the capital stock directly corresponds to the number of capital varieties. Moreover, the different varieties are imperfect substitutes, determined by the substitution elasticity parameter κ . The imperfect substitutability then creates monopoly rents and additional incentives for investing. The growth rate of output γ_Y , therefore, depends on two factors: an exogenously specified capital growth rate γ_K and an endogenous growth part determined by κ . The relationship between these parameters satisfies:

$$1 + \gamma_Y = (1 + \gamma_K)^{\frac{1}{\kappa}}. \quad (3.29)$$

Whenever $0 < \kappa < 1$, the output growth rate γ_Y exceeds the capital growth rate γ_K . To make this difference in growth rates compatible with the balanced growth path and to avoid situations where investments grow faster than the actual stock of capital, we also make the base depreciation rate time-dependent. This assumption is necessary for the balanced growth path to exist, but there is also an appealing intuition behind the adjustment. Namely, as the economies develop further, their capital stock grows more specialized and more susceptible to depreciation. The base depreciation rate is:

$$\delta_{r,t} = \left(\frac{1 + \gamma_Y}{1 + \gamma_K} \right)^t \delta_{r,0} + \gamma_K \left(\left(\frac{1 + \gamma_Y}{1 + \gamma_K} \right)^t - 1 \right), \quad (3.30)$$

which collapses to a constant value when $\kappa = 1$ and the gains from specialization are switched off.

Finally, to obtain the baseline consumption growth rate $g_r = \frac{C_{r,t+1}}{C_{r,t}}$ on the balanced growth path calibration, we can maximize Eq. (3.21) with respect to Eq. (3.22) to obtain the standard Keynes-Ramsey rule:

$$g_r \equiv \left[\frac{1 + \bar{r}}{1 + \rho} \right]^{\frac{1}{\theta}}. \quad (3.31)$$

According to Eq. (3.31), a higher interest rate \bar{r} boosts growth by inducing more saving, whereas a higher discount rate ρ incentivizes present consumption, therefore reducing the rate of growth. A higher intertemporal substitution elasticity $1/\theta$ also increases growth rates as the households become more willing to tolerate consumption variability in response to interest rate changes. In our setting, Eq. (3.31) also implicitly pins down the temporal discount rate ρ .

3.A.5 Numerical implementation

We follow Ref. [162] and formulate the general equilibrium economy as a mixed complementary problem (MCP). The formulation includes three types of inequality constraints: market-clearing conditions, zero-profit conditions, and income balance conditions. Each equilibrium condition f has a complementary variable z , such that the following conditions always hold: $f(z) \geq 0, z \geq 0, z^T f(z) = 0$. For instance, we can write the market-clearing condition as $f(p) = S(p) - D(p)$, where we use the price level p as the complementary variable, and supply and demand functions S and D , respectively. When the market clears, $f(p) = 0$, the equilibrium prices are positive. However, if supply exceeds demand, the complementary variable (prices) becomes zero. Similarly, for the zero-profit conditions, the complementary variable is the output level. As long as sectoral profits are non-negative, the output level

is positive. With negative profits, however, firms exit the market and the output becomes zero.

Although the theoretical model considers an infinite time horizon, the numerical implementation requires using a finite approximation. This introduces the risk of horizon effects affecting the equilibrium outcome as we approach the terminal period. To remedy the risk around the terminal period, we employ the method from Ref. [160]. This technique imposes an additional constraint on capital accumulation at the terminal period T to approximate the infinite horizon equilibrium. We introduce the post-terminal capital stock as an additional variable and require that the growth rate of investments in the terminal period mirror the output growth rate:

$$\frac{I_T}{I_{T-1}} = \frac{Y_T}{Y_{T-1}}. \quad (3.32)$$

That is, we only fix the growth rate of investments, and do not have to fix the actual growth rate, nor the terminal level, of capital stock. To further reduce terminal effects, we always discard a set of years from the end of the simulation results.

We use the programming language GAMS (General Algebraic Modeling System) as well as the MPSGE (Mathematical Programming System for General Equilibrium, [163]) sub-system to implement the economic model. To solve the model, use the PATH numerical solver [164].

3.A.6 Regional aggregation

Aggregate region	GTAP region
USA	United States of America
Japan	Japan
Philippines	Philippines
Caribbean	Rest of Caribbean ¹
China	China, Hong Kong
Rest of the World	Australia, New Zealand, Rest of Oceania, Republic of Korea, Mongolia, Taiwan, Rest of East Asia, Cambodia, Indonesia, Lao PDR, Malaysia, Singapore, Thailand, Viet Nam, Rest of Southeast Asia, Bangladesh, India, Nepal, Pakistan, Sri Lanka, Rest of South Asia, Canada, Mexico, Rest of North America, Argentina, Bolivia, Brazil, Chile, Colombia, Ecuador, Paraguay, Peru, Uruguay, Venezuela, Rest of South America, Costa Rica, Guatemala, Honduras, Nicaragua, Panama, El Salvador, Rest of Central America, Austria, Belgium, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Poland, Portugal, Slovakia, Slovenia, Spain, Sweden, United Kingdom, Switzerland, Norway, Rest of European Free Trade Association, Albania, Bulgaria, Belarus, Croatia, Romania, Russian Federation, Ukraine, Rest of Eastern Europe, Rest of Europe, Kazakhstan, Kyrgyzstan, Rest of Former Soviet Union, Armenia, Azerbaijan, Georgia, Bahrain, Iran, Israel, Kuwait, Oman, Qatar, Saudi Arabia, Turkey, United Arab Emirates, Rest of Western Asia, Egypt, Morocco, Tunisia, Rest of North Africa, Cameroon, Côte d'Ivoire, Ghana, Nigeria, Senegal, Rest of Western Africa, Rest of Central Africa, South Central Africa, Ethiopia, Kenya, Madagascar, Malawi, Mauritius, Mozambique, Tanzania, Uganda, Zambia, Zimbabwe, Rest of Eastern Africa, Botswana, Namibia, South Africa, Rest of South African Customs Union, Rest of the World

¹ Includes: Anguilla, Antigua and Barbuda, Aruba, Bahamas, Barbados, British Virgin Islands, Cayman Islands, Cuba, Dominica, Dominican Republic, Grenada, Haiti, Jamaica, Montserrat, Netherlands Antilles, Puerto Rico, Saint Kitts and Nevis, Saint Lucia, Saint Vincent and Grenadines, Trinidad and Tobago, Turks and Caicos Islands, Virgin Islands.

Table 3.A.1: Aggregation of countries and regions.

3.A.7 Sectoral and production factor aggregation

Aggregate variable	GTAP variable
Goods and sectors	
Manufacturing	Textiles, Wearing apparel, Leather products, Wood products, Motor vehicles, Other transport equipment, Water, Construction, Paper products, publishing, Chemical, rubber, plastic products, Minerals, Ferrous metals, Other metals, Metal products, Electronic equipment, Other machinery and equipment, Other manufactures
Services	Trade, Communication, Financial services, Insurance, Business services, Recreation, Dwellings, Public Administration, Defense, Education, Health
Transport	Water transport, Air transport, Other transport
Agriculture	Paddy rice, Wheat, Cereal grains, Vegetables, fruits, nuts, Oil seeds, Sugar cane, sugar beet, Plant-based fibers, Other crops, Bovine cattle, Other animal products, Raw milk, Wool, Forestry, Fishing, Bovine meat products, Other meat products, Vegetable oils and fats, Dairy products, Processed rice, Sugar, Other food products, Beverages and tobacco
Electricity	Electricity
Coal	Coal
Natural gas	Gas, Gas manufacture, distribution
Crude oil	Oil
Refined oil	Petroleum, coal products
Factors of production	
Resources	Land, Natural resources
Labour	Skilled labour, Unskilled labour
Capital	Capital

Table 3.A.2: Aggregation of sectors and production factors.

3.A.8 Main parameter values

Parameter	Description	Value
Elasticities of substitution for production activities		
$\sigma_{i,r}$	Intermediate composite Q and inputs B from other sectors	0.5
$\nu_{i,r}$	Labour L and energy E in intermediate good production	1.0
$\epsilon_{i,r}$	Energy type Z in the energy aggregate	0.5
$\eta_{i,r}$	Imports and domestic goods	$\in [1.9, 6.0]$
$\phi_{i,r}$	Import regions	$\in [3.8, 12]$
κ	Intermediate varieties	$\in \{1.0, 0.86\}$
Elasticities of substitution for private consumption		
$\xi_{i,r}^{agg}$	Transportation and other consumption goods	1.0
$\xi_{i,r}^{oth}$	Energy and non-energy consumption goods	0.25
$\xi_{i,r}^{ene}$	Energy varieties	0.4
$\xi_{i,r}^{non}$	Non-energy consumption goods	0.25
Other parameters		
$1/\theta$	Intertemporal elasticity of substitution	0.5
$\delta_{i,r,t}$	Baseline capital depreciation	0.07
\bar{r}	Baseline nominal interest rate	0.05
γ_K	Capital growth rate	0.02

Table 3.A.3: Default parameter values used in numerical simulations. Based on Refs. [110, 146, 147, 165].

3.B Additional results




Figure 3.B.1: Aggregate investment with recurring shocks. Solid lines represent the means of 500 Monte Carlo simulations relative to a no-shock steady-state economy. The shaded areas denote the 5th to 95th percentile ranges.




Figure 3.B.2: GDP levels with recurring shocks. Solid lines represent the means of 500 Monte Carlo simulations relative to a no-shock steady-state economy. The shaded areas denote the 5th to 95th percentile ranges.




Figure 3.B.3: Aggregate industry output with recurring shocks. Solid lines represent the means of 500 Monte Carlo simulations relative to a no-shock steady-state economy. The shaded areas denote the 5th to 95th percentile ranges.

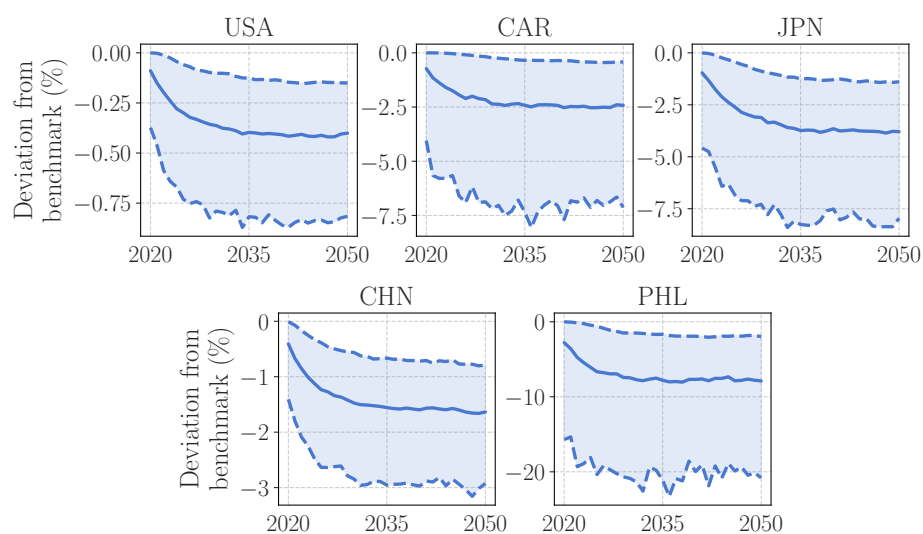


Figure 3.B.4: Capital intensity with recurring shocks. Solid lines represent the means of 500 Monte Carlo simulations relative to a no-shock steady-state economy. The shaded areas denote the 5th to 95th percentile ranges.




Figure 3.B.5: Aggregate capital stock with recurring shocks. Solid lines represent the means of 500 Monte Carlo simulations relative to a no-shock steady-state economy. The shaded areas denote the 5th to 95th percentile ranges.

3.C Disaster impact distributions




Figure 3.C.1: Empirical cumulative distribution functions of tropical cyclone damages from the historical sample (blue) versus the synthetic years (red) under constant climate conditions.

3.D Cyclone damages under climate change

The cyclone damage estimation with climate change follows similar steps as outlined in Section 3.2.1 for constant climate conditions. The difference is that under climate change, the cyclone damage distribution changes in ten-year intervals as climate change progresses under different assumed exogenous scenarios. More precisely, climate change affects the occurrence rate of cyclones (frequency) and the damage that each cyclone may cause (intensity).

We again start from the historical tracks and extend our sample with synthetic years. We use 2020 as a reference year. We generate a new pool of 5'000 synthetic years of damages for every decadal step from 2025 to 2095. At each step, we use the CLIMADA framework to calculate the cyclone intensity change. CLIMADA relies

on Ref. [155] as a reference for its climate change scenarios. It uses a dynamical downscaling of models from the Coupled Model Intercomparison Project Phase 5 (CMIP5) under the RCP4.5 emissions concentration scenario to project tropical cyclone activity for the end-of-century years 2081–2100. However, Ref. [155] only presents results for the RCP4.5 scenario. On top, CLIMADA interpolates other RCP scenarios from the RCP4.5 values according to the relative radiative forcing of each scenario.

RCP4.5

Ref. [155] considers the change in tropical cyclone activity under the RCP4.5 scenario. The authors find no significant change in tropical cyclones' frequency in the North Atlantic basin for late twenty-first century projections. However, they find a statistically significant decrease in tropical cyclone frequency in the Northwestern Pacific basin of 34.5%. We, therefore, keep the average yearly number of cyclones in the Caribbean islands and the US unchanged but specify a 34.5% drop in end-of-century cyclone frequency for Japan, China, and the Philippines. We interpolate the occurrence rate linearly at each decadal step from 2025 to 2095. The cyclone intensity in the North Atlantic basin increases by 4.5% in 2100 compared to the present day's climate and by 5-7% in the Northwestern Pacific basin [155].

Figure 3.D.1 shows the cumulative distribution of damages from the synthetic years without climate change, with 2020 as the reference climate, versus the synthetic yearly damages in 2100 under an RCP4.5 scenario. Overall, climate change gradually increases cyclone damage for the US and Caribbean islands. For the Northwestern Pacific regions in our study, cyclone damage slightly increases for the first decades with the higher intensity. However, damages fall over time below the 2020 reference year as the frequency effect dominates.




Figure 3.D.1: Empirical cumulative distribution functions of (synthetic) cyclone damages in 2020 (blue) versus their potential damages in 2100 under the RCP4.5 climate scenario (red).

RCP8.5

As Ref. [155] only considers the RCP4.5 scenario, we also use the results in Ref. [143] for constructing the high-emission RCP8.5 damages. The authors again downscale models from the CMIP5 ensemble to project end-of-century changes in tropical cyclone activity but also include global effects under the RCP8.5 scenario. The results suggest a frequency increase of 10 – 40% depending on the CMIP5 model. The changes are largest in the Northwestern Pacific basin but also present in the North Atlantic. We take a conservative calibration of a 5% and 10% increase in cyclone frequency for the North Atlantic and Northwestern Pacific basins, respectively. For the cyclone intensity values, we again use the linear interpolation based on Ref. [155] with radiative forcings as scaling terms.

Figure 3.D.2 shows the cumulative distribution of damages from the synthetic years without climate change, with 2020 as the reference climate, versus the synthetic damages in 2100 under the RCP8.5 scenario. The mean damages and stan-

dard deviations now increase for all regions, resulting in significantly higher projected damages in the economic model.




Figure 3.D.2: Empirical cumulative distribution functions of (synthetic) cyclone damages in 2020 (blue) versus their potential damages in 2100 under the RCP8.5 climate scenario (red).

Table 3.D.1 shows the corresponding changes in regional impact statistics.

Statistic (% change)	USA	CAR	JPN	CHN	PHL
RCP4.5					
Mean damage, $\bar{\delta}_{TC}$	30.16	24.03	-9.73	-12.66	-13.47
Std. of damages, σ_{TC}	25.94	18.72	5.86	-1.22	-2.18
RCP8.5					
Mean damage, $\bar{\delta}_{TC}$	112.92	87.95	166.83	151.50	126.55
Std. of damages, σ_{TC}	93.11	66.39	115.51	103.22	81.92

Table 3.D.1: Percentage change in tropical cyclone impacts in 2100 under different climate scenarios compared to estimates under constant climate conditions.

Chapter 4

Learning-based model predictive control for climate policy analysis

Abstract

This chapter explores the applicability of learning-based model predictive control for solving stochastic economic decision problems. The high computational cost of sequential decision-making under uncertainty and the long planning horizons encountered in many economic applications severely restrict the usability of conventional dynamic programming techniques in high-dimensional settings. Learning-based control techniques provide a promising alternative approach as they can often deliver high-quality approximate solutions with a significantly smaller computational budget. In addition, stochastic model predictive control enables an intuitive formulation of *cautious* behaviours that complement the conventional risk-averse decision rules. An integrated climate-economy model provides a challenging environment to evaluate the proposed method.

This chapter is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 870245.

4.1 Introduction

The climate-economy interaction is one of the most challenging dynamical systems to consider. Integrated assessment models (IAMs) that couple economic dynamics with a stylized Earth system model can approximate this relationship and provide valuable insight into climate policy design. However, despite pervasive uncertainties in both human and climate systems, most IAM implementations are entirely deterministic due to the high computational cost of stochastic optimization. Meanwhile, recent advances in learning-based control have introduced powerful approximate solution techniques for complex tasks from autonomous racing [32, 33, 166, 167] and robotic locomotion [168] to games [169–171].

This chapter sets up a general framework for solving stochastic economic decision problems using learning-based model predictive control (MPC). Learning-based control allows inferring unknown system properties from a self-generated database of interactions with the system. The learned components can enhance controller quality, for instance, in safety-critical applications [172]. On the other hand, MPC approximates infinite-horizon optimal policies by repeatedly solving a surrogate problem over a finite planning window, executing only the first action in the obtained control sequence, and replanning forward into a receding horizon after each step. Although more common in industrial applications, MPC also lends itself to dynamical decision problems in economics [173] and finance [174] and provides a natural way to deal with stochastic systems. Climate-economy IAMs provide another challenging environment for evaluating the proposed technique.

