Estimating the Climate–Economy Relationship in Europe for Over a Century

A Model Re-analysis and Extension Using an Extended Time Frame

Thesis

Sterre Maingay

Delft University of Technology

Estimating the Climate–Economy Relationship in Europe for Over a Century

A Model Re-analysis and Extension Using an Extended Time Frame

by

Sterre Maingay

Chair:	Prof. Dr. T. Filatova
Advisor:	Dr. T. Chatzivasileiadis
Supervisor:	Dr. E. Schröder
Project Duration:	February, 2025 - August, 2025
Faculty:	Faculty of Technology, Policy and Management, Delft

Cover: Own photo
Style: TU Delft Report Style

Executive Summary

This thesis estimates the economic effects of temperature, precipitation, and relative sea level rise (SLR) on regional economic growth in Europe from 1900 to 2015. A central starting point is the model developed by Burke, Hsiang, et al. (2015) (BHM), which provided the first global evidence that aggregate economic growth responds non-linearly to temperature, with growth peaking at an annual average temperature of 13°C and declining sharply at higher temperatures. Rising temperatures can reduce productivity and agricultural yields, especially in warmer regions (Somanathan et al., 2021). Shifts in precipitation patterns can cause longer dry periods or heavier rainfall, increasing the risk of floods and stressing water systems and infrastructure (Kotz et al., 2022; Malhi et al., 2021). SLR increases the risk of flooding and loss of land in low–lying coastal areas (Cortés Arbués et al., 2024; Chatzivasileiadis et al., 2023). This thesis addresses key gaps in the climate econometric literature by re–analysing and extending the BHM model with a regional dataset and additional climate variables for over a century. In doing so, it moves beyond theory–based models by relying on observed historical data to estimate the economic effects of climate variables with greater spatial and temporal detail.

This research uses a quantitative approach, applying an econometric model to regional (NUTS–2 level) climate from the Climatic Research Unit, University of East Anglia (n.d.) and economic data by Rosés et al. (2021). This thesis performs a re-analysis and extension, as posted by Clemens (2017), of the model by BHM to examine whether its findings remain valid. It tests the stability of the concave relationship between temperature and economic growth while also including precipitation and SLR to capture the combined effects of multiple climate variables.

The results from this research support the overall concave shape of the relation between temperature and economic growth. Meaning that even over a longer time frame, at the regional level, and with the inclusion of an additional climate variable, temperature and economic growth are linked in a non–linear way, with growth peaking at an optimal temperature of 11.8 °C. Temperature has the strongest effect on economic growth among the climate variables included in the model. Precipitation does not show a statistically significant effect on its own, but it is jointly significant when included with other climate variables. SLR has a smaller (compared to temperature) but statistically significant impact, suggesting that it plays a meaningful role alongside temperature in shaping economic outcomes.

However, the estimated coefficients for temperature in this research are approximately six times larger than those in the model by BHM, indicating stronger temperature sensitivity at the regional level. The sensitivity analysis further shows that these coefficients are not stable, particularly across different benchmark years, suggesting that the estimated turning point is shaped by historical context and should not be interpreted as a fixed economic threshold.

Future research should build on this work by using micro-level climate and economic data to examine the effects of seasonal variation, heatwaves, extreme rainfall, and other short-term events (Kotz et al., 2024; Somanathan et al., 2021). This could be combined with models that include adaptation processes or delayed responses to climate change (Mérel et al., 2021). The use of only twelve benchmark years in this study limits the ability to identify such short-run effects. Increasing temporal resolution would improve understanding of both short-term shocks and long-term structural patterns.

It will also be essential to better understand how climate variables interact. Temperature is a known driver of both precipitation and SLR, through its influence on atmospheric moisture and ice melt (Malhi et al., 2021). These dependencies should be explicitly modelled in future studies to avoid misattributing indirect effects or underestimating compound risks (Kotz et al., 2022). Finally, moving beyond the reduced-form approach by BHM would allow for more flexible models that capture causal mechanisms and account for the possibility of multiple optima, thresholds, or plateaus that a fixed quadratic form may overlook. These steps are essential to ensure that econometric models can support meaningful and context-specific design of climate policies.