The MPC approach has several advantages. First, the computational cost is moderate compared to dynamic programming (DP) algorithms. DP requires iterating on an optimization problem at every point in a gridded state space, causing the numerical effort to grow exponentially fast with the state dimension.¹ MPC avoids the gridding procedure by truncating the problem in time and only approximating (or sometimes neglecting) the future periods beyond the planning horizon.

A model predictive controller can handle uncertain systems directly as the constant re-optimization enables reacting to model errors in a closed loop. Moreover,

¹Adaptive sparse grids [175] can attenuate the increase in computational complexity.

it is possible to encode uncertainty explicitly into the planning algorithm in tasks requiring cautious control. We model risk awareness via *chance-constrained programming*, reformulating the original system constraints in terms of an admissible constraint violation probability. Chance constraints provide an intuitive risk-performance tradeoff in many economic decision problems. For instance, in the climate context, they allow specifying a maximum probability for breaching a given temperature target. Compared to deterministic IAMs, this uncertainty treatment leads to significantly more aggressive emission cuts, as the planner must leave a sufficiently wide safety margin not to violate the temperature constraint.

The main MPC drawback is that the policies are necessarily myopic due to the truncated optimization horizon. That is in stark contrast to DP, where referencing an optimal value function allows characterizing exact solutions based on the *principle of optimality* [176]. Myopia is particularly problematic when applying MPC to IAMs. Climate dynamics and emission abatement costs can be strongly nonlinear, and greenhouse gases might linger in the atmosphere for centuries. Therefore, ignoring the full consequences of today’s (in)actions can lead to costly future readjustments. A longer prediction horizon reduces shortsightedness but at a computational cost. As a remedy, we consider learning an approximate terminal value function to provide intermediate targets for the controller. The value function propagates long-term information into the shortsighted algorithm, allowing MPC to approximate optimal decisions even in multi-century economic planning problems.

In summary, this chapter works towards bridging the literatures on dynamic economic decision problems, risk-aware MPC [33], and data-driven value function approximation [34, 177]. The resulting framework enables using local search methods like MPC over arbitrarily long planning horizons in repetitive environments with nonlinear dynamics. In addition, the chance-constrained formulation allows an intuitive treatment of risk when the system dynamics are uncertain. Although widely used in control engineering, to our knowledge, we are the first to apply the learning-based MPC approach in a climate-economy context.

The structure of this chapter is as follows. After a brief review of previous work, Section 2 introduces the necessary concepts in learning-based model predictive control. Section 3 presents the integrated climate-economy model, and Section 4 the numerical experiments. Section 5 concludes.

Related work

Integrated assessment models (IAMs) provide a tool for analyzing the costs and benefits of climate policy alternatives by coupling stylized Earth system dynamics to a simple representation of the global economy. One of the earliest and most influential IAMs is the Dynamic Integrated model of Climate and the Economy (DICE) [178–180]. A global decision-maker sets consumption levels and greenhouse gas mitigation rates in each period. More current consumption yields a higher immediate utility, but the resulting cut in savings suppresses future consumption possibilities. Similarly, mitigating emissions today is costly but reduces future climate impacts.

Decision making under uncertainty has long posed a computational challenge for solving IAMs. Ref. [181] uses a DICE-like model with a binomial tree structure to represent unknown future climate damages. Another approach is to carefully reduce the dimensionality of the original state space to afford a DP solution [31, 182–185]. Ref. [186] develops an approximate rolling-horizon DP algorithm under parametric uncertainty. Ref. [187] and Ref. [30] break the curse of dimensionality of DP with a supercomputer.

Model predictive control (MPC) is a well-established control technique for dynamical systems subject to input and state constraints. Traditionally, MPC considers stabilizing linear systems to some predetermined reference equilibrium, which yields a convex optimization problem with an efficient numerical implementation using standard gradient-based solvers [188]. Recently, with the increase in computational power, economic model predictive control (EMPC) has received growing interest for problems with general nonlinear objectives and dynamics [189]. Sampling-based optimization techniques extend the use of MPC to non-smooth problems [32]. Notably, gradient-free sampling techniques enable easy use of deep neural networks as function approximators in data-driven control.

Inferring unknown system properties from data can enhance control performance in repetitive tasks and safety-critical systems. In this regard, learning-based MPC shares close synergies with reinforcement learning (RL). Ref. [166] combines cautious MPC with Gaussian process regression to improve the prediction model in an autonomous racing application where the true system dynamics are unknown. Ref. [190] employs MPC as a safety filter on top of an RL algorithm, and Ref. [191]

uses deep quantile regression to learn uncertainty tubes around MPC trajectories. Ref. [34] and Ref. [192] integrate RL-style value function approximation to sampling-based MPC.

Several authors have applied MPC to sequential decision-making problems in economics. Ref. [173] provides a convergence analysis and implementation details for many standard classes of optimal growth problems. Ref. [193] applies MPC to green investment risks and Ref. [194] to emissions trading. Prior works on climate-economy IAMs include Refs. [195–199]. However, only a few of these studies include uncertainty, and none of them considers learning-based extensions.

Finally, chance-constrained MPC strikes a balance between performance and safety requirements and provides an appealing risk-aware formulation for many practical applications. Common examples are portfolio optimization [200–202], power system management [203], and inventory problems [204]. Still, MPC is just one among many techniques for approaching complex sequential decision problems. For instance, several authors have recently considered deep learning methods as an alternative to conventional dynamic programming [205–207]. However, it is the natural extensions to uncertainty quantification and probabilistic constraint satisfaction, in particular, that make MPC a valuable addition to the economic modelling toolbox.

4.2 Preliminaries

4.2.1 Notation

We denote time t variables in the true system with parenthesis, $\mathbf{x}(t)$, and the i^{th} -step variables along the MPC prediction horizon with subscripts, $\mathbf{x}_{i|t}$, conditioned on the starting time instance. The Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ is $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$. The Jacobian matrix of a vector-valued function f evaluated at location \mathbf{x} is $\nabla f(\mathbf{x})$. The Pontryagin set difference is $\mathcal{A} \ominus \mathcal{B} = \{a \mid a + b \in \mathcal{A} \forall b \in \mathcal{B}\}$. The i^{th} element in a vector \mathbf{x} is $[\mathbf{x}]_i$. For a matrix M , the i^{th} row is $[M]_i$ and the element (i, j) is $[M]_{i,j}$. A diagonal matrix with main diagonal entries \mathbf{x} is $\text{diag}(\mathbf{x})$.

4.2.2 Model predictive control

This section describes the use of MPC to approximate a general discrete-time infinite-horizon stochastic optimal control problem commonly encountered in economics. The focus here is on MPC variants based on deterministic optimal control, leaving out, for instance, scenario optimization [208] and tree-based methods [209].

The state of the system evolves according to a nonlinear, continuous dynamics model $\mathbf{x}(t+1) = f(\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t))$, where $\mathbf{x}(t) \in \mathcal{X} \subset \mathbb{R}^{n_x}$ is the state vector, $\mathbf{u}(t) \in \mathcal{U} \subset \mathbb{R}^{n_u}$ the control input, and \mathcal{X} and \mathcal{U} the respective constraint sets. The system is subject to a sequence of additive stochastic disturbances $\mathbf{w}(t) \in \mathbb{R}^{n_w}$, which we assume to be independent and identically distributed (i.i.d.). The control objective is to maximize the expected sum over a real-valued reward function $r(\mathbf{x}, \mathbf{u})$, discounted with $\beta \in (0, 1]$. Ideally, for a given initial state \mathbf{x}_0 and a noise distribution Q^w , we would like to solve the following problem:

$$\begin{aligned} & \max_{\pi} \mathbb{E}_w \left(\sum_{t=0}^{\infty} \beta^t r(\mathbf{x}(t), \mathbf{u}(t)) \right) \\ \text{s.t. } & \mathbf{x}(t+1) = f(\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t)) \\ & \mathbf{u}(t) = \pi(\mathbf{x}(t)) \\ & \mathbf{x}(0) = \mathbf{x}_0 \\ & \mathbf{w}(t) \sim Q^w \\ & \mathbf{x}(t) \in \mathcal{X}, \mathbf{u}(t) \in \mathcal{U} \forall t. \end{aligned} \tag{4.1}$$

The search space in Eq. (4.1) is over general policy rules $\pi : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_u}$, mappings from any going state to the optimal control action. As such, the problem is intractable and requires approximation with an easier surrogate problem.

The main idea behind MPC is simple. Instead of tackling the entire infinite horizon at once, the controller repeatedly solves the problem for a finite $N \in \mathbb{N}$ steps into the future. It takes the first element in the obtained control sequence as an immediate control input, implements it in the system, and observes the resulting next state. Finally, it updates the initial conditions and re-optimizes for another forecast window of N steps forward in time. Starting at any time instance k , the

reduced search space is therefore over control trajectories $\bar{\mathbf{u}}_k = [\mathbf{u}_{0|k}, \dots, \mathbf{u}_{N-1|k}]$ and the corresponding states $\bar{\mathbf{x}}_k = [\mathbf{x}_{0|k}, \dots, \mathbf{x}_{N|k}]$. For an optimal control sequence $\bar{\mathbf{u}}_k^*$, the MPC policy rule then becomes $\mathbf{u}(k) = \bar{\pi}(\mathbf{x}(k)) = [\bar{\mathbf{u}}_k^*]_0$. The full $\bar{\mathbf{u}}_k$ and $\bar{\mathbf{x}}_k$ are referred to as time- k *open-loop* predictions, and the sequences constructed by following the MPC policy $\bar{\pi}$ as *closed-loop* trajectories.

Truncating the optimization horizon, however, leads to approximation errors. MPC often considers a terminal value function $V(\mathbf{x}_{N|k})$ or a terminal constraint set $\mathbf{x}_{N|k} \in \mathcal{X}_{N|k}$ at the end of the prediction window as an intermediate target to account for the missing tail of steps. These terminal ingredients guide decision making towards long-term optimal actions, although their design can be challenging for general nonlinear systems.

The second MPC approximation step is to replace the original state-space model with a deterministic prediction function $\mathbf{x}_{i+1|k} = \bar{f}(\mathbf{x}_{i|k}, \mathbf{u}_{i|k})$. For instance, \bar{f} can be an approximation of the underlying system based on first principles. Decision problems in economics typically assume access to a perfect simulator, in which case the prediction function corresponds to the nominal noise-free transition dynamics.

A fundamental design choice in model predictive control is how to deal with uncertainty. There are three main paradigms. *Nominal* MPC ignores all uncertainties during planning, but repeatedly re-solving the problem still enables reacting to disturbances in a feedback loop. On the other hand, *robust* MPC assumes that all uncertainty realizations come from a known bounded set and only deploys control inputs that guarantee constraint satisfaction under any possible outcome. Although the robust approach can deliver strong theoretical properties, it can lead to conservative behaviour by always preparing for the worst-case realization and not exploiting the distributional information of the random events. Finally, *stochastic* MPC specifies a noise distribution, treating the (possibly unbounded) sequence $\{\mathbf{w}(t)\}_{t=0}^{\infty}$ as a stochastic process with independent and identically distributed terms. The system constraints only need to hold in probability, which reduces conservatism.

We can define chance constraints for the state and control trajectories as:

$$\Pr(\bar{\mathbf{x}}_k \in \mathcal{X}) \geq p_x$$

$$\Pr(\bar{\mathbf{u}}_k \in \mathcal{U}) \geq p_u,$$

where $p = 1$ corresponds to the robust approach with hard constraints and $0 < p < 1$ to the stochastic setting. However, solving the MPC problem using standard numerical optimizers still requires converting the chance constraints into a deterministic format. That involves two steps. First, *uncertainty propagation* models the accumulation of uncertainty over time and yields a distribution of possible states along every step of the prediction horizon. Then, using these uncertainty estimates, *constraint tightening* defines new constraint sets $\bar{\mathcal{X}} \subseteq \mathcal{X}$ and $\bar{\mathcal{U}} \subseteq \mathcal{U}$. Keeping the nominal open-loop trajectories within the tightened bounds corresponds to satisfying the original system constraints with the desired probability. For more details on implementing the chance constraints, see Section 4.4.2.

The final simplifying step concerns the expectation operator. In general, there is no reason to assume that MPC will maximize the expected sum in Eq. (4.1). For many types of functional forms, quadratic in particular, it is possible to decompose and evaluate the expectation analytically based on mean and variance information. However, for the general case considered here, we simply ignore the expectation and convert the original objective into a sequence of deterministic finite-horizon problems, invoking an approximate form of certainty equivalence [173, 210]. First, the noise process in many economic applications is simple, typically assuming additive zero-mean i.i.d. Gaussian disturbances with relatively small magnitudes. Therefore, locally, dropping the expectation yields a reasonable surrogate objective while maintaining the useful MPC properties. Second, a well-tuned value function, approximated from a sufficient amount of prior experience, can guide the open-loop MPC trajectories towards terminal states with high expected reward-to-go, resembling the standard Bellman equation [176] in a stochastic setting.

At every time instance k , we can then approximate the original problem in Eq. (4.1) by considering the following surrogate system:

$$\begin{aligned}
& \max_{\bar{\mathbf{u}}_k} \beta^N V(\mathbf{x}_{N|k}) + \sum_{i=0}^{N-1} \beta^i r(\mathbf{x}_{i|k}, \mathbf{u}_{i|k}) \\
& \text{s.t. } \mathbf{x}_{i+1} = \bar{f}(\mathbf{x}_{i|k}, \mathbf{u}_{i|k}) \\
& \quad \mathbf{x}_{0|k} = \mathbf{x}(k) \\
& \quad \bar{\mathbf{u}}_k = [\mathbf{u}_{0|k}, \dots, \mathbf{u}_{N-1|k}] \\
& \quad \bar{\mathbf{x}}_k = [\mathbf{x}_{0|k}, \dots, \mathbf{x}_{N|k}] \\
& \quad \mathbf{x}_{N|k} \in \bar{\mathcal{X}}_N, \bar{\mathbf{x}}_k \in \bar{\mathcal{X}}, \bar{\mathbf{u}}_k \in \bar{\mathcal{U}}.
\end{aligned} \tag{4.2}$$

Repeatedly solving Eq. (4.2) generates a sequence of control inputs, emulating an infinite horizon by shifting the prediction window forward after each iteration, always maintaining a distance of N steps to the termination. Under certain conditions in a deterministic setting, a formal convergence analysis between the original infinite-horizon solution and the MPC approximation is also possible [173].

4.2.3 Sampling-based trajectory optimization

Gradient-based numerical solvers are appealing for MPC due to their accuracy and stability. However, their use can be restrictive in some data-driven applications. Gradient-based methods require the optimization landscape to be sufficiently smooth, which can be difficult to guarantee under high-dimensional function approximators like deep neural networks. On the other hand, methods such as linear basis function regression and Gaussian processes can deliver smooth function approximations but have limited expressive power in higher dimensions and high computational costs, respectively. Sampling-based optimization can provide more flexibility, placing no shape requirements on the system dynamics or objective functions.

Given the prediction model $\mathbf{x}_{t+1} = \bar{f}(\mathbf{x}_t, \mathbf{u}_t)$ and an instantaneous reward function $r(\mathbf{x}_t, \mathbf{u}_t)$, the simplest approach to sampling-based trajectory optimization is to generate Z possible control trajectories over a prediction horizon of N steps purely at random and pick the sample with highest cumulative reward as an approximate solution. Although sufficient for some applications [211], such random shooting rarely generates meaningful enough patterns for high-precision control.

To add more structure to the algorithm, we follow the model predictive path integral method (MPPI) for sampling-based MPC [32, 212]. The approach brings two improvements compared to simple random shooting. First, instead of choosing a single best trajectory as an approximate solution, the method takes a weighted average over all generated samples according to their relative cumulative performance. Second, instead of sampling action candidates as unstructured noise, an additional filtering step is applied to produce smooth action trajectories for more guided search directions. Figure 4.2.1 illustrates the process in a simple one-dimensional economic growth model.

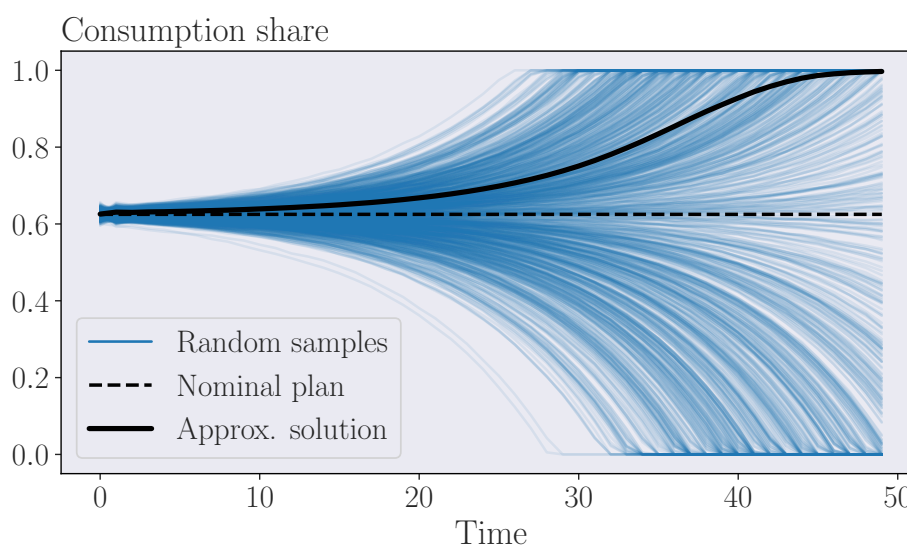


Figure 4.2.1: Illustration of sampling-based trajectory optimization in a one-dimensional economic growth model. The task is to find an optimal output consumption share over a period of fifty steps. We initialize the algorithm with a constant nominal trajectory, denoted with the black dashed line. Blue lines denote 1000 randomly generated smooth perturbations. The solid black line marks the obtained approximate solution, a weighted average over all samples. The optimal open-loop sequence ends with the consumption share reaching 100% as the agent wants to leave no capital behind at the end of the optimization period.

In practical terms, we initialize the algorithm with a nominal control sequence $\mathbf{u} = \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$ for a time horizon of N steps. Then, we sample Z zero-mean Gaussian perturbations $\boldsymbol{\varepsilon}^z = \{\boldsymbol{\varepsilon}_0^z, \boldsymbol{\varepsilon}_1^z, \dots, \boldsymbol{\varepsilon}_{N-1}^z\} \forall z \in \{0, \dots, Z-1\}$, and apply the

noisy control sequences $\mathbf{u} + \boldsymbol{\varepsilon}^z$ by recursively using the prediction model \bar{f} from a known initial state. Denoting the resulting cumulative rewards of all sample trajectories with $R = \{R_0, R_1, \dots, R_{Z-1}\}$, we use the following weights to average over the Z candidates:

$$w_z = \frac{\exp\{\lambda(R_z - \max[R])\}}{\sum_{j=0}^{Z-1} \exp\{\lambda(R_j - \max[R])\}}. \quad (4.3)$$

The term λ is a tunable temperature parameter. Setting $\lambda = 0$ gives equal weight to all trajectories. With $\lambda > 0$, the best sample gets the highest weight, while the weights for the remaining samples decay faster with worse relative performance. Finally, following the usual MPC formulation, we choose the first element of the weighted control sequence as the next immediate control input and update the old nominal sequence \mathbf{u} based on the remainder of the newly weighted sample. The process of generating and evaluating individual trajectories is easily parallelizable, which allows using a large sample size to approximate the optimal solution.

More formally, under the information-theoretic MPC approach outlined above, the original trajectory optimization problem becomes one of probability matching [32, 213]. The idea is first to derive a performance upper bound based on the concept of free energy. The second step is to establish the equivalence between finding an optimal control sequence by maximizing the original objective and reaching the upper performance bound by sampling controls from an abstract optimal control distribution \mathcal{Q}^* . However, we cannot sample from the abstract optimal distribution directly. Instead, we can minimize the Kullback–Leibler divergence between \mathcal{Q}^* and an open-loop control distribution \mathcal{Q} by using the analytic expressions for their density functions, pushing the proposal distribution \mathcal{Q} as close as possible to the optimal, \mathcal{Q}^* . Still, this only expresses the optimal controls in terms of an expectation over \mathcal{Q}^* , resulting in an intractable integral. However, we can use importance sampling with weights w_z from Eq. (4.3) to approximate the expectation with respect to \mathcal{Q}^* using Monte Carlo draws from the proposal distribution \mathcal{Q} .

The resulting control law allows selecting actions based on a weighted average of sampled paths and is applicable to systems with arbitrary objective functions and nonlinear dynamics. However, there are some caveats. First, being a local search method, the algorithm relies on a suitable proposal distribution for the importance

sampling and a good initial guess to warm-start the iteration. Moreover, matching the expectations between the optimal and proposal distributions suffices for finding the optimal controls only under control-affine system dynamics. For more general systems, the optimal control distribution might, for instance, be multi-modal and difficult to describe using mean information alone. However, empirical experiments suggest that the method also works well in more general settings [32]. We validate the performance of the sampling-based algorithm against a standard gradient-based controller in Appendix 4.A. Additional practical improvements for sampling-based MPC algorithms include, among others, smoothing the obtained control sequence [213], adaptive importance sampling [214], and gradient descent steps to refine the sampled rollouts [215].

4.2.4 Uncertainty-aware neural network ensembles

Data-driven controllers can gradually improve performance by collecting more experience through interactions with the system. Often it is also helpful to quantify the uncertainty in the learned components. For instance, uncertainty can guide the agent to explore unknown regions for maximal information gain or enforce cautious behaviour by avoiding the uncertain areas altogether. The standard taxonomy classifies uncertainty as either *epistemic* for the lack of data or *aleatoric* for the inherent, irreducible system noise.

Neural networks are powerful and scalable nonlinear function approximators. However, traditional feed-forward neural networks do not have a natural representation of predictive uncertainty. Common techniques for considering uncertainty include Bayesian networks [216] and training an ensemble of multiple networks in parallel [217]. Bayesian networks learn a distribution over the model parameters instead of only point estimates, whereas using an ensemble model allows exploiting the distribution over individual network outputs to quantify uncertainty. We use the ensembling approach due to its simple implementation and easy integration into sequential decision-making problems.

Ensembling multiple neural networks into a single model is common across different fields of machine learning [34, 218, 219]. The idea is to inject variability into the networks by either randomly initializing the model parameters, bootstrapping

the training data, or a combination thereof. After training, the ensemble members should agree in regions with enough training examples, but their disagreement elsewhere can be used to proxy predictive (epistemic) uncertainty. Figure 4.2.2 illustrates the process for a one-dimensional toy regression task. The value of ensembling also goes beyond uncertainty quantification, as averaging over the ensemble members can sometimes provide more accurate predictions compared to individual model outputs [218].

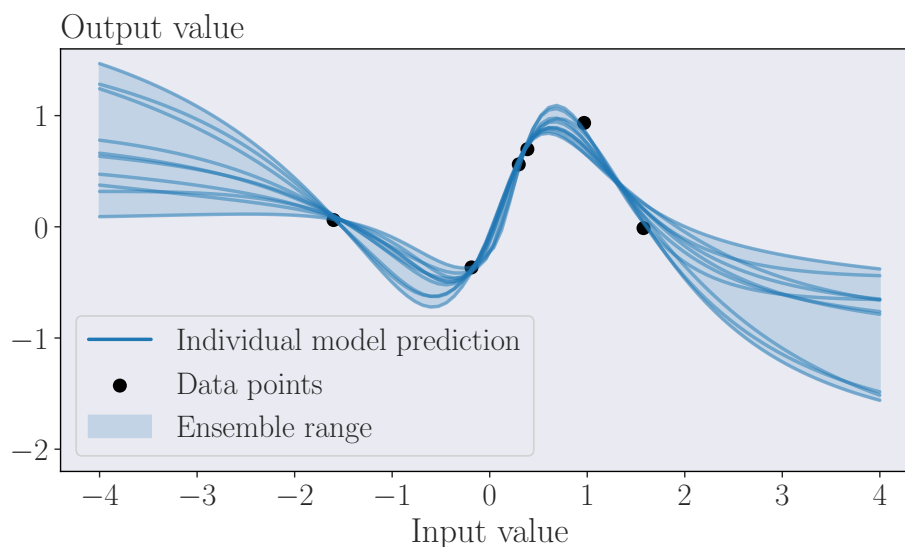


Figure 4.2.2: Predictions of a neural network ensemble in a toy regression problem. The ensemble consists of ten randomly initialized networks. Solid blue lines denote the predictions for individual models. After training, the models agree near the training points but scatter quickly in regions with no available data. The disagreement within the ensemble, measured, for instance, as the empirical variance across all predictions, provides a convenient estimate for epistemic uncertainty.

4.2.5 Value function approximation

The dynamical system considered in this chapter conforms with the standard formalism of an infinite-horizon discounted Markov Decision Processes [220]. In this setting, given a decision rule π , the corresponding value function for any going state

\mathbf{x} is the expected sum of discounted future rewards, $V^\pi(\mathbf{x}) = \mathbb{E}[\sum_{t=0}^{\infty} \beta^t r(\mathbf{x}_t, \pi(\mathbf{x}_t)) \mid \mathbf{x}_0 = \mathbf{x}, \pi]$. The expectation is over state transitions and decision rules, both of which might contain randomness. Denoting the next state from \mathbf{x} with \mathbf{x}' , we can define the usual Bellman optimality operator as:

$$\mathcal{T}V(\mathbf{x}) = \max_{\mathbf{u}} \mathbb{E}[r(\mathbf{x}, \mathbf{u}) + \beta V(\mathbf{x}')]. \quad (4.4)$$

The optimal value function V^* is the fixed point of this operator, which implies $\mathcal{T}V^*(\mathbf{x}) = V^*(\mathbf{x})$. To learn a parameterized approximate value function \hat{V}_θ , we can generate target values $\mathcal{T}\hat{V}_{\theta_i}(\mathbf{x})$ and iteratively update the parameters θ using:

$$\theta_{i+1} = \arg \min_{\theta} \mathbb{E}_{\mathbf{x} \sim \mathcal{Q}_{\mathbf{x}}} [(\hat{V}_\theta(\mathbf{x}) - \mathcal{T}\hat{V}_{\theta_i}(\mathbf{x}))^2]. \quad (4.5)$$

While training the weights, we sample possible states from a distribution $\mathcal{Q}_{\mathbf{x}}$. For more details on generating the regression targets in the context of learning-based MPC, see Section 4.4.1.

4.3 Environment

This section describes the dynamical system used in our experiments. A simple growth module specifies the drivers of long-run economic development and the links between economic activity and greenhouse gas emissions. In addition, two climate modules form a stylized, globally averaged Earth system model. First, a carbon cycle module determines how new emissions add up to atmospheric CO₂ concentrations and how these concentrations evolve with natural sinks slowly absorbing the excess carbon. Second, a temperature response module translates carbon concentrations into radiative forcings and temperature anomalies relative to a preindustrial baseline. Full details on calibrating the system are available in Appendix 4.B.

4.3.1 Model economy

The economic dynamics follow the stochastic general equilibrium model of Ref. [30]. The model differs from the standard DICE implementation of Ref. [221] mainly by splitting the time horizon into annual steps instead of five-year increments and by

including a two-dimensional stochastic process for persistent business cycle fluctuations. The time horizon for our implementation is 500 years.²

The primary economic state variable is the capital stock K , measured in monetary units. World population L and the base economic productivity A grow according to deterministic trends. The gross economic output before climate change impacts is a combination of technology, capital, and labour:

$$\Gamma(\zeta, A, K, L) = \exp(\zeta)AK^\alpha L^{1-\alpha}, \quad (4.6)$$

where the constant $0 < \alpha < 1$ determines the factor shares between labour and capital, and ζ is the stochastic productivity component. A random variable χ models the persistence of the productivity shocks. The random technology terms evolve as:

$$\begin{bmatrix} \zeta(t+1) \\ \chi(t+1) \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \zeta(t) \\ \chi(t) \end{bmatrix} + \nu(t). \quad (4.7)$$

The constant $0 < \rho < 1$ determines the shock persistence and $\nu(t) \sim \mathcal{N}(0, \Sigma^\nu)$ with a diagonal covariance matrix Σ^ν the fluctuation magnitude. The global average atmospheric temperature anomaly T_{AT} relative to a preindustrial equilibrium affects economic output through the following damage function:

$$\Psi(T_{AT}) = 1 + \left(\frac{T_{AT}}{\psi_1}\right)^{\psi_2} + \left(\frac{T_{AT}}{\psi_3}\right)^{\psi_4}. \quad (4.8)$$

The annual output is then $Y(t) = \Gamma(\zeta(t), A(t), K(t), L(t))/\Psi(T_{AT}(t))$ for non-negative ψ_i .³ Total annual CO₂ emissions $E(t)$ are the sum of exogenous land emissions $E_{\text{land}}(t)$ and industrial process emissions. The industrial emissions are proportional to the gross output, $E_{\text{ind}}(t) = \sigma(t)(1 - \mu(t))\Gamma(\zeta(t), A(t), K(t), L(t))$. The carbon intensity term $\sigma(t) > 0$ falls exogenously over time reflecting improvements in energy efficiency, and the planner sets an emissions mitigation rate $0 \leq \mu(t) \leq \mu_{\text{max}}$ in each period. The emissions abatement cost is:

²We do not consider climate tipping points or recursive Epstein-Zin preferences, which have a central role in Ref. [30].

³This damage formulation accommodates both the quadratic damage function from Ref. [179], the calibration from Ref. [222] with higher additional damages from a catastrophic temperature increase, as well as the more recent meta-analysis on damage estimates in Ref. [223].

$$\Omega(t) = \omega_1(t)\mu(t)^{\omega_2}Y(t), \quad (4.9)$$

for an exogenously decreasing $\omega_1(t) > 0$ and a constant cost term $\omega_2 > 1$. The aggregate resource constraint of the economy is:

$$Y(t) = C(t) + \Omega(t) + I(t), \quad (4.10)$$

where $C(t)$ and $I(t)$ denote consumption and investment levels, respectively. Given a constant depreciation rate $0 < \delta < 1$, the capital stock law of motion is:

$$K(t+1) = (1 - \delta)K(t) + I(t). \quad (4.11)$$

The agent's instantaneous utility u depends on per-capita consumption with:

$$u(C, L) = \frac{(C/L)^{1-1/\xi}}{1 - 1/\xi} L. \quad (4.12)$$

The constant ξ determines the intertemporal elasticity of substitution, which, with the model horizon spanning multiple centuries, captures tendencies for intergenerational consumption smoothing. The agent's goal is to maximize societal welfare W , defined as the expected sum of utility over time. The expectation is with respect to all economic and climate uncertainties. With a discounting term β close to 1, the welfare expression becomes:

$$W = \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t u(C(t), L(t)) \right]. \quad (4.13)$$

An essential IAM output quantity is the *social cost of carbon* (SCC), describing the marginal societal impact of an additional ton of carbon released into the atmosphere. A standard definition that gives the SCC in suitable monetary units is as a ratio between the marginal welfare loss from an additional unit of emissions and the marginal welfare gain of an additional unit of consumption [221]:

$$SCC_t = -\frac{\partial W}{\partial E(t)} \bigg/ \frac{\partial W}{\partial C(t)}. \quad (4.14)$$

A quantity closely related to the SCC is the carbon tax τ_E . Since emissions are the

only market distortion in our simplified world economy, the externality of environmental damages can be fully internalized with a Pigouvian carbon tax. Notably, whenever the input constraint $\mu(t) \leq \mu^{\max}$ is not binding, an optimal carbon tax per one tonne of carbon exactly equals the SCC with:

$$\tau_E(t) = \frac{\omega_1(t)\omega_2\mu(t)^{\omega_2-1}}{0.001\sigma(t)}. \quad (4.15)$$

When the abatement constraint is binding, the SCC can far exceed the optimal carbon tax level [30].⁴

4.3.2 Carbon cycle

The carbon cycle dynamics follow the Finite Amplitude Impulse Response (FAIR) model [224–227]. FAIR provides a transparent and lightweight emulator of the emissions response in full-scale Earth system models. Still, the model is detailed enough to capture the nonlinearities from climate feedback loops. The model considers all variables as global averages, which makes it an excellent submodule for integrated assessment [228–231].⁵

The carbon state variable consists of four reservoirs, G_i , with different rates of carbon uptake and time decay. The reservoirs approximately correspond to the concentration anomalies in the geological, deep ocean, biospheric, and ocean mixed layer processes [225]. The carbon content in each of the pools evolves according to:

$$G_i(t+1) = a_i\tau_i\gamma(t)E(t)\left(1 - \exp\left[\frac{-1}{\tau_i\gamma(t)}\right]\right) + \exp\left[\frac{-1}{\tau_i\gamma(t)}\right]G_i(t), \quad (4.16)$$

which is a discretized representation of the original FAIR dynamics. All reservoirs are empty at the preindustrial equilibrium. The parameter a_i , with $\sum_{i=1}^4 a_i = 1$,

⁴We can derive the carbon tax level $\tau_E(t)$ as follows. Optimally behaving producers will curb emissions up to the point where marginal abatement cost reaches the carbon price, which in the Pigouvian setting should be exactly equal to the social cost of emissions. From Eq. (4.9), the marginal abatement cost per one unit of output is $\omega_1\omega_2\mu_t^{\omega_2-1}$. The corresponding output quantity that produces exactly one unit of emissions is $\tilde{Y} = \frac{1}{\sigma}$, which directly gives the tax level in Eq. (4.15).

⁵We only consider carbon dioxide emissions. However, the FAIR model can capture the impact of a range of greenhouse gases with a unified functional representation [227].

captures the share of annual aggregate emissions $E(t)$ entering each of the reservoirs, and τ_i determines the corresponding lifetimes in the cycle. The carbon dynamics are nonlinear, capturing the effect of weakening carbon sinks (positive climate feedbacks) with higher cumulative emissions and temperature. The nonlinearity term $\gamma(t)$ models these effects and acts as a scaler on the time decay rates, with a higher value leading to a larger share of new emissions remaining in the atmosphere. The total carbon concentration at time t is $M(t) = M_0 + \sum_{i=1}^4 G_i(t)$, with M_0 denoting the preindustrial concentration. Writing $M_u(t)$ as all historical emissions already removed from the carbon cycle and $T_{AT}(t)$ for the current mean temperature, we can calibrate the nonlinearity with the following system of equations:

$$M_u(t) = \sum_{s=t_0}^t E(s) - \sum_i G_i(t) \quad (4.17)$$

$$\gamma(t) = g_0 \exp\left(\frac{r_0 + r_M M_u(t) + r_T T_{AT}(t)}{g_1}\right) \quad (4.18)$$

$$g_1 = \sum_{i=1}^4 a_i \tau_i [1 - (1 + 100/\tau_i) \exp(-100/\tau_i)] \quad (4.19)$$

$$g_0 = \exp\left(-\frac{\sum_{i=1}^4 a_i \tau_i [1 - \exp(-100/\tau_i)]}{g_1}\right) \quad (4.20)$$

The exponentiated numerator in Eq. (4.18) describes variations in the 100-year integrated impulse response function (iIRF), a measure for the average fraction of a carbon pulse remaining in the atmosphere over the given period [226]. Parameters r_0, r_M, r_T denote the preindustrial iIRF, and the increase in iIRF with cumulative carbon uptake and warming, respectively.