Acknowledgements

This thesis has been both an academic and personal process. I chose this topic because I found it interesting, and it sparked my curiosity. Climate and its connection to society is something I care about. Many people of my generation worry about how future climate impacts will shape our lives. I wanted to explore whether this concern could be made more concrete through the use of data and models, and whether I could do this. This brought me into a field I was not familiar with, and at first, I struggled. Over time, I learned how to work through unfamiliar material and started to enjoy the process. Therefore, this thesis has taught me, next to technical knowledge, also much more.

I am grateful for my friends, family and fellow students who have always stood by my side on this journey. They encouraged me to really internalise this last eventful period of my study. I treasure their trust in me, which gave me the space to grow, to take ownership of my goals, and to complete something I had committed to.

I want to thank my thesis committee. First and foremost, Dr. Theodorus Chatzivasileiadis, thank you for the guidance and support you have given me throughout this project. I appreciate how sincerely you wanted to help me move forward whilst I ventured into unknown territory. Because of this, I was able to explore and find my way through this process. To Prof. Dr. Tatiana Filatova and Dr. Enno Schröder, thank you for your thoughtful feedback and rigorous comments on my work. Your suggestions helped me strengthen not only the content but also the structure of this thesis. This project, and the guidance you all have given me throughout, has taught me more than I could have anticipated. I will carry these teachings with me as I begin a new chapter of my life.

I also want to thank the TU Delft, specifically the EPA programme and its teaching staff, for offering a stimulating academic environment over the past two years. The programme challenged me to approach complex societal problems from different angles and gave me the tools to do so.

Finally, I want to acknowledge the work of Burke, Hsiang, et al. (2015), Rosés et al. (2021), and Climatic Research Unit, University of East Anglia (n.d.). Without their research, this thesis could not have been written.

Sterre Maingay
Delft, July 2025

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
SLR	Sea Level Rise
BHM	Burke et al., 2015
FE	Fixed Effects
GDP	Gross Domestic Product
EPA	Engineering and Policy Analysis

1
Introduction

A growing body of evidence shows that climate change can influence the functioning of modern human societies (Burke, Hsiang, et al., 2015). Yet, this relationship between climate and economy is complex and difficult to comprehend, let alone study and model (Tol, 2021). Changes in temperature, precipitation, and sea level rise (SLR) have the potential to shape economic outcomes, both separately and combined (Cortés Arbués et al., 2024; Chang et al., 2023). Temperature, precipitation, and SLR are among the main variables through which climate change affects the economy (Somanathan et al., 2021; Chen et al., 2019). Changes in these variables can influence how and where economic activity takes place, with consequences for economic growth (Tol, 2018). The paper by Burke, Hsiang, et al. (2015) (hereafter referred to as BHM) provided the first global evidence that aggregate economic growth responds non-linearly to temperature, with growth peaking at an annual average temperature of 13°C and declining sharply at higher temperatures. Their empirical model marked a shift away from theory–based integrated assessment models by estimating climate impacts directly from historical economic data. Integrated assessment models had until then relied on stylised (meaning simplified) damage functions that were structured based on economic theory and calibrated using rough estimates or expert judgment, rather than empirical data (Burke, Hsiang, et al., 2015). These models often involved strong assumptions and simplifications, and were acknowledged to have substantial limitations. Early research on temperature–GDP relationships, for instance, had little empirical basis for its response functions and empirical validation of these functions had remained scarce (Chang et al., 2023).

This research builds upon the work done by BHM and significantly extends the time frame of analysis from 1960–2010 (in the analysis by BHM) to 1900–2015, enabling a deeper investigation into long-run patterns of climate–growth effects over a century. By applying the model to regional–level data across Europe, this study provides new empirical insight into a spatial and temporal context that had not yet been covered at this scale. The use of a regional panel allows for the identification of local economic responses that are can be masked in national aggregates (Rosés et al., 2021). The analysis also includes SLR as an additional climate variable, which expands the model beyond the focus of the model by BHM on temperature and precipitation. This makes the results more comprehensive in describing the mechanisms through which climate conditions influence economic performance (Kotz et al., 2024). Ultimately, this long-run, regional analysis enhances the understanding of the effects of temperature, precipitation and SLR on economic growth in Europe over the past century, using observed historical data.