4.3.3 Temperature dynamics

The temperature dynamics follow the two-layer energy balance model in Ref. [232]. Higher carbon concentration $M(t)$ feeds into a higher level of radiative forcing with a logarithmic relationship:

$$\mathcal{F}(t) = \frac{\mathcal{F}_{2\text{XCO}_2}}{\ln(2)} \ln\left(\frac{M(t)}{M_0}\right) + \mathcal{F}_{\text{EX}}(t), \quad (4.21)$$

where $\mathcal{F}_{2\text{XCO}_2}$ is the forcing corresponding to the doubling of CO_2 and \mathcal{F}_{EX} stands for exogenous forcings from non- CO_2 sources. The temperature anomaly vector $\mathbf{T}(t) = [T_{\text{AT}}(t), T_{\text{OC}}(t)]$ for the mean surface air and deep ocean temperatures evolves as a linear function in terms of radiative forcing:

$$\mathbf{T}(t+1) = \begin{bmatrix} 1 - \Xi^{-1}(\eta + \Lambda) & \Xi^{-1}\eta \\ \eta\Xi_0^{-1} & 1 - \eta\Xi_0^{-1} \end{bmatrix} \mathbf{T}(t) + [\Xi\mathcal{F}(t), 0]^\top + v(t). \quad (4.22)$$

The parameters Ξ and Ξ_0 are per-area heat capacity coefficients for upper and deep oceans, respectively, and η is the between-layer heat exchange term. The parameter Λ controls the CO_2 perturbation feedback. In addition to the calibration in Ref. [232], we add a noise term $v(t) \sim \mathcal{N}(0, \Sigma^v)$ to the atmospheric temperature process to model the variability in historical temperature observations [198, 233].

In total, the state variable $\mathbf{x} = [K; \zeta; \chi; \mathbf{G}; \mathbf{T}; E_{\text{cum}}]^\top$ has ten dimensions: capital stock, two stochastic productivity states, a four-dimensional carbon cycle, two-dimensional temperature system, and cumulative emissions for easily keeping track of Eq. (4.17). For the numerical implementation, we also append time to the state vector as it provides valuable information for inferring the terminal value function (see details below). The control vector $\mathbf{u} = [c; \mu]^\top$ consists of the output consumption share $c(t) \in (0, 1]$ and the emissions control rate $\mu(t) \in [0, \mu_{\text{max}}]$. Defining the admissible controls as state-independent variables substantially simplifies the implementation. For instance, the optimal consumption *levels* can vary within the model horizon by several orders of magnitude due to centuries of exponential growth. The implicit assumption when limiting the consumption share to unity is that all investments are irreversible, as the consumption level can never exceed the current output.

The cost-benefit IAM described above has many desirable properties. First, with several interacting subsystems, maintaining a modular structure enables easily changing individual model components. When solving the model using MPC, the underlying climate system can even be a black-box simulator with the mis-

match between the prediction model and true dynamics learned from experience. The benchmark model used here also has room for new components without the computational burden becoming restrictively high. For instance, additional state and control variables might capture path dependencies in the energy system or the potential role of geoengineering technologies in the climate policy mix. Finally, the current setup is highly transparent. The model builds on a stylized dynamic general equilibrium representation of the economy and a lightweight climate subsystem, making it easy to analyze and extend to further applications.

4.4 Experiments

This section presents numerical experiments to evaluate the proposed solution technique in the climate-economy setting. The analysis comes in two parts. First, we address the issue of MPC myopia by learning an approximate terminal value function from simulated data to guide the shortsighted nominal controller. We run the controller in a deterministic benchmark environment where a long-term optimal reference path is easy to compute. The goal is to verify that the learning-based MPC method can approximate this benchmark policy and is thus comparable with other existing IAM solution approaches.

In the second experiment, we document the uncertainty quantification and constraint tightening steps that enable constructing risk-aware climate policy paths. We simulate the stochastic IAM under a climate policy target of limiting the global mean atmospheric temperature increase to two degrees Celcius and interpret *risk* as the probability of exceeding this temperature threshold. We then show how the optimal carbon price trajectories evolve under different risk appetite levels. Importantly, the second experiment does not consider value function learning. This ensures that the resulting carbon price paths for this illustrative example only vary because of the differences in risk treatment and not due to differences in learning performance. We can then analyze the *relative* carbon prices under different risk levels, assuming that value learning would similarly affect all results. However, switching on the learning-based components is straightforward also in the stochastic setting.

4.4.1 Value function learning

Addressing shortsightedness. Without terminal value information, MPC approximations of the long-term optimal trajectories can deteriorate rapidly if the prediction horizon is not sufficiently long for the task at hand. The application to climate-economy models provides a clear example. Figure 4.4.1 plots the relative errors of the MPC solution under different prediction horizon lengths in a deterministic setting where an optimal benchmark trajectory is easy to compute. The consumption shares are reasonably accurate for as short prediction horizons as 25 steps, which corresponds to only 5% of the entire model horizon of 500 years. In contrast, for the emissions control rate, bringing the relative error down to less than one percent requires an MPC prediction window of approximately 350 steps. This can be problematic as the computational time increases rapidly with a longer MPC horizon.⁶




Figure 4.4.1: Relative errors and solution times for MPC under various prediction horizon lengths compared to a non-myopic deterministic benchmark. The shortest considered horizon is 25 steps, and the longest one covers the full episode of 500 steps. **Left:** Errors in the consumption share are small already for short prediction horizons, never exceeding one percent. **Middle:** Errors for the emissions control rate are substantial, starting from up to 75% relative to the benchmark. As a result, the MPC solution consistently lags behind the long-run optimal mitigation policy. **Right:** Solution times grow rapidly with the prediction horizon.

⁶For these error illustrations, we use an MPC implementation on a gradient-based numerical solver instead of a sampling-based technique. It makes the analysis more robust against hyperparameter selection and removes all randomness from the comparison. Solution times on a ThinkPad X1 Carbon Gen 6 16GB laptop using the open source IPOPT optimizer [58].

Myopia in terms of emissions mitigation is intuitive. Due to the nonlinearity in carbon dynamics, mitigation costs, and climate damages, the controller does not see the total cost of excessive carbon emissions beyond the planning horizon. As a result, the emission cuts under MPC consistently lag behind the long-run optimal level.

The gap between the infinite-horizon optimal policy and the myopic MPC solution has an appealing economic interpretation of decision making under information constraints or rational inattention [173, 234]. The interpretation is especially fitting in the climate context, describing the mismatch between the observed global climate action and the level of ambition needed to avoid catastrophic levels of warming [7]. Still, being able to replicate the optimal long-term outcomes is crucial for comparing the MPC performance with other existing solution methods. To improve the accuracy of our approach while maintaining a reasonably short prediction horizon, we construct an approximate terminal value function for replacing V in Eq. (4.2).

Adding a terminal value function or a terminal constraint set is a standard procedure in the MPC toolkit to account for the remaining time steps beyond the prediction horizon [188]. However, designing these terminal components can be difficult for general nonlinear systems. For iterative tasks, a common approach is to learn an approximate terminal value function offline by evaluating the discounted reward-to-go empirically from a given state. The approximate value function then guides the open-loop MPC trajectories online to intermediate goals with a high expected reward [177, 210].

More specifically, we follow the approximation scheme outlined in Ref. [34]. The authors use an ensemble of M neural networks to learn a parameterized value function \hat{V}_θ . The ensemble consists of parameter sets $\theta_1, \dots, \theta_M$. After each interaction with the environment, we store the observed (state, action, next state, reward) tuple into a replay buffer \mathcal{D} , overriding old entries after reaching the maximum buffer capacity. Every Z steps, we update the value ensemble by randomly querying the buffer for starting states \mathbf{x} and constructing regression targets by evaluating the N -step return:

$$y(\mathbf{x}) = \beta^N \hat{V}_\theta(\mathbf{x}_N) + \sum_{t=0}^{N-1} \beta^t r(\mathbf{x}_t, \mathbf{u}_t), \quad \mathbf{x}_0 = \mathbf{x}. \quad (4.23)$$

Ensuring sufficient exploration is crucial in the training phase to avoid bad local optima in the non-convex loss minimization landscape. We evaluate the neural network ensemble optimistically using a softmax function:

$$\hat{V}_\theta(\mathbf{x}) = \varphi^{-1} \log \left(\sum_{i=1}^M \exp \left\{ \varphi \hat{V}_{\theta_i}(\mathbf{x}) - \log(M) \right\} \right). \quad (4.24)$$

Above, $\varphi > 0$ is a hyperparameter controlling the level of optimism. A low value yields a prediction close to the ensemble mean, whereas a high value pushes the prediction towards the ensemble maximum [235]. Whenever there is disagreement within the ensemble, the controller, therefore, overweights the optimistic ensemble members. Such optimism results in exploratory behaviour in uncertain regions.

Figure 4.4.2 illustrates the impacts of value function learning on the MPC performance. Without the terminal value function, the MPC emissions mitigation trajectory under a 30-year prediction horizon lags far behind the long-term optimal benchmark solution with maximum relative errors exceeding 65%. However, after learning the approximate value function, the controller can closely replicate the long-run optimal mitigation trajectory, with relative errors capped to less than 4%. Without the approximate terminal value function, a similar level of accuracy would require a prediction horizon of approximately 200 steps.

The result complements the dynamic stochastic general equilibrium example in Ref. [173]. The authors show that imposing the steady-state capital level as a terminal condition can guide the MPC algorithm to the optimal closed-loop behaviour in situations where the controller would otherwise fail. The learning-based approximation is particularly valuable for models where the analytical steady-state expression cannot be easily derived.

In simulated environments such as IAMs, the MPC sampling rates are less of an issue than in real-time industrial processes. Therefore, for individual IAM experiments, increasing the prediction horizon and ignoring the terminal value approximation might yield approximate solutions faster. However, when running multiple experiments, pretraining a baseline model once and fine-tuning it to specific use cases becomes more attractive, especially if the experiment involves Monte Carlo simulations over thousands of model runs. For instance, in the deterministic setting above, value learning can significantly improve MPC performance even after a few




Figure 4.4.2: Emissions control rates before and after terminal value learning. Without learning, the mitigation rate in blue lags below the long-run optimal values, leading to excessive greenhouse gas emissions and temperature increase. With the approximate value function in place, the MPC trajectory closely replicates the optimal mitigation path. Each error bar in grey shows the maximum relative error within a ten-step interval between the benchmark in black and the smoothed learning-based MPC policy in orange.

dozen training episodes. Stochastic system dynamics typically require more data — and training time — to generalize well.

Finally, determining a sufficient level of accuracy can be difficult. As such, IAMs provide a highly inaccurate description of the coupled human-climate system. Absent any political dynamics, for instance, deriving an exactly optimal carbon tax level from any IAM has little informative value for practical policy work. However, optimization-based IAMs can still provide valuable insight into what reasonable policy alternatives could look like, especially regarding the interplay between different policy variables over time. Therefore, we consider MPC solutions that largely eliminate myopia as sufficiently accurate for climate policy analysis.

4.4.2 Uncertainty propagation and constraint tightening

Risk-aware planning. A nominal MPC scheme that ignores uncertainties during planning is easy to implement and can provide reasonable approximate solutions in many applications. However, it can also lead to undesirable constraint violations. Consider an example of limiting the mean atmospheric temperature anomaly to 2°C, a binding constraint in most IAMs. In addition, assume that the transition dynamics are now stochastic, as described in Section 4.3. A deterministic prediction model without any notion of risk would drive temperatures to meet the constraint exactly, but while at the boundary, any economic or climate shock might throw the system unexpectedly over the temperature threshold. Figure 4.4.3 illustrates the example by running a nominal model predictive controller in a stochastic IAM for 1000 simulations under a maximum temperature increase constraint. While the temperature constraint holds on average, the violations are frequent, with sometimes substantial temperature overshoot.

This section introduces risk awareness to the agent via chance constraints. Since the system dynamics are uncertain, the state trajectories also become random variables. Assuming that the state constraints determine a safe operating region in our system, such as a temperature threshold for catastrophic climate impacts, we can define *risk* in terms of constraint violation probabilities. Our goal is then to modify the controller such that it can systematically avoid the constraint-breaching behaviour in Figure 4.4.3 by determining how close to the threshold it can safely operate. Risk-seeking behaviour implies tolerating a higher probability of failure and, thus, a narrower safety margin.

We proceed in two parts. The first task in cautious control is uncertainty quantification. More precisely, given an initial state, noisy transition dynamics, and a nominal input trajectory, we want to propagate the uncertainty forward in time and estimate a distribution of possible states at each step of the prediction horizon. The second task is constraint tightening. That is, based on the uncertain state trajectories, we want to find a deterministic equivalent expression for the chance constraints $\Pr(\bar{\mathbf{x}}_k \in \mathcal{X}) \geq p_x$. The main difficulties are propagating probability distributions over arbitrary nonlinear system dynamics and evaluating the joint constraint satisfaction probabilities, two operations that generally have no closed-form expression.

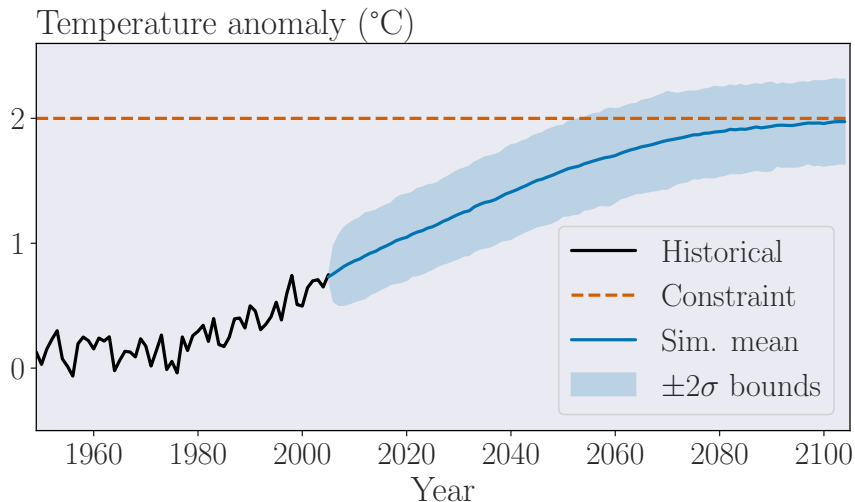


Figure 4.4.3: MPC constraint violations in an integrated assessment model. The solid black line denotes the historical temperature anomaly observations from the 1900 level [236]. In solid blue, the mean temperature anomaly over 1000 model runs with climate and economic shocks. The dashed line indicates the temperature constraint. The constraint holds on average, but the stochastic disturbances lead to frequent constraint violations with substantial temperature overshoot. The shaded area corresponds to the two standard deviation bounds from the mean.

Instead, we use approximation schemes that exploit the properties of the Gaussian noise process and the geometry of the polytopic system constraints [33, 237, 238].

Uncertainty propagation

An affine transformation of a normal distribution yields another normal distribution. Nonlinear dynamics, however, render the distribution non-Gaussian after the first step. In our setting, the simplest way to approximate the state distributions along the prediction horizon is, therefore, to assume all uncertainty to be Gaussian and linearize the system as in an extended Kalman filter [239]. We exploit the iterative MPC structure and linearize the dynamics at each step around the optimal nominal open-loop plan at the previous sampling instance. We can write the mean and covariance dynamics for the uncertain state trajectory in closed form as [238]:

$$\begin{aligned}\boldsymbol{\mu}^x(t+1) &= \bar{f}(\boldsymbol{\mu}^x(t), \mathbf{u}(t)) \\ \Sigma^x(t+1) &= J(t)\Sigma^x(t)J^\top(t) + \Sigma^w(t).\end{aligned}\tag{4.25}$$

Above, \bar{f} is the nominal nonlinear transition model, $\boldsymbol{\mu}^x$ the mean vector and Σ^x the covariance matrix of the system states, and $J(t) = \nabla \bar{f}(\boldsymbol{\mu}^x(t), \mathbf{u}(t))$ the Jacobian of the dynamics. Σ^w is the known covariance of the additive zero-mean noise vector $\mathbf{w}(t)$ that includes both the business cycle fluctuations and the temperature shocks.

The linearization strategy naturally produces some approximation errors. However, in our integrated climate-economy model, the nonlinearities mainly materialize over long periods. Therefore, the linearized dynamics provide a relatively accurate system approximation over the short MPC prediction horizon.

Constraint tightening

Once we have quantified the state uncertainty over the prediction horizon, we can perform the constraint modification. We follow the methodology detailed in Ref. [237] and Ref. [166], summarized below. The key idea is to represent the state variable at the k^{th} prediction step as $\mathbf{x}_k = \mathbf{z}_k + \mathbf{e}_k$, consisting of a deterministic nominal system state \mathbf{z}_k and an error term \mathbf{e}_k . After defining the time- k error sets, we can construct a safety margin around the nominal prediction and tighten the original constraints accordingly. Higher system uncertainty or a lower tolerance for constraint violation results in wider safety margins, inducing more cautious actions when operating close to the system boundaries. Figure 4.4.4 illustrates the procedure for a one-dimensional example.