One of the symptoms of climate change is the rise in temperature (Burke, Hsiang, et al., 2015). Higher temperatures can reduce productivity among workers, called heat stress, which affects all sectors of an economy (Carleton et al., 2016). Higher temperatures can reduce agricultural output due to changes in growing conditions, leading to unsuccessful harvests. It can also reduce industrial output due to heat stress on machinery and industrial capital, lowering capital productivity (Zhang et al., 2018). Kalkuhl et al. (2020) find that an increase in mean global temperature of 3.5°C by 2100 could reduce total global output by 7% to 14%. BHM analysed the economic effects of temperature over the period 1980 to 2010. They estimate that, under a scenario of unmitigated warming with a projected global temperature increase of 3.7 °C by 2100, average global incomes could decline by more than 20%. BHM state that this would result in a 23% reduction in global GDP by 2100 compared to a scenario without further warming. Kahn et al. (2021) project that a continuous increase in temperature of 0.004°C per year would lead to a 7.22% decline in global GDP by 2100. They further estimate that if the temperature rise is limited to 0.001°C per year, in line with the Paris Agreement, the resulting GDP loss would be much smaller, only 1.07%. These findings show that projected economic outcomes depend on different scenarios of temperature increases. While the estimates of the different studies differ in magnitude, they consistently show that increases in temperature lead to with long–term losses in global economic output.

Another symptom of climate change is the change in precipitation patterns (Kotz et al., 2022; Malhi et al., 2021). Precipitation includes all forms of water that fall from the sky, such as rain, snow, or hail. The most damages to the economy come from extreme rainfall situations, such as both too little rain (droughts) and too much (floods), which are most damaging to the economy (Kotz et al., 2022). Such extreme rainfall situations reduce agricultural outputs and put pressure on infrastructure (Kotz et al., 2022; Malhi et al., 2021). It is important to note that only recently, in studies using data with a higher spatial resolution (firm-level data over national aggregate data) and smaller time frames (daily precipitation over annual precipitation), researchers have found significant impacts of precipitation on economic outcomes (Kotz et al., 2022). Historically, most macroeconomic research into the effects of annual averages of precipitation on a national scale did not find a significant impact of precipitation on economic growth (Burke, Hsiang, et al., 2015; Chen et al., 2019; Dell et al., 2012). Precipitation–related damages are smaller than those caused by temperature increases but remain important, with an estimated global cost of around 1.2 trillion USD by 2049 (Kotz et al., 2024). This damage estimation arises from three components: total annual precipitation, the number of wet days, and extreme daily rainfall. Total annual precipitation contributes a change of around 0.01 percentage points, the number of wet days adds about 0.34 percentage points (0.07 to 0.90), and extreme daily rainfall contributes around 0.36 percentage points (0.13 to 0.65) to the reduction in economic growth. This shows that not only the timing and intensity of rainfall events, but also the frequency of wet days, are better estimates for estimating economic damages than total annual precipitation.

Next to the economic damages caused by precipitation alone, temperature and precipitation often have a natural negative correlation, changes in one can be linked to changes in the other (Burke and Emerick, 2016; Blanc et al., 2017). Fro example, drier years tend to be hotter, and hotter years tend to be drier (Blanc et al., 2017). Therefore, if a model only includes temperature, it can wrongly attribute the effect of both temperature and precipitation to temperature alone (Blanc et al., 2017). In such a model, this can lead to inaccurate parameter estimates, which in turn will overestimate the damage of temperature on the economy (Blanc et al., 2017).