We define the errors in terms of a *k-step probabilistic reachable set* \mathcal{R}_k [240], which satisfies:

$$\mathbf{e}_0 = 0 \implies \Pr(\mathbf{e}_k \in \mathcal{R}_k) \geq p,$$

that is, conditioned on the starting value, \mathcal{R}_k contains the k -step error term with at least the predefined probability level p . Once we obtain an estimate for \mathcal{R}_k , we can define the tightened state constraint set \mathcal{Z} for the nominal system as:




Figure 4.4.4: Illustration of the constraint tightening procedure. The system is at time k , with a planned trajectory for 30 steps forward. The original state constraint is denoted with the black dashed line. The nominal state trajectory is in solid blue, with the shaded area derived during the uncertainty propagation step. The tightened constraint in orange depends on the specified constrain violation probability p_x together with the uncertainty estimate. Keeping the nominal (mean) trajectory deterministically within the tightened constraint corresponds to meeting the original system bounds with the desired probability, as the uncertainty accumulates over the prediction horizon.

$$\mathbf{z}_k \in \mathcal{Z} = \mathcal{X} \ominus \mathcal{R}_k,$$

where \mathcal{X} denotes the original system bounds and \ominus is the Pontryagin set difference. If the nominal prediction \mathbf{z}_k stays within the tightened set, the uncertain trajectory satisfies the original system constraints with the desired probability [166]:

$$\mathbf{z}_k \in \mathcal{Z} \implies \Pr(\mathbf{x}_k = \mathbf{z}_k + \mathbf{e}_k \in \mathcal{X}) \geq \Pr(\mathbf{e}_k \in \mathcal{R}_k) \geq p_x.$$

Further, assuming a Gaussian distribution on the state evolution, \mathcal{R}_k only depends on the state covariance matrix Σ_k^x .

We focus on polytopic state constraints consisting of n_c individual half-spaces:

$$\mathcal{X} = \{\mathbf{x} \mid H\mathbf{x} \leq \mathbf{b}\}, \quad H \in \mathbb{R}^{n_c \times n_x}, \quad \mathbf{b} \in \mathbb{R}^{n_c}.$$

The assumption is not restrictive for most economic applications that typically require simple constraint sets such as constant state upper and lower bounds and rate-of-change limits. However, evaluating the joint constraints $\Pr(H\mathbf{x} \leq \mathbf{b})$ can still be prohibitively costly for online planning algorithms. A closed-form expression does not generally exist even when only working with Gaussian random variables. A typical approximation scheme is to rewrite the joint chance constraints in terms of individual half-spaces using Boole's inequality [237, 241–244]. Boole's inequality upper bounds the probability of a union of m events A_i by the sum of individual event probabilities:

$$\Pr\left(\bigcup_{i=1}^m A_i\right) \leq \sum_{i=1}^m \Pr(A_i). \quad (4.26)$$

Considering the satisfaction of each constraint as an individual event and the original joint chance constraint as an intersection of these events, we can write the complementary outcome of constraint violation as $\Pr(H\mathbf{x} > \mathbf{b}) \leq 1 - p_x$. Using Boole's union bound as a conservative approximation, we get that:

$$\Pr(H\mathbf{x} > \mathbf{b}) = \sum_{i=1}^{n_c} \Pr([H]_i^\top \mathbf{x} > b_i). \quad (4.27)$$

Then, we can evaluate each individual half-space constraint $\Pr([H]_i^\top \mathbf{x} > b_i)$ efficiently by using the Gaussian cumulative distribution function.

Finally, to ensure the satisfaction of *all* individual chance constraints with the predefined probability p_x , we need to distribute the tolerated failure probability across all constraints. We distribute the failure probability uniformly, such that each constraint must hold with $\bar{p} = 1 - (1 - p_x)n_c^{-1}$. More sophisticated optimization-based risk allocation schemes also exist [245]. Using ϕ^{-1} to denote the standard Gaussian quantile function, the k -step probabilistic reachable set becomes [166]:

$$\mathcal{R}_k = \{\mathbf{e} \mid [H]_i^\top \mathbf{e} \leq \phi^{-1}(\bar{p}) \sqrt{[H]_i^\top \Sigma_k^x [H]_i} \quad \forall i = 1, \dots, n_c\},$$

as we marginalize the error distribution according to each half-space. The resulting tightened constraint set is $\mathcal{Z} = \{\mathbf{z} \mid [H]_i^\top \mathbf{z} \leq \tilde{b}_i \quad \forall i = 1, \dots, n_c\}$ where $\tilde{b}_i = b_i - \phi^{-1}(\bar{p}) \sqrt{[H]_i^\top \Sigma_k^x [H]_i}$. We only apply the tightenings to state constraints and leave the input constraints untouched.

Simulation results

This section describes the climate policy outcomes generated under the cautious model predictive controller. The fixed constraint considered through all simulations is to limit the temperature increase to two degrees Celsius relative to a preindustrial benchmark. Figure 4.4.5 illustrates how modifying the permitted constraint violation probability affects optimal carbon price paths over the first 100 years of simulation. Intuitively, lower tolerance for risk, defined here as a lower tolerance for constraint violation, leads to more aggressive emissions mitigation policies and, thus, higher carbon taxes. For reference, the figure also shows scenarios without constraint tightening and temperature constraints.

The constraint tightening adds notable variation to the carbon price paths.⁷ Without the temperature constraint, climate policy only depends on the relatively weak climate damages, leading to low tax levels. Similarly, with the two-degree policy in place without constraint tightening, carbon emissions are still relatively cheap, leading to frequent constraint violations. In contrast, with the temperature target and constraint tightenings in place, the relatively strong temperature disturbances lead to substantial emissions cuts early on to maintain an adequate safety margin to the temperature limit. The carbon tax slowly declines due to exogenous technological improvements as time progresses.

We can interpret the results as how much additional confidence a higher carbon price brings in meeting a specific climate policy goal. For instance, over the first 60 years of simulation, gaining a 95% probability of staying below two degrees of warming translates to an average carbon tax increase of a few hundred dollars per tonne compared to the non-tightened case.

⁷The conversion factor between tC and tCO₂ is 3.666.

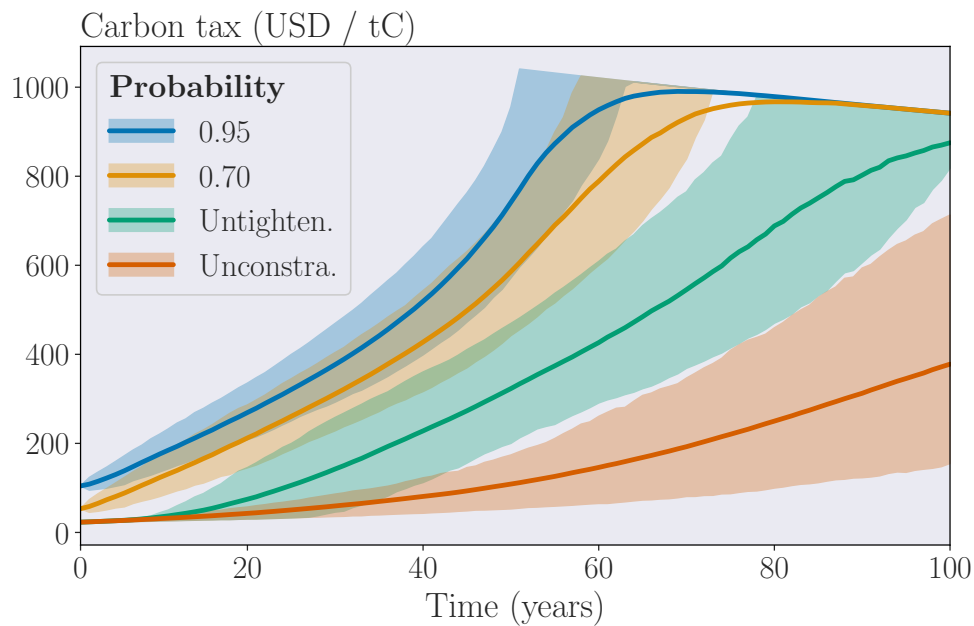


Figure 4.4.5: The evolution of carbon taxes when limiting the mean temperature anomaly to two degrees Celsius under different constraint satisfaction probabilities. The solid lines denote the mean tax levels after 1000 simulations for each scenario. The shaded region shows the likely range, defined as the 17th to 83rd percentile region. In green is a scenario without constraint tightening. The red line is a reference scenario without a temperature constraint. Lower tolerance for constraint violation leads to more aggressive emissions cuts and, therefore, higher carbon taxes.

The carbon prices presented here are comparable to previous results both under the same temperature constraint but deterministic dynamics [230, 246] and to those with substantial uncertainties but no explicit temperature constraint [30, 31, 247]. However, studies combining stochastic system dynamics with a specific temperature target are less common. The MPC-based approach, therefore, provides a way for modelling risk-aware decision rules for climate policy and complements the prior works that commonly describe risk aversion using recursive Epstein-Zin preferences [30, 31, 181].

Figure 4.4.6 illustrates the evolution of the four carbon boxes under the probabilistic two-degree policy. Without the constraint, the atmospheric temperature increases by more than four degrees. Policies that meet the target with at least

70% probability remove most emissions over the 60 first years of simulation. As a result, the accumulation in the slowest carbon boxes (top row) gradually peters out and starts declining for the fastest-evolving reservoirs (bottom row).

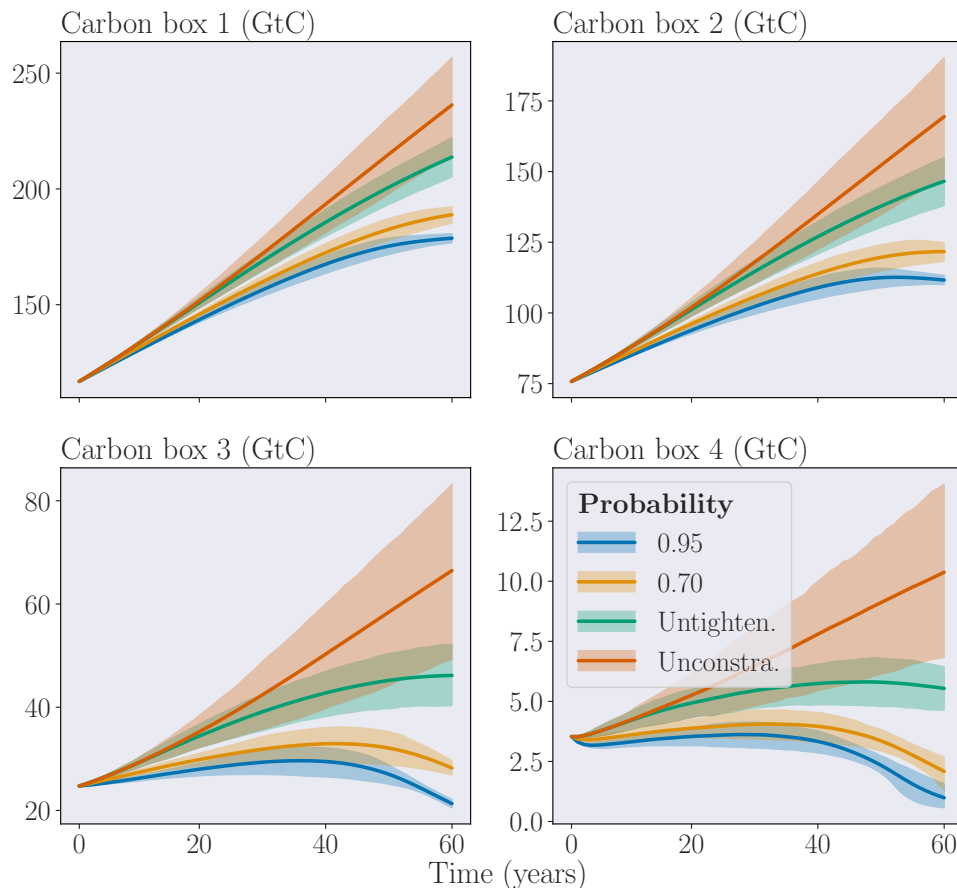


Figure 4.4.6: Carbon concentrations when limiting the mean temperature anomaly to two degrees Celsius under different constraint satisfaction probabilities. Each subplot corresponds to one of the four carbon reservoirs with different time decay rates. The solid lines denote the mean concentration levels after 1000 simulations for each scenario. The shaded region shows the likely range, defined as the 17th to 83rd percentile region. In green is a scenario without constraint tightening. The red line is a reference scenario without a temperature constraint.

The purpose of these results is to illustrate the additional insights that the MPC approach can bring to the standard IAM framework. However, while analyzing these results, it is crucial to stress the limitations of the cost-benefit IAM approach itself.

Using the above risk-aware formulation, we might argue, for instance, how much an increase in carbon price improves our chances of meeting some pre-specified temperature target. It also allows us to estimate how adjusting the risk tolerance affects the optimal carbon price relative to other temperature targets and risk tolerance levels. The overarching assumption is that we can operate under marginal adjustments, such that a somewhat looser climate policy leads to slightly higher climate damages and vice versa. However, the simplicity of the model hides away the deep structural uncertainties surrounding catastrophic climate outcomes and potentially fat-tailed damage distributions [248]. Alas, the cost-benefit comparisons based on IAM damage functions alone might be *arbitrarily inaccurate* [249].

4.5 Conclusion

This chapter works towards bridging the literatures on dynamic economic decision problems, risk-aware model predictive control (MPC), and data-driven value function approximation. Learning a value function from prior experience enables using local search methods like MPC over arbitrarily long planning horizons in repetitive tasks. The approximate value function works as a terminal reward for the control algorithm, reducing myopia. Whereas several techniques exist for solving sequential decision problems under uncertainty, the MPC approach is especially appealing for the intuitive treatment of risk via probabilistic constraint satisfaction. We consider a climate-economy integrated assessment model (IAM) as a running example. IAMs require optimizing climate policy pathways over multiple centuries with non-linear system dynamics, often coupled with considerable uncertainty, making them a challenging environment to test the proposed method.

Several avenues remain to extend the proposed framework. Throughout the chapter, we have assumed that all uncertainty stems from additive disturbances in the system dynamics. Another interesting approach is to consider parametric uncertainty and build a chance-constrained MPC algorithm based on efficient polynomial sampling techniques [250–252]. There is also further room for increasing the model size without the curse of dimensionality becoming an issue. For instance, in the climate context, this might include additional state variables for modelling geoengineering with the injection of aerosols into the stratosphere [63], climate adaptation

policies with adaptive capital stocks [156], or other forms of contingency plans for catastrophic climate change outcomes. In addition to the climate-economy setting, the proposed framework can be applied to any long-term resource allocation problems where we can represent risks in terms of constraint violation probabilities.

Appendix

4.A Sampling-based MPC validation

Whether the sampling-based MPC algorithm performs well in general non-control-affine systems remains a task for empirical validation. Especially under stochastic disturbances, it is unclear whether the sampling distribution can recover to a reasonable trajectory after a shock. We also separately evaluate the algorithm's performance in the case of binding constraints and constraint tightening. This is because the tightening procedure relies on linearizing the system dynamics around the solution at the previous sampling instance to approximate the uncertainty propagation over time. The goal is to verify that the linearization remains reasonably accurate and that penalizing the constraint-violating sample trajectories works as intended.

To perform this validation, we run the sampling-based MPC algorithm in parallel with a standard gradient-based method, fixing the random seed so that both algorithms receive the same shock realizations. We can then verify that the resulting solution paths are approximately similar. Fig. 4.A.1 shows the results for the emissions mitigation rates and Fig. 4.A.2 for the consumption rates. The sampling-based method can replicate the gradient-based outcomes well for both control variables. The mean trajectories produced by the different methods are almost indistinguishable.

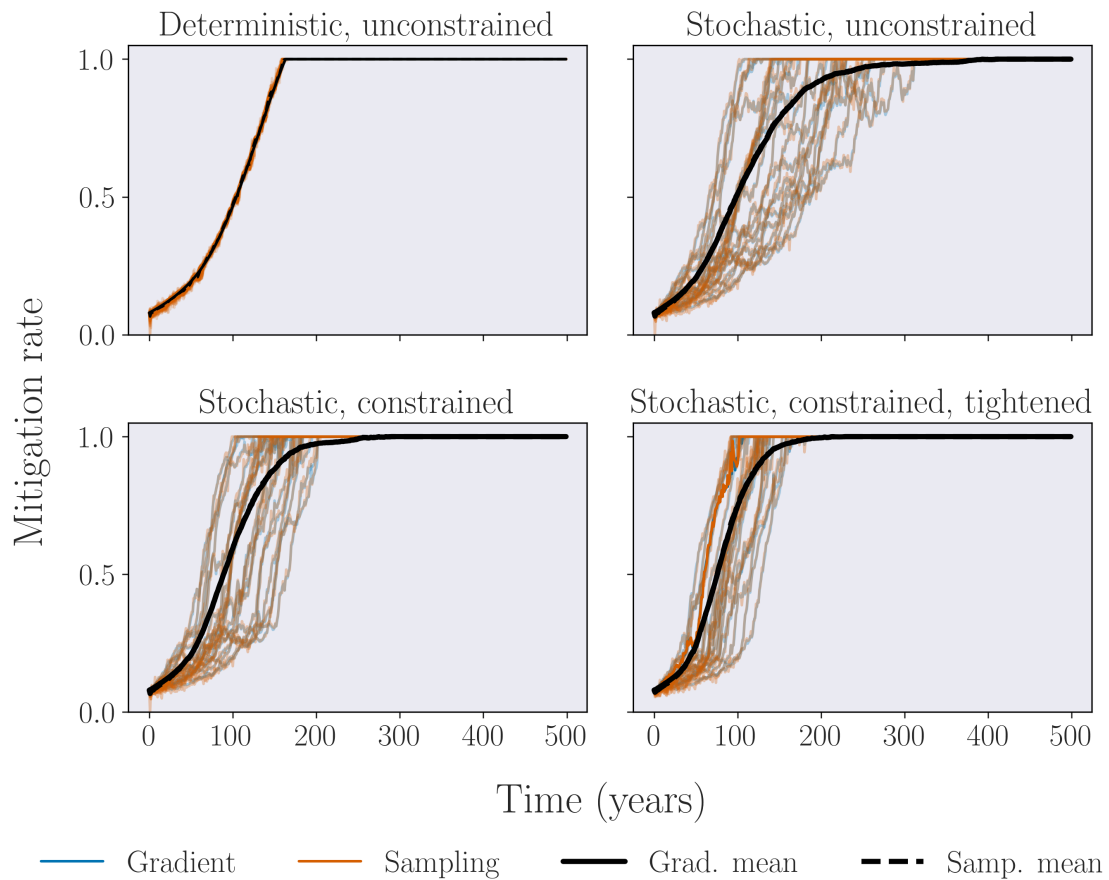


Figure 4.A.1: A comparison between gradient-based and sampling-based model-predictive control algorithms under deterministic and stochastic system dynamics. The random seed is fixed to ensure identical disturbances for all scenarios and algorithms. All data correspond to emissions mitigation rates under a 30-year prediction horizon. The black lines denote the mean trajectories of 100 sample paths, the solid line for the gradient-based method and the dashed line for the sampling-based alternative. The coloured lines indicate the first 20 samples for both methods, blue for the gradient-based and orange for the sampling-based method. The considered constraint is to limit the temperature increase to 3 °C, and the tightening procedure considers a 90% probability of satisfying this constraint. The stochastic environment involves disturbances in economic productivity and the temperature system. The mean trajectories produced by the two algorithms are almost indistinguishable.