Lastly, the third symptom of climate change that will be discussed in this research is SLR (Nováková et al., 2018). Due to SLR, the water levels around the world's coasts are rising (Nicholls et al., 2021). As a result, the frequency of flooding is expected to double within this century, with floods that currently occur once every 100 years projected to happen as often as once every 10 years by 2100, primarily due to SLR (Kirezci et al., 2020). This is important since SLR is a climate-driven symptom which causes, mostly through flooding, direct damages to physical capital stocks, with these impacts leading to significant economic consequences at regional and sectoral levels (Cortés Arbués et al., 2024). Cortés Arbués et al. (2024) compared a scenario with SLR impacts (direct and indirect) to a baseline scenario with no SLR impact and found a total loss of 1.26% of GDP for the whole of Europe and the United Kingdom, equalling a total of 872 billion Euro by 2100. This loss is felt much more in coastal regions, which could lose up to 6.3% and 20.8% of regional GDP by 2100. These losses are calculated relative to a scenario in which regional economies grow at a constant annual exogenous growth rate of 2%. This finding underscores the uneven distribution of SLR risks to regional economic growth between coastal and non–coastal regions in Europe and the United Kingdom (Cortés Arbués et al., 2024). Recent work by Chatzivasileiadis et al. (2023) shows that, historically, SLR has caused a cumulative GDP loss of 4.7% in European coastal regions. This will have a long–term impact on annual GDP growth ranging from -0.02% to 0.04% per year (Chatzivasileiadis et al., 2023). Chatzivasileiadis et al. (2023) mention that in previous research by the OECD (2019), global coastal flood damage could amount to approximately 4% of global GDP, or around USD 50 trillion by 2100. This damage estimate is from a higher–end scenario in which SLR is 1.3m and compared to a scenario without further SLR. The paper by Aral et al. (2016) states that SLR is positively correlated with temperature over time. The estimated SLR by 2100 in their study ranges from 60 to 132 centimetres, depending on different scenarios, with corresponding global surface temperature increases between approximately 1.8°C and 4.5°C. These projections reflect a strong but delayed effect of an increase in temperature on SLR. Compared to temperature, the damage from SLR is smaller in terms of the reduction in economic output, but more concentrated in specific coastal regions (Chatzivasileiadis et al., 2023; Cortés Arbués et al., 2024). Overall, these findings show that SLR can cause serious economic damage, especially in coastal areas.

1.1. Modelling the climate-economy relationship

In recent years, researchers in the field of climate econometrics have developed statistical methods to examine how climate influences economic performance (Burke, Hsiang, et al., 2015). These methods are used to make climate econometric models. To use these model and meaningfully interpret the outputs, researchers make a distinction between weather, which describes short-term conditions, and climate, which refers to long-term patterns of these conditions (Blanc et al., 2017). Weather events (such as a heatwave or heavy rainfall) can be seen as specific instances within the broader climate distribution (Blanc et al., 2017). While economic agents (such as farms, firms, institutions, governments, and individuals) cannot plan around the variability of weather, they can and do adapt to climate, which forms the basis for long-term decisions (Mérel et al., 2021). However, adaptation is often slow and takes place over longer periods (Mérel et al., 2021). Temperature, precipitation, and SLR are among the main variables through which climate change affects the economy (Burke, Hsiang, et al., 2015; Somanathan et al., 2021; Chen et al., 2019). Changes in these variables can influence how and where economic activity takes place, with consequences for economic growth (Tol, 2018).

Existing studies have provided insights into the relationship between temperature, precipitation, SLR and economic performance (Dell et al., 2012; Cortés Arbués et al., 2024). However, the study by BHM and other studies alike typically focus on shorter periods of 50 years (Burke, Hsiang, et al., 2015; Dell et al., 2012). These studies also mostly use country-level data and annual averages of climate and economic variables (Kotz et al., 2022; Burke, Hsiang, et al., 2015). This limits the ability of the climate econometric models to capture causal relationships (Uhlig, 2012; Burke and Emerick, 2016; Mérel et al., 2021). Because climate is a long-term process, short datasets cannot detect these slow shifts and may instead pick up effects driven by weather fluctuations. Economic agents can adapt to persistent changes in climate but not to variable weather, so adaptation may appear over the long–run even when it is invisible in the outcomes of short–term models. Modelling the relation between climate and the economy on the country–level causes the differences between regions to be aggregated, revealing an oversimplified relationship (Rosés et al., 2021; Kotz et al., 2022). Modelling this relation at the regional level provides a more detailed view of where and how climate affects economic performance (Kotz et al., 2022).

The explored literature tends to isolate temperature, without accounting for the combined effects of other environmental variables on the economy, like precipitation and/or SLR. This is problematic because it may lead to inaccurate estimates of climate change impacts, particularly in settings where temperature, precipitation and SLR all influence the economy (Kotz et al., 2022). As Hsiang (2016) point out, capturing such interactions requires careful attention to both spatial and temporal variation, so to where and when. These dimensions are essential to identifying the mechanisms through which temperature, precipitation and SLR influence economic outcomes (Hsiang, 2016).

1.2. Research approach

This research aims to contribute to the understanding of how long-term shifts in climate conditions influence economic outcomes across time and space. Existing work has shown that temperature, precipitation, and SLR can affect economic outcomes, but besides the fact that most of this literature is based on relatively short periods and/or national averages, it often focuses on an isolated climate factor, such as temperature (Kotz et al., 2022). As a result, patterns such as variation across regions may go undetected or be inaccurate or incomplete (Rosés et al., 2021; Chatzivasileiadis et al., 2023). But also, variations in temperature tend to be highly correlated with variations in other climate variables (Blanc et al., 2017). If an econometric model soley investigates the effect of temperature on economic outcomes, it will wrongfully attribute the effect of the other climate variables to temperature (Blanc et al., 2017). To reduce such bias, this research, in line with the work by BHM, controls for unobserved regional differences using region–specific fixed effects (FE). Year FE capture sudden global events such as economic crises or wars. Flexible region–specific time trends control for the slow changes in regional growth patterns, such as shifts in population, demography or policy. These trends allow growth rates to evolve over time in a non-linear way and help to avoid wrongly attributing long–run economic changes to climate variables.

To address this, first, this research will examine whether the estimated relationship between climate and economic growth, as posted by BHM, holds when extended to a period of 115 years, drawing on the regional historical GDP dataset developed by Rosés et al. (2021) and climate data by the Climatic Research Unit, University of East Anglia (n.d.). This dataset allows for long-run analysis of economic growth at the regional (NUTS-2) level. To this end, the statistical model by BHM is used to see if the estimated relationship between temperature and economic outcomes holds when extending the time frame from 1960–2010, to 1900–2015. Second, this research investigates how the relationship between climate and economic growth varies across regions, as national averages can mask important regional differences and if, estimating the original model by BHM with regional data changes the temperature and economic growth relationship, as opposed to national–level data.

This research will combine temperature, precipitation, and SLR rather than looking at them in isolation. Including an additional climate variable, SLR, is important because climate variables often interact, and their compound effects may be different from their effects in isolation (Kotz et al., 2022). Given the complexity of these relationships, it is important to assess whether the empirical results are sensitive to modelling choices. If these compound effects are not captured, estimates of climate impacts risk being biased or misleading, which can cause misinformed and ineffective action against climate change, as well as misinterpretation of the scale and distribution of climate-related damages (Blanc et al., 2017; Mérel et al., 2021).

To this end, a re-analysis and an extension of the original model by BHM will be carried out, following the approach outlined by Clemens (2017). This is done to test how the findings respond to changes in model specification, dataset composition, and the inclusion of SLR. The model developed by BHM will serve as the analytical baseline. Testing its performance within an extended historical and regional framework helps to determine whether its conclusions hold under alternative assumptions.

1.3. The EPA relevance of this research

The relevance of this research to the master's program Engineering and Policy Analysis (EPA) lies in its alignment with addressing climate change, a global challenge and one of the 17 United Nations Sustainable Development Goals (SDGs), particularly SDG 13: Climate Action (United Nations, n.d.). Climate change is a wicked problem characterised by complex interdependencies to other variables, uncertainty, and the involvement of multiple stakeholders on all levels of society (Incropera, 2015). This research focuses on the intersection of climate and economy, using econometric methods, specifically a panel data regression model that combines regional economic data from the Rosés et al. (2021) dataset and the climatic dataset from the Climatic Research Unit Climatic Research Unit, University of East Anglia (n.d.). This model will explore how long-term climate variables influence economic outcomes across multiple regions in Europe. As Hsiang (2016) notes, climate policy requires a clear understanding of the full economic burden of climate change. Empirically estimated relationships, such as those in this research, can inform projections of future climate change damages. While methodological challenges remain, empirical climate-economy research offers valuable input for policy design. Therefore, this study aligns with EPAs emphasis on evidence-based policymaking by combining analytical tools with real–world policy relevance to create insights into the economic consequences of climate variables by analysing models. Furthermore, this research adopts a systems-thinking approach by integrating econometric models and socio-economic analysis. This is relevant for policy since it addresses the broader societal challenges, such as the uneven exposure to climate risks across time and space, which complicates efforts to understand long-term impacts but could inform effective regional action (Burke, Hsiang, et al., 2015; Mérel et al., 2021; Andersen et al., 2016; Dell et al., 2008; Rosés et al., 2021).

1.4. Main and sub research questions

From the literature review, discussed in Chapter 2, it is clear that estimating the economic impacts of climate change in Europe requires a regional economic analysis based on an extended time frame, including the variable SLR. This has led to the formulation of the following main research question:

To what extent are the estimated economic impacts of climate, identified in the model by BHM, robust?