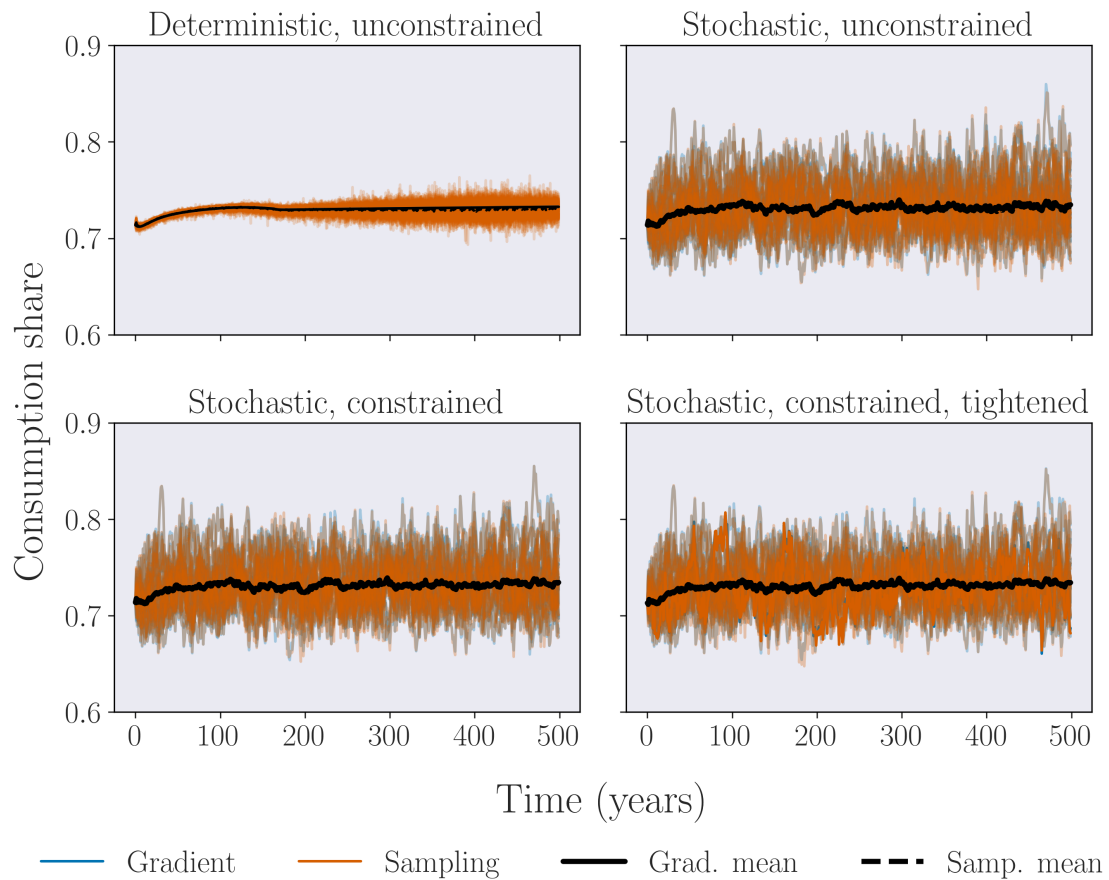


Figure 4.A.2: A comparison between gradient-based and sampling-based model-predictive control algorithms under deterministic and stochastic system dynamics. The random seed is fixed to ensure identical disturbances for all scenarios and algorithms. All data correspond to consumption shares under a 30-year prediction horizon. The black lines denote the mean trajectories of 100 sample paths, the solid line for the gradient-based method and the dashed line for the sampling-based alternative. The coloured lines indicate the first 20 samples for both methods, blue for the gradient-based and orange for the sampling-based method. The considered constraint is to limit the temperature increase to 3 °C, and the tightening procedure considers a 90% probability of satisfying this constraint. The stochastic environment involves disturbances in economic productivity and the temperature system. The mean trajectories produced by the two algorithms are almost indistinguishable.

4.B Environment calibration

Economic module

The calibration of the economic module follows the model implementations in Ref. [30] and Ref. [179]. The deterministic technology productivity trend evolves according to:

$$A(t) = A_0 \exp\left(\frac{\alpha_1(1 - \exp(-\alpha_2 t))}{\alpha_2}\right). \quad (4.28)$$

The carbon intensity of unit economic output is:

$$\sigma(t) = \sigma_0 \exp\left(\frac{\sigma_1(1 - \exp(-\sigma_2 t))}{\sigma_2}\right). \quad (4.29)$$

The world population grows as:

$$L(t) = L_0 \exp(-\Delta_L t) + L_{\max}(1 - \exp(-\Delta_L t)). \quad (4.30)$$

Land emissions from biological processes are:

$$E_{\text{land}}(t) = E_0 \exp(-\Delta_E t). \quad (4.31)$$

The equation capturing the exogenous non-CO₂ radiative forcings is:

$$\mathcal{F}_{EX}(t) = \begin{cases} \mathcal{F}_0 + \Delta_{\mathcal{F}} t & \text{if } t \leq 100 \\ \mathcal{F}_{\max} & \text{otherwise.} \end{cases} \quad (4.32)$$

The backstop technology cost path evolves according to:

$$\omega_1(t) = \frac{p_0 \sigma(t) (1 + \exp(-\Delta_p t)) \tilde{p} - 1}{\omega_2 \tilde{p}}. \quad (4.33)$$

Parameter	Description	Value
α	Production factor shares	0.30
A_0	Initial total factor productivity	0.02722
α_1	Initial productivity growth rate	0.0092
α_2	Productivity growth rate decay	0.001
σ_0	Initial carbon intensity	0.13418
σ_1	Carbon intensity change rate	-0.00730
σ_2	Carbon intensity change rate	0.003
L_0	Initial population	6514 (M.)
L_{\max}	Asymptotic population	8600 (M.)
Δ_L	Population growth rate	0.035
E_0	Initial land emissions	1.1 (GtC)
Δ_E	Land emissions change rate	0.01
ρ	Technology shock persistence	0.775
ω_2	Mitigation cost exponent	2.8
p_0	Backstop cost	1.17 (kUSD / tC)
Δ_p	Backstop cost change rate	0.005
\tilde{p}	Initial to final backstop cost ratio	2
δ	Capital depreciation rate	0.1
ξ	Intertemporal substitution elasticity	1.5
β	Discount factor	0.985
Σ^ν	Technology shock covariance	diag(0.035 ² , 0.008 ²)
ψ_1	Damage function coefficient	20.46
ψ_2	Damage function coefficient	2.0
ψ_3	Damage function coefficient	6.081
ψ_4	Damage function coefficient	6.754
\mathcal{F}_0	Initial non-CO ₂ forcings	-0.06
$\Delta_{\mathcal{F}}$	Change in non-CO ₂ forcings	0.0036
\mathcal{F}_{\max}	Maximum non-CO ₂ forcings	0.30

Table 4.B.1: Main parameters for the model economy.

Carbon cycle module

Parameter	Description	Value
a_1	Carbon reservoir fraction	0.2173
a_2	Carbon reservoir fraction	0.2240
a_3	Carbon reservoir fraction	0.2824
a_4	Carbon reservoir fraction	0.2763
τ_1	Carbon reservoir time decay	10^6
τ_2	Carbon reservoir time decay	394.9
τ_3	Carbon reservoir time decay	36.54
τ_4	Carbon reservoir time decay	4.304
r_0	Preindustrial iIRF term	35.00
r_M	Carbon impact on iIRF	0.019
r_T	Temperature impact in iIRF	4.165
M_0	Preindustrial carbon concentration	596.4

Table 4.B.2: Main parameters for the carbon cycle module.

Temperature response module

Parameter	Description	Value
$\bar{\mathcal{F}}_{2XCO_2}$	Radiative forcing parameter	3.6813
Ξ	Upper ocean heat capacity	7.3
$\bar{\Xi}_0$	Deep ocean heat capacity	106.0
η	Heat exchange parameter	0.73
Λ	CO ₂ perturbation feedback	1.13
Σ^v	Temperature shock covariance	diag(0.01284, 0)

Table 4.B.3: Main parameters for the temperature response dynamics.

4.C Implementation details

Neural network ensemble	
Network layers	[64, 64]
Activation	tanh
Num. models	6
Weight decay	0.0005
Polyak target update rate	0.005
Batch size	256
Buffer capacity	1e6
Update frequency	1 (ep.)
Gradient steps	500
Temperature (φ)	10
MPPI controller	
Prediction horizon	30
Num. samples	1024
Temperature (λ)	25

Table 4.C.1: Controller parameter values.

4.D Soft constraints

A stochastic disturbance leading to a constraint violation renders the entire solution infeasible. To ensure numerical stability for the chance-constrained simulations, we soften the original constraints to allow for temporary violations, although with a cost. That is, instead of considering a general optimization problem:

$$\begin{aligned} \min_z h(z) \\ \text{s.t. } g(z) \leq 0, \end{aligned}$$

we add an additional decision variable ϵ and a constraint violation penalty function q , and instead solve for:

$$\begin{aligned} \min_{z, \epsilon} h(z) + q(\epsilon) \\ \text{s.t. } g(z) \leq \epsilon \\ \epsilon \geq 0. \end{aligned}$$

An important property of the penalty function is that if a feasible solution to the original problem exists, the softened problem will produce the same outcome with the slack variable $\epsilon = 0$. We set a quadratic constraint violation penalty $q(\epsilon)$ to keep the slack variables as small as possible. The softening procedure only applies to state variables. The input constraints often correspond to physical rules, e.g. nonnegative consumption levels, whereas the state constraints determine the desired system properties where occasional violations can be permitted.

References

- [1] Eunice Newton Foote. Circumstances affecting the heat of the Sun’s rays. *The American Journal of Science and Arts*, 22(66):383–384, 1856.
- [2] John Tyndall. On the absorption and radiation of heat by gases and vapours, and on the physical connexion of radiation, absorption, and conduction. *Philosophical Transactions of the Royal Society of London*, 151:1–36, 1861.
- [3] Svante Arrhenius. On the influence of carbonic acid in the air upon the temperature of the ground. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(251):237–276, 1896.
- [4] IPCC. *Climate Change 2022: Impacts, adaptation, and vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, 2022.
- [5] IPCC. *Global warming of 1.5°C. An IPCC special report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. Cambridge University Press, Cambridge, UK and New York, NY, USA, 2018.
- [6] Malte Meinshausen, Jared Lewis, Christophe McGlade, Johannes Gütschow, Zebedee Nicholls, Rebecca Burdon, Laura Cozzi, and Bernd Hackmann. Realization of Paris Agreement pledges may limit warming just below 2°C. *Nature*, 604(7905):304–309, 2022.
- [7] UNEP. Emissions gap report 2021: The heat is on – A world of climate promises not yet delivered. United Nations Environment Programme, 2021.
- [8] Arthur Cecil Pigou. *The economics of welfare*. Macmillan, 1920.
- [9] Lucas Bretschger and Aimilia Pattakou. As bad as it gets: How climate damage functions affect growth and the social cost of carbon. *Environmental and Resource Economics*, 72(1):5–26, 2019.
- [10] Moritz A. Drupp, Mark C. Freeman, Ben Groom, and Frikk Nesje. Discounting disentangled. *American Economic Journal: Economic Policy*, 10(4):109–34, 2018.

- [11] World Bank. State and trends of carbon pricing, 2022. Washington, DC.
- [12] Joshua Busby. After Paris: Good enough climate governance. *Current History*, 115(777):3–9, 2016.
- [13] Simon Evans, Josh Gabbatiss, Robert McSweeney, Aruna Chandrasekhar, Ayesha Tandon, Giuliana Viglione, Zeke Hausfather, Xiaoying You, Joe Goodman, and Sylvia Hayes. COP26: Key outcomes agreed at the UN climate talks in Glasgow, 2022. Accessed May 31 2022.
- [14] Martin L. Weitzman. A voting architecture for the governance of free-driver externalities, with application to geoengineering. *The Scandinavian Journal of Economics*, 117(4):1049–1068, 2015.
- [15] Kyle C. Meng and Ashwin Rode. The social cost of lobbying over climate policy. *Nature Climate Change*, 9(6):472–476, 2019.
- [16] Kirill Orach, Andreas Duit, and Maja Schlüter. Sustainable natural resource governance under interest group competition in policy-making. *Nature Human Behaviour*, 4(9):898–909, 2020.
- [17] Aaditya Mattoo, Nadia Rocha, and Michele Ruta. *Handbook of Deep Trade Agreements*. Washington, DC: World Bank, 2020.
- [18] Eurostat. Trade by NACE Rev. 2 activity and enterprise size class, 2022. Online data code: EXT_TEC01. Accessed May 31 2022.
- [19] Michael Blanga-Gubbay, Paola Conconi, and In Song Kim. Lobbying by firms on deep trade agreements: How large firms dominate politics, 2021.
- [20] Giovanni Maggi and Ralph Ossa. Are trade agreements good for you? Working Paper 27252, National Bureau of Economic Research, 2020.
- [21] Gene M. Grossman, Phillip McCalman, and Robert W. Staiger. The “new” economics of trade agreements: From trade liberalization to regulatory convergence? *Econometrica*, 89(1):215–249, 2021.
- [22] Scott Barrett. The incredible economics of geoengineering. *Environmental and Resource Economics*, 39(1):45–54, 2008.
- [23] Joshua B. Horton. Geoengineering and the myth of unilateralism: Pressures and prospects for international cooperation. *Stanford Journal of Law, Science & Policy*, 4(1):56–69, 2011.
- [24] J. G. Shepherd. Geoengineering the climate: Science, governance and uncertainty. Monograph, Royal Society, 2009.
- [25] Debraj Ray. *A game-theoretic perspective on coalition formation*. Oxford University Press, 2007.

- [26] Kyle Hyndman and Debraj Ray. Coalition formation with binding agreements. *The Review of Economic Studies*, 74(4):1125–1147, 2007.
- [27] EM-DAT. The emergency events database - Université Catholique de Louvain (UCL) - CRED, D. Guha-Sapir, Brussels, Belgium, 2022. Accessed on 02/06/2022.
- [28] Chris Shughrue, BT Werner, and Karen C. Seto. Global spread of local cyclone damages through urban trade networks. *Nature Sustainability*, 3(8):606–613, 2020.
- [29] WMO. Atlas of mortality and economic losses from weather, climate and water extremes (1970–2019), 2021. WMO-No. 1267. World Meteorological Organization.
- [30] Yongyang Cai and Thomas S. Lontzek. The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734, 2019.
- [31] Christian P. Traeger. A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics*, 59(1):1–37, 2014.
- [32] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, 2017.
- [33] Lukas Hewing, Alexander Liniger, and Melanie N. Zeilinger. Cautious NMPC with Gaussian process dynamics for autonomous miniature race cars. *2018 European Control Conference (ECC)*, 2018.
- [34] Kendall Lowrey, Aravind Rajeswaran, Sham Kakade, Emanuel Todorov, and Igor Mordatch. Plan online, learn offline: Efficient learning and exploration via model-based control, 2019.
- [35] World Bank. World Development Indicators. Tariff rate, applied, weighted mean, all products (%), 2022. Data code: TM.TAX.MRCH.WM.AR.ZS. Accessed June 31 2022.
- [36] Dani Rodrik. What do trade agreements really do? *Journal of Economic Perspectives*, 32(2):73–90, 2018.
- [37] Jean-Luc Angot, Geneviève Bastid Burdeau, Christophe Bellmann, Sophie Devienne, Lionel Fontagné, Roger Genet, Géraud Guibert, Sabrina Robert-Cuendet, and Katheline Schubert. L’impact de l’accord économique et commercial global entre l’Union Européenne et le Canada (AECG/CETA) sur l’environnement, le climat et la santé, 2017.

- [38] Todd Allee, Manfred Elsig, and Andrew Lugg. Is the European Union trade deal with Canada new or recycled? A text-as-data approach. *Global Policy*, 8(2):246–252, 2017.
- [39] Scott Barrett. Strategic environmental policy and international trade. *Journal of Public Economics*, 54(3):325 – 338, 1994.
- [40] Robert W. Staiger and Alan O. Sykes. International trade, national treatment, and domestic regulation. *The Journal of Legal Studies*, 40(1):149–203, 2011.
- [41] Georges A. Tanguay. Strategic environmental policies under international duopolistic competition. *International Tax and Public Finance*, 8(5):793–811, 2001.
- [42] Roberto Burguet and Jaume Sempere. Trade liberalization, environmental policy, and welfare. *Journal of Environmental Economics and Management*, 46(1):25 – 37, 2003.
- [43] Yu-Bong Lai and Chia-Hsien Hu. Trade agreements, domestic environmental regulation, and transboundary pollution. *Resource and Energy Economics*, 30(2):209 – 228, 2008.
- [44] Azim Essaji. Trade liberalization, standards and protection. *The B.E. Journal of Economic Analysis & Policy*, 10(1), 2010.
- [45] Ronald Fischer and Pablo Serra. Standards and protection. *Journal of International Economics*, 52(2):377 – 400, 2000.
- [46] Mathieu Parenti and Gonzague Vannoorenberghe. A simple theory of deep trade integration. *Working Paper*, 2019.
- [47] Gene M. Grossman and Elhanan Helpman. Protection for sale. *The American Economic Review*, 84(4):833–850, 1994.
- [48] Kyle Bagwell and Robert W. Staiger. Reciprocity, non-discrimination and preferential agreements in the multilateral trading system. *European Journal of Political Economy*, 17(2):281 – 325, 2001.
- [49] Nuno Limão. Trade policy, cross-border externalities and lobbies: Do linked agreements enforce more cooperative outcomes? *Journal of International Economics*, 67(1):175 – 199, 2005.
- [50] Kristy Buzard. Trade agreements in the shadow of lobbying. *Review of International Economics*, 25(1):21–43, 2017.
- [51] Andreas Polk, Armin Schmutzler, and Adrian Müller. Lobbying and the power of multinational firms. *European Journal of Political Economy*, 36:209–227, 2014.

- [52] Andreas Polk and Armin Schmutzler. Lobbying against environmental regulation vs. lobbying for loopholes. *European Journal of Political Economy*, 21(4):915–931, 2005.
- [53] James A. Brander. Intra-industry trade in identical commodities. *Journal of International Economics*, 11(1):1 – 14, 1981.
- [54] James Brander and Paul Krugman. A ‘reciprocal dumping’ model of international trade. *Journal of International Economics*, 15(3):313 – 321, 1983.
- [55] Richard Baldwin. Politically realistic objective functions and trade policy profs and tariffs. *Economics Letters*, 24(3):287 – 290, 1987.
- [56] Kyle Bagwell and Robert W. Staiger. Enforcement, private political pressure, and the general agreement on tariffs and trade/world trade organization escape clause. *The Journal of Legal Studies*, 34(2):471–513, 2005.
- [57] David Vogel. *Trading up : Consumer and environmental regulation in a global economy*. Harvard University Press, 1995.
- [58] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [59] Iain Dunning, Joey Huchette, and Miles Lubin. JuMP: A modeling language for mathematical optimization. *SIAM Review*, 59(2):295–320, 2017.
- [60] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- [61] Andrey Stoyanov. Trade policy of a free trade agreement in the presence of foreign lobbying. *Journal of International Economics*, 77(1):37–49, 2009.
- [62] Daniel Heyen and Jere Lehtomaa. Solar geoengineering governance: A dynamic framework of farsighted coalition formation. *Oxford Open Climate Change*, 1(1), 2021.
- [63] National Academies of Sciences, Engineering, and Medicine. *Reflecting sunlight: Recommendations for solar geoengineering research and research governance*. The National Academies Press, Washington, DC, 2021.
- [64] Anthony R. Harding, Katharine Ricke, Daniel Heyen, Douglas G. MacMartin, and Juan Moreno-Cruz. Climate econometric models indicate solar geoengineering would reduce inter-country income inequality. *Nature Communications*, 11(1), 2020.
- [65] David W. Keith and Douglas G. MacMartin. A temporary, moderate and responsive scenario for solar geoengineering. *Nature Climate Change*, 5(3):201–206, 2015.

- [66] David G. Victor, M. Granger Morgan, Jay Apt, John Steinbruner, and Katharine Ricke. The geoengineering option: A last resort against global warming? *Foreign Affairs*, 88(2):64–76, 2009.
- [67] Todd Sandler. Collective action and geoengineering. *The Review of International Organizations*, pages 1–21, 2017. 00000.
- [68] Steve Rayner, Clare Heyward, Tim Kruger, Nick Pidgeon, Catherine Redgwell, and Julian Savulescu. The Oxford Principles. *Climatic Change*, 121(3):499–512, 2013.
- [69] Janos Pasztor, Cynthia Scharf, and Kai-Uwe Schmidt. How to govern geoengineering? *Science*, 357(6348):231–231, 2017.
- [70] Robert N. Stavins and Robert C. Stowe. Governance of the deployment of solar geoengineering. Technical report, Harvard Project on Climate Agreements, Cambridge, Massachusetts, 2019.
- [71] Jesse L. Reynolds. *The governance of solar geoengineering: Managing climate change in the Anthropocene*. Cambridge University Press, 2019.
- [72] Janos Pasztor and Nicholas Harrison. Governing climate-altering approaches. *Global Policy*, 2021.
- [73] Debraj Ray and Rajiv Vohra. Coalition formation. In *Handbook of Game Theory with Economic Applications*, volume 4, pages 239–326. Elsevier, 2015.
- [74] Katharine L Ricke, Juan B Moreno-Cruz, and Ken Caldeira. Strategic incentives for climate geoengineering coalitions to exclude broad participation. *Environmental Research Letters*, 8(1):014021, 2013.
- [75] Johannes Emmerling and Massimo Tavoni. Exploration of the interactions between mitigation and solar radiation settings management in cooperative and non-cooperative international governance settings. *Global Environmental Change*, 53:244–251, 2018.
- [76] Daniel Heyen, Joshua Horton, and Juan Moreno-Cruz. Strategic implications of counter-geoengineering: Clash or cooperation? *Journal of Environmental Economics and Management*, 95:153–177, 2019.
- [77] Scott Barrett. Self-enforcing international environmental agreements. *Oxford Economic Papers*, 46:878–894, 1994.
- [78] Carlo Carraro and Domenico Siniscalco. Strategies for the international protection of the environment. *Journal of Public Economics*, 52(3):309–328, 1993.
- [79] Johan Eyckmans and Michael Finus. Coalition formation in a global warming game: How the design of protocols affects the success of environmental treaty-making. *Natural Resource Modeling*, 19(3):323–358, 2006.

- [80] Michael Finus. Game theoretic research on the design of international environmental agreements: Insights, critical remarks, and future challenges. *International Review of Environmental and Resource Economics*, 2(1):29–67, 2008.
- [81] Parkash Chander. The gamma-core and coalition formation. *International Journal of Game Theory*, 35(4):539–556, 2007.
- [82] Aart de Zeeuw. Dynamic effects on the stability of international environmental agreements. *Journal of Environmental Economics and Management*, 55(2):163–174, 2008.
- [83] Dritan Osmani and Richard Tol. Toward farsightedly stable international environmental agreements. *Journal of Public Economic Theory*, 11(3):455–492, 2009.
- [84] Jobst Heitzig and Ulrike Kornek. Bottom-up linking of carbon markets under far-sighted cap coordination and reversibility. *Nature Climate Change*, 8(3):204–209, 2018.
- [85] Johannes Emmerling, Ulrike Kornek, Valentina Bosetti, and Kai Lessmann. Climate thresholds and heterogeneous regions: Implications for coalition formation. *The Review of International Organizations*, 2020.
- [86] Muhammet A. Bas and Aseem Mahajan. Contesting the climate. *Climatic Change*, 162(4):1985–2002, 2020.
- [87] Marshall Burke, Solomon M. Hsiang, and Edward Miguel. Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239, 2015.
- [88] Scott Barrett. International cooperation for sale. *European Economic Review*, 45(10):1835–1850, 2001.
- [89] Wilfried Rickels, Martin F. Quaas, Katharine Ricke, Johannes Quaas, Juan Moreno-Cruz, and Sjak Smulders. Who turns the global thermostat and by how much? *Energy Economics*, 91:104852, 2020.
- [90] Juan B. Moreno-Cruz. Mitigation and the geoengineering threat. *Resource and Energy Economics*, 41:248–263, 2015.
- [91] Vassiliki Manoussi and Anastasios Xepapadeas. Cooperation and competition in climate change policies: Mitigation and climate engineering when countries are asymmetric. *Environmental and Resource Economics*, 66(4):605–627, 2017.
- [92] Anna Lou Abatayo, Valentina Bosetti, Marco Casari, Riccardo Ghidoni, and Massimo Tavoni. Solar geoengineering may lead to excessive cooling and high strategic uncertainty. *Proceedings of the National Academy of Sciences*, 117(24):13393–13398, 2020.

- [93] A. Parker, J. B. Horton, and D. W. Keith. Stopping solar geoengineering through technical means: A preliminary assessment of counter-geoengineering. *Earth's Future*, 6(8):1058–1065, 2018.
- [94] Edward A. Parson and David W. Keith. End the deadlock on governance of geoengineering research. *Science*, 339(6125):1278–1279, 2013.
- [95] Andy Parker. Governing solar geoengineering research as it leaves the laboratory. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2031):20140173, 2014.
- [96] Daniel Heyen. Strategic conflicts on the horizon: R&D incentives for Environmental technologies. *Climate Change Economics*, 7(4):1650013, 2016.
- [97] Martin F. Quaas, Johannes Quaas, Wilfried Rickels, and Olivier Boucher. Are there reasons against open-ended research into solar radiation management? A model of intergenerational decision-making under uncertainty. *Journal of Environmental Economics and Management*, 84:1–17, 2017.
- [98] Armando Gomes and Philippe Jehiel. Dynamic processes of social and economic interactions: On the persistence of inefficiencies. *Journal of Political Economy*, 113(3):626–667, 2005. Publisher: The University of Chicago Press.
- [99] Vítor V. Vasconcelos, Phillip M. Hannam, Simon A. Levin, and Jorge M. Pacheco. Coalition-structured governance improves cooperation to provide public goods. *Scientific Reports*, 10(1):9194, 2020.
- [100] Kalyan Chatterjee, Bhaskar Dutta, Debraj Ray, and Kunal Sengupta. A noncooperative theory of coalitional bargaining. *The Review of Economic Studies*, 60(2):463–477, 1993.
- [101] Francis Bloch. Sequential formation of coalitions in games with externalities and fixed payoff division. *Games and Economic Behavior*, 14(1):90–123, 1996.
- [102] Christopher Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica*, 53(3):613–628, 1985.
- [103] Jere Lehtomaa and Clément Renoir. Catastrophes naturelles et dynamiques économiques: Application aux cyclones tropicaux. *Revue française d'économie*, Forthcoming, 2022.
- [104] Solomon M Hsiang and Amir S Jina. The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. Working Paper 20352, National Bureau of Economic Research, 2014.
- [105] Tatyana Deryugina, Laura Kawano, and Steven Levitt. The economic impact of Hurricane Katrina on its victims: Evidence from individual tax returns. *American Economic Journal: Applied Economics*, 10(2):202–33, 2018.

- [106] Michael Bourdeau-Brien and Lawrence Kryzanowski. Natural disasters and risk aversion. *Journal of Economic Behavior & Organization*, 177:818–835, 2020.
- [107] Andrew Gettelman, David N Bresch, Chihchieh C Chen, John E Truesdale, and Julio T Bacmeister. Projections of future tropical cyclone damage with a high-resolution global climate model. *Climatic Change*, 146(3-4):575–585, 2018.
- [108] IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2021.
- [109] Gabriela Aznar Siguan and David N Bresch. Climada v1: A global weather and climate risk assessment platform. *Geoscientific Model Development*, 12:3085–3097, 2019.
- [110] Lucas Bretschger, Filippo Lechthaler, Sebastian Rausch, and Lin Zhang. Knowledge diffusion, endogenous growth, and the costs of global climate policy. *European Economic Review*, 93:47–72, 2017.
- [111] Kenneth R Knapp, Michael C Kruk, David H Levinson, Howard J Diamond, and Charles J Neumann. The international best track archive for climate stewardship (IBTrACS) unifying tropical cyclone data. *Bulletin of the American Meteorological Society*, 91(3):363–376, 2010.
- [112] Miguel O. Román, Zhuosen Wang, Qingsong Sun, Virginia Kalb, Steven D. Miller, Andrew Molthan, Lori Schultz, Jordan Bell, Eleanor C. Stokes, Bharatendu Pandey, Karen C. Seto, Dorothy Hall, Tomohiro Oda, Robert E. Wolfe, Gary Lin, Navid Golpayegani, Sadashiva Devadiga, Carol Davidson, Sudipta Sarkar, Cid Praderas, Jeffrey Schmaltz, Ryan Boller, Joshua Stevens, Olga M. Ramos González, Elizabeth Padilla, José Alonso, Yasmín Detrés, Roy Armstrong, Ismael Miranda, Yasmín Conte, Nitza Marrero, Kytt MacManus, Thomas Esch, and Edward J. Masuoka. NASA’s Black Marble nighttime lights product suite. *Remote Sensing of Environment*, 210:113 – 143, 2018.
- [113] Stéphane Hallegatte and Valentin Przulski. *The economics of natural disasters: Concepts and methods*. The World Bank, 2010.
- [114] Preeya S. Mohan, Bazoumana Ouattara, and Eric Strobl. Decomposing the macroeconomic effects of natural disasters: A national income accounting perspective. *Ecological Economics*, 146:1–9, 2018.
- [115] Mark Skidmore and Hideki Toya. Do natural disasters promote long-run growth? *Economic Inquiry*, 40(4):664–687, 2002.
- [116] Eric Strobl. The economic growth impact of natural disasters in developing countries: Evidence from hurricane strikes in the Central American and Caribbean regions. *Journal of Development Economics*, 97(1):130–141, 2012.
- [117] Ken-Ichi Akao and Hiroaki Sakamoto. A theory of disasters and long-run growth. *Journal of Economic Dynamics and Control*, 95:89–109, 2018.

- [118] Laura Bakkensen and Lint Barrage. Climate shocks, cyclones, and economic growth: Bridging the micro-macro gap. *NBER Working Paper No. 24893.*, 2021.
- [119] Thomas Douenne. Disaster risks, disaster strikes, and economic growth: The role of preferences. *Review of Economic Dynamics*, 38:251–272, 2020.
- [120] Masako Ikefuji and Ryo Horii. Natural disasters in a two-sector model of endogenous growth. *Journal of Public Economics*, 96(9-10):784–796, 2012.
- [121] Holger Strulik and Timo Trimborn. Natural disasters and macroeconomic performance. *Environmental and Resource Economics*, 72(4):1069–1098, 2019.
- [122] Yasuhide Okuyama, Geoffrey J. D. Hewings, and Michael Sonis. *Measuring Economic Impacts of Disasters: Interregional Input-Output Analysis Using Sequential Interindustry Model*, pages 77–101. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [123] Jesús Crespo Cuaresma, Jaroslava Hlouskova, and Michael Obersteiner. Natural disasters as creative destruction? Evidence from developing countries. *Economic Inquiry*, 46(2):214–226, 2008.
- [124] Norman V Loayza, Eduardo Olaberria, Jamele Rigolini, and Luc Christiaensen. Natural disasters and growth: Going beyond the averages. *World Development*, 40(7):1317–1336, 2012.
- [125] Stéphane Hallegatte and Patrice Dumas. Can natural disasters have positive consequences? Investigating the role of embodied technical change. *Ecological Economics*, 68(3):777–786, 2009.
- [126] Matthew E Kahn. The death toll from natural disasters: The role of income, geography, and institutions. *Review of Economics and Statistics*, 87(2):271–284, 2005.
- [127] Hideki Toya and Mark Skidmore. Economic development and the impacts of natural disasters. *Economics Letters*, 94(1):20–25, 2007.
- [128] Gabriel Felbermayr and Jasmin Gröschl. Natural disasters and the effect of trade on income: A new panel IV approach. *European Economic Review*, 58:18–30, 2013.
- [129] Gabriel Felbermayr and Jasmin Gröschl. Naturally negative: The growth effects of natural disasters. *Journal of Development Economics*, 111:92–106, 2014.
- [130] Ilan Noy. The macroeconomic consequences of disasters. *Journal of Development Economics*, 88(2):221–231, 2009.
- [131] Robert Mendelsohn, Kerry Emanuel, Shun Chonabayashi, and Laura Bakkensen. The impact of climate change on global tropical cyclone damage. *Nature Climate Change*, 2(3):205–209, 2012.

- [132] D. N. Bresch and G. Aznar-Siguan. Climada v1.4.1: Towards a globally consistent adaptation options appraisal tool. *Geoscientific Model Development*, 14(1):351–363, 2021.
- [133] S. Eberenz, D. Stocker, T. Rössli, and D. N. Bresch. Exposure data for global physical risk assessment. *Earth System Science Data Discussions*, 2019:1–19, 2019.
- [134] Xi Chen and William D. Nordhaus. Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594, 2011.
- [135] CIESIN. Gridded population of the world, version 4 (GPWv4): Administrative unit center points with population estimates. Population Count, v4.10., 2016. Center for International Earth Science Information Network. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC).
- [136] World Bank. Building the world bank’s wealth accounts: Methods and data, 2018. Environment and Natural Resources Global Practice, World Bank.
- [137] Stéphane Hallegatte. The use of synthetic hurricane tracks in risk analysis and climate change damage assessment. *Journal of Applied Meteorology and Climatology*, 46(11):1956–1966, 2007.
- [138] Roger A Jr Pielke. Future economic damage from tropical cyclones: sensitivities to societal and climate changes. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1860):2717–2729, 2007.
- [139] William D Nordhaus. The economics of hurricanes and implications of global warming. *Climate Change Economics*, 1(01):1–20, 2010.
- [140] Luisito Bertinelli, Preeya Mohan, and Eric Strobl. Hurricane damage risk assessment in the Caribbean: An analysis using synthetic hurricane events and nightlight imagery. *Ecological Economics*, 124:135–144, 2016.
- [141] Kerry Emanuel. Global warming effects on us hurricane damage. *Weather, Climate, and Society*, 3(4):261–268, 2011.
- [142] Kathleen Sullivan Sealy and Eric Strobl. A hurricane loss risk assessment of coastal properties in the Caribbean: Evidence from the Bahamas. *Ocean & Coastal Management*, 149:42–51, 2017.
- [143] Kerry Emanuel. Downscaling CMIP5 climate models shows increased tropical cyclone activity over the 21st century. *Proceedings of the National Academy of Sciences*, 110(30):12219–12224, 2013.
- [144] Lucas Bretschger, Roger Ramer, and Florentine Schwark. Growth effects of carbon policies: applying a fully dynamic cge model with heterogeneous capital. *Resource and Energy Economics*, 33(4):963–980, 2011.

- [145] P Armington. A theory of demand for products distinguished by place of production. *IMF Staff Working Papers*, 1969.
- [146] Badri Narayanan, Angel Aguiar, and Robert McDougall. Global trade, assistance, and production: The GTAP 8 data base, 2012.
- [147] S. Paltsev, J.M. Reilly, H.D. Jacoby, R.S. Eckaus, J. McFarland, M. Sarofim, M. Asadoorian, and M. Babiker. The MIT emissions prediction and policy analysis (EPPA) model: Version 4. *Joint Program Report Series Report 125*, 2005.
- [148] Aaron B. Gertz, James B. Davies, and Samantha L. Black. A CGE framework for modeling the economics of flooding and recovery in a major urban area. *Risk Analysis*, 39(6):1314–1341, 2019.
- [149] Tatyana Deryugina. The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy*, 9(3):168–98, 2017.
- [150] James M. Shultz, Jill Russell, and Zelde Espinel. Epidemiology of tropical cyclones: The dynamics of disaster, disease, and development. *Epidemiologic Reviews*, 27(1):21–35, 2005.
- [151] Robbie M. Parks, G. Brooke Anderson, Rachel C. Nethery, Ana Navas-Acien, Francesca Dominici, and Marianthi-Anna Kioumourtzoglou. Tropical cyclone exposure is associated with increased hospitalization rates in older adults. *Nature Communications*, 12(1):1545, 2021.
- [152] Jesse Anttila-Hughes and Solomon Hsiang. Destruction, disinvestment, and death: Economic and human losses following environmental disaster, 2013. SSRN Working Paper.
- [153] CBO. Potential increases in hurricane damage in the United States: Implications for the federal budget. Congressional Budget Office, 2016.
- [154] Paul M Romer. Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102, 1990.
- [155] Thomas R Knutson, Joseph J Sirutis, Ming Zhao, Robert E Tuleya, Morris Bender, Gabriel A Vecchi, Gabriele Villarini, and Daniel Chavas. Global projections of intense tropical cyclone activity for the late twenty-first century from dynamical downscaling of CMIP5/RCP4.5 scenarios. *Journal of Climate*, 28(18):7203–7224, 2015.
- [156] Antony Millner and Simon Dietz. Adaptation to climate change and economic growth in developing countries. *Environment and Development Economics*, 20(3):380–406, 2015.
- [157] Stéphane Hallegatte and Adrien Vogt-Schilb. Are losses from natural disasters more than just asset losses? The role of capital aggregation, sector interactions, and investment behaviors, 2016. Policy Research Working Paper; No. 7885. World Bank, Washington, DC.

- [158] Lorenzo Carrera, Gabriele Standardi, Francesco Bosello, and Jaroslav Mysiak. Assessing direct and indirect economic impacts of a flood event through the integration of spatial and computable general equilibrium modelling. *Environmental Modelling & Software*, 63:109–122, 2015.
- [159] Daron Acemoglu. *Introduction to Modern Economic Growth*. Princeton University Press, 2009.
- [160] Morten I. Lau, Andreas Pahlke, and Thomas F. Rutherford. Approximating infinite-horizon models in a complementarity format: A primer in dynamic general equilibrium analysis. *Journal of Economic Dynamics and Control*, 26(4):577 – 609, 2002.
- [161] Sergey Paltsev. Moving from static to dynamic general equilibrium economic models. *MIT Joint Program on the Science and Policy of Global Change, Technical note No. 4*, 2004.
- [162] Lars Mathiesen. *Computation of economic equilibria by a sequence of linear complementarity problems*, pages 144–162. Springer Berlin Heidelberg, Berlin, Heidelberg, 1985.
- [163] Thomas F. Rutherford. Applied general equilibrium modeling with MPSGE as a GAMS subsystem: An overview of the modeling framework and syntax. *Computational Economics*, 14:1–46, 1999.
- [164] Michael C Ferris and Todd S Munson. Complementarity problems in GAMS and the PATH solver. *Journal of Economic Dynamics and Control*, 24(2):165–188, 2000.
- [165] Fuad Hasanov. Housing, household portfolio, and intertemporal elasticity of substitution: Evidence from the consumer expenditure survey, 2007.
- [166] Lukas Hewing, Juraj Kabzan, and Melanie N. Zeilinger. Cautious model predictive control using Gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6):2736–2743, 2020.
- [167] Lukas Hewing, Kim P. Wabersich, Marcel Menner, and Melanie N. Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):269–296, 2020.
- [168] Tuomas Haarnoja, Sehoon Ha, Aurick Zhou, Jie Tan, George Tucker, and Sergey Levine. Learning to walk via deep reinforcement learning, 2019.
- [169] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

- [170] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [171] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [172] Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based model predictive control for safe exploration, 2018.
- [173] Lars Grüne, Willi Semmler, and Marleen Stieler. Using nonlinear model predictive control for dynamic decision problems in economics. *Journal of Economic Dynamics and Control*, 60:112–133, 2015.
- [174] F. Herzog, S. Keel, G. Dondi, L.M. Schumann, and H.P. Geering. Model predictive control for portfolio selection. In *2006 American Control Conference*, pages 8 pp.–, 2006.
- [175] Johannes Brumm and Simon Scheidegger. Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612, 2017.
- [176] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, USA, 1957.
- [177] Dimitri P. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4):310–334, 2005.
- [178] William D. Nordhaus. An optimal transition path for controlling greenhouse gases. *Science*, 258(5086):1315–1319, 1992.
- [179] W.D. Nordhaus. *A question of balance: Weighing the options on global warming policies*. Yale University Press, 2008.
- [180] William Nordhaus. Evolution of modeling of the economics of global warming: Changes in the DICE model, 1992–2017. *Climatic Change*, 148(4):623–640, 2018.
- [181] Kent D. Daniel, Robert B. Litterman, and Gernot Wagner. Declining CO₂ price paths. *Proceedings of the National Academy of Sciences*, 116(42):20886–20891, 2019.
- [182] Derek Lemoine and Christian Traeger. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–66, 2014.

- [183] Svernn Jensen and Christian P. Traeger. Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, 69:104–125, 2014.
- [184] Derek Lemoine and Christian P. Traeger. Economics of tipping the climate dominoes. *Nature Climate Change*, 6(5):514–519, 2016.
- [185] Derek Lemoine and Christian P. Traeger. Ambiguous tipping points. *Journal of Economic Behavior & Organization*, 132:5–18, 2016.
- [186] Soheil Shayegh and Valerie M. Thomas. Adaptive stochastic integrated assessment modeling of optimal greenhouse gas emission reductions. *Climatic Change*, 128(1):1–15, 2015.
- [187] Yongyang Cai, Timothy M. Lenton, and Thomas S. Lontzek. Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nature Climate Change*, 6(5):520–525, 2016.
- [188] James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model predictive control theory, computation, and design*. Nob Hill Publishing, 2017.
- [189] Timm Faulwasser, Lars Grüne, and Matthias A. Müller. *Economic Nonlinear Model Predictive Control*. Now Foundations and Trends, 2018.
- [190] Kim Peter Wabersich and Melanie N. Zeilinger. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, 129:109597, 2021.
- [191] David D. Fan, Ali Agha, and Evangelos A. Theodorou. Deep learning tubes for tube MPC, 2020.
- [192] Mohak Bhardwaj, Sanjiban Choudhury, and Byron Boots. Blending MPC & value function approximation for efficient reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- [193] Joao Paulo Braga, Willi Semmler, and Dieter Grass. De-risking of green investments through a green bond market – Empirics and a dynamic model. *Journal of Economic Dynamics and Control*, 131:104201, 2021.
- [194] Simon Quemin and Raphaël Trotignon. Emissions trading with rolling horizons. *Journal of Economic Dynamics and Control*, 125:104099, 2021.
- [195] Bing Chu, Stephen Duncan, Antonis Papachristodoulou, and Cameron Hepburn. Using economic model predictive control to design sustainable policies for mitigating climate change. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 406–411, 2012.
- [196] Bing Chu, Stephen Duncan, Antonis Papachristodoulou, and Cameron Hepburn. Analysis and control design of sustainable policies for greenhouse gas emissions. *Applied Thermal Engineering*, 53(2):420–431, 2013.

- [197] Thierry Brechet, Carmen Camacho, and Vladimir Veliov. Model predictive control, the economy, and the issue of global warming. *Annals of Operations Research*, 220(1):25–48, 2014.
- [198] Steven R. Weller, Salman Hafeez, and Christopher M. Kellett. A receding horizon control approach to estimating the social cost of carbon in the presence of emissions and temperature uncertainty. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5384–5390, 2015.
- [199] Christopher M. Kellett, Steven R. Weller, Timm Faulwasser, Lars Grüne, and Willi Semmler. Feedback, dynamics, and optimal control in climate economics. *Annual Reviews in Control*, 47:7–20, 2019.
- [200] Florian Herzog, Gabriel Dondi, and Hans P. Geering. Stochastic model predictive control and portfolio optimization. *International Journal of Theoretical and Applied Finance*, 10(02):203–233, 2007.
- [201] James A. Primbs. Portfolio optimization applications of stochastic receding horizon control. In *2007 American Control Conference*, pages 1811–1816, 2007.
- [202] Minyong Shin, Joo Hyung Lee, and James A. Primbs. Constrained stochastic mpc under multiplicative noise for financial applications. In *49th IEEE Conference on Decision and Control (CDC)*, pages 6101–6106, 2010.
- [203] Tobias Gybel Hovgaard, Lars F. S. Larsen, and John Bagterp Jørgensen. Robust economic MPC for a power management scenario with uncertainties. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 1515–1520, 2011.
- [204] J. M. Maestre, P. Velarde, I. Jurado, C. Ocampo-Martínez, I. Fernández, B. Isla Tejera, and J. R. del Prado. An application of chance-constrained model predictive control to inventory management in hospitalary pharmacy. In *53rd IEEE Conference on Decision and Control*, pages 5901–5906, 2014.
- [205] Victor Duarte. Machine learning for continuous-time economics, 2018.
- [206] Jesús Fernández-Villaverde, Galo Nuño, George Sorg-Langhans, and Maximilian Vogler. Solving high-dimensional dynamic programming problems using deep learning, 2020.
- [207] Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *International Economic Review*, forthcoming, 2022.
- [208] Marco C. Campi and Simone Garatti. *Introduction to the Scenario Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2018.
- [209] J. M. Maestre, L. Raso, P. J. van Overloop, and B. de Schutter. Distributed tree-based model predictive control on an open water system. In *2012 American Control Conference (ACC)*, pages 1985–1990, 2012.

- [210] Dimitri Bertsekas. Lessons from AlphaZero for optimal, model predictive, and adaptive control, 2021.
- [211] Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning, 2017.
- [212] Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 1101–1112. PMLR, 2020.
- [213] Grady Williams, Paul Drews, Brian Goldfain, James M. Rehg, and Evangelos A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *Trans. Rob.*, 34(6):1603–1622, 2018.
- [214] Dylan M. Asmar, Ransalu Senanayake, Shawn Manuel, and Mykel J. Kochenderfer. Model predictive optimized path integral strategies, 2022.
- [215] Homanga Bharadhwaj, Kevin Xie, and Florian Shkurti. Model-predictive control via cross-entropy and gradient-based optimization. In Alexandre M. Bayen, Ali Jadbabaie, George Pappas, Pablo A. Parrilo, Benjamin Recht, Claire Tomlin, and Melanie Zeilinger, editors, *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 277–286. PMLR, 2020.
- [216] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622, Lille, France, 2015. PMLR.
- [217] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [218] Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David Crandall, and Dhruv Batra. Why m heads are better than one: Training a diverse ensemble of deep networks, 2015.
- [219] Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [220] Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley series in probability and statistics. John Wiley & Sons Inc., Hoboken, NJ, 2005.

- [221] William D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, 2017.
- [222] Martin L. Weitzman. GHG targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory*, 14(2):221–244, 2012.
- [223] Peter H. Howard and Thomas Sterner. Few and not so far between: A meta-analysis of climate damage estimates. *Environmental and Resource Economics*, 68(1):197–225, Sep 2017.
- [224] F. Joos, R. Roth, J. S. Fuglestedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. Segsneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmospheric Chemistry and Physics*, 13(5):2793–2825, 2013.
- [225] R. J. Millar, Z. R. Nicholls, P. Friedlingstein, and M. R. Allen. A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions. *Atmospheric Chemistry and Physics*, 17(11):7213–7228, 2017.
- [226] C. J. Smith, P. M. Forster, M. Allen, N. Leach, R. J. Millar, G. A. Passerello, and L. A. Regayre. FAIR v1.3: A simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297, 2018.
- [227] N. J. Leach, S. Jenkins, Z. Nicholls, C. J. Smith, J. Lynch, M. Cain, T. Walsh, B. Wu, J. Tsutsui, and M. R. Allen. FAIR v2.0.0: A generalized impulse response model for climate uncertainty and future scenario exploration. *Geoscientific Model Development*, 14(5):3007–3036, 2021.
- [228] Timm Faulwasser, Robin Nydestedt, Christopher M. Kellett, and Steven R. Weller. Towards a FAIR-DICE IAM: Combining DICE and FAIR models. *IFAC-PapersOnLine*, 51(5):126–131, 2018. 1st IFAC Workshop on Integrated Assessment Modelling for Environmental Systems IAMES 2018.
- [229] Martin C. Hänsel, Moritz A. Drupp, Daniel J. A. Johansson, Frikk Nesje, Christian Azar, Mark C. Freeman, Ben Groom, and Thomas Sterner. Climate economics support for the UN climate targets. *Nature Climate Change*, 10(8):781–789, 2020.
- [230] Simon Dietz, Frederick van der Ploeg, Armon Rezai, and Frank Venmans. Are economists getting climate dynamics right and does it matter? *Journal of the Association of Environmental and Resource Economists*, 8(5):895–921, 2021.

- [231] Wilfried Rickels and Jörg Schwinger. Implications of temperature overshoot dynamics for climate and carbon dioxide removal policies in the DICE model. *Environmental Research Letters*, 16(10):104042, 2021.
- [232] O. Geoffroy, D. Saint-Martin, D. J. L. Olivié, A. Voldoire, G. Bellon, and S. Tytéca. Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments. *Journal of Climate*, 26(6):1841 – 1857, 2013.
- [233] Angelo Carlino, Matteo Giuliani, Massimo Tavoni, and Andrea Castelletti. Multi-objective optimal control of a simple stochastic climate-economy model. *IFAC-PapersOnLine*, 53(2):16593–16598, 2020. 21st IFAC World Congress.
- [234] Christopher A. Sims. Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163, 2006.
- [235] Kevin Lu, Igor Mordatch, and Pieter Abbeel. Adaptive online planning for continual lifelong learning, 2019.
- [236] Colin P. Morice, John J. Kennedy, Nick A. Rayner, and Phil D. Jones. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The hadcrut4 data set. *Journal of Geophysical Research: Atmospheres*, 117(D8), 2012.
- [237] Michael P. Vitus. *Stochastic control via chance constrained optimization and its application to unmanned aerial vehicles*. PhD thesis, Stanford University, 2012.
- [238] Ashkan Jasour, Allen Wang, and Brian C. Williams. Moment-based exact uncertainty propagation through nonlinear stochastic autonomous systems, 2021.
- [239] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT Press, Cambridge, Massachusetts, 2005.
- [240] Lukas Hewing and Melanie N. Zeilinger. Stochastic model predictive control for linear systems using probabilistic reachable sets. *2018 IEEE Conference on Decision and Control (CDC)*, 2018.
- [241] Lars Blackmore and Masahiro Ono. Convex chance constrained predictive control without sampling. *AIAA Guidance, Navigation, and Control Conference*, 2009.
- [242] Lars Blackmore, Masahiro Ono, and Brian C. Williams. Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics*, 27(6):1080–1094, 2011.
- [243] J.M. Grosso, C. Ocampo-Martínez, V. Puig, and B. Joseph. Chance-constrained model predictive control for drinking water networks. *Journal of Process Control*, 24(5):504–516, 2014.

- [244] Kyri Baker and Bridget Toomey. Efficient relaxations for joint chance constrained AC optimal power flow. *Electric Power Systems Research*, 148:230–236, 2017.
- [245] Masahiro Ono and Brian C. Williams. Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint. In *2008 47th IEEE Conference on Decision and Control*, pages 3427–3432, 2008.
- [246] Simon Dietz and Frank Venmans. Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management*, 96:108–129, 2019.
- [247] Yongyang Cai and Kenneth Judd. A simple but powerful simulated certainty equivalent approximation method for dynamic stochastic problems. Working Paper 28502, National Bureau of Economic Research, 2021.
- [248] Martin L. Weitzman. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, 5(2):275–292, 2011.
- [249] Martin L Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1):1–19, 2009.
- [250] Ali Mesbah, Stefan Streif, Rolf Findeisen, and Richard D. Braatz. Stochastic nonlinear model predictive control with probabilistic constraints. In *2014 American Control Conference*, pages 2413–2419, 2014.
- [251] Robin Nydestedt. Application of polynomial chaos expansion for climate economy assessment, 2018.
- [252] Alena Miftakhova. Global sensitivity analysis for optimal climate policies: Finding what truly matters. *Economic Modelling*, 105:105653, 2021.

Curriculum Vitae

Born on July 20, 1990 in Helsinki, Finland.

2018 – 2022 D.Sc. in Economics, ETH Zurich

2015 – 2017 M.Sc. in Environmental Economics, University of Helsinki

2011 – 2015 B.Sc. in Environmental Economics, University of Helsinki