To answer this main research question, three subquestions need to be answered first:

1. How robust is the model by BHM when extending the time frame to 115 years?
2. What are the compound effects of temperature, precipitation, and SLR on economic outcomes, and to what extent do these compounding effects alter the climate-economy relationship identified in the model by BHM?
3. Is the estimated climate effect of the model by BHM robust when tested for sensitivity to spatial and temporal exclusions?

Each of these subquestions addresses a knowledge gap as identified in Chapter 2.

1.5. Thesis structure

Chapter 2 provides the theoretical background. In this Chapter, the econometric literature is discussed, focusing on the relationship between temperature, precipitation and SLR and economic outcomes. Chapter 3 presents the climate-economy model by BHM. In this Chapter, the model by BHM will be explained and how the model captures the relationship between temperature and economic growth. Chapter 4, describes the data preparation process. In this Chapter, the data sources, the steps taken to clean and structure the data, and how the data are prepared for analysis using the climate-economy model are outlined. Chapter 5 presents the results of the analysis. In this Chapter, the estimated relationship between climate variables and economic performance and discusses the main findings, and gives answers to the first and second subquestion. Chapter 6 answers the third subquestion. In this Chapter, a spatial and temporal sensitivity analysis is carried out. Chapter 7 provides a discussion of the results. In this Chapter, the findings are discussed in relation to existing literature, and the limitations of this study are considered. Chapter 8 provides the conclusion of this study. In this Chapter, the main findings and reflecting on their implications for understanding the economic impacts of climate change are summarised.

Several appendices are at the end of this thesis as supporting material. Appendix A contains additional details on data cleaning and preparation. Appendix B offers technical details of the climate-economy model by BHM. Appendix C presents further details of the results. Appendix D shows the outcomes of using a 30-year rolling mean for temperature and precipitation in the climate-economy model.

2 Theoretical Background

Climate change is no longer a distant possibility but a measurable and accelerating global fact (Dell et al., 2008; Malhi et al., 2021; Rahimi et al., 2020). Average temperatures have increased, precipitation patterns have become less predictable, and sea levels continue to rise (Duchenne-Moutien et al., 2021). These developments are largely linked to human activity and are placing pressure on the economy (Magnan et al., 2021). Climate change is expected to come with substantial economic costs, particularly through its effects on productivity, agriculture, infrastructure, and eventually economic growth (Burke, Hsiang, et al., 2015; Tol, 2018). Understanding how temperature, precipitation, and SLR affect the economy is essential for assessing the impact of climate change. Without understanding how climate relates to the economy, it becomes difficult to anticipate climate damages, and how adaptation might mitigate these effects (Hsiang, 2016; Newell et al., 2021).

To understand how climate variables damage the economy, and how much the economic damage will be, an important distinction needs to be made. This is the distinction between climate and weather. Climate refers to the long-term patterns of these climate variables as opposed to weather, which describes short-term events (Blanc et al., 2017). Climate provides therefore the context in which economic agents make strategic decisions (Mérel et al., 2021). By contrast, weather consists of the day-to-day variations and individual events that are inherently unpredictable (Blanc et al., 2017). As Blanc et al. (2017) explains, each weather event is just one realisation drawn from the broader climate distribution. Since economic agents can, therefore, only adjust to climate, adaptation can only be a response to climate and not to transient weather fluctuations (Mérel et al., 2021; Blanc et al., 2017). A farmer will not change to a different farming method after one year of drought (Kolstad et al., 2020). For this, the investment is too big. The farmer will change the method only after prolonged drought (Kolstad et al., 2020).

Understanding the distinction between weather and climate is fundamental for formulating effective policy to address the impacts of climate change (Burke and Emerick, 2016). Effective policy depends not only on recognising the economic risks posed by a changing climate but also on quantifying those risks in ways that guide targeted responses. Without the clear distinction between climate and weather, projecting the future economic impacts of climate change becomes problematic (Chang et al., 2023; Hsiang, 2016). This research focuses on the relation between annual temperature, precipitation, and SLR and economic outcomes in benchmark years. Unlike temperature and precipitation, SLR is not measured annually. Instead, it reflects the difference in relative SLR between 1900 and the selected benchmark year. Because of this, the results of this study will reflect the impact of a combination of weather and climate events on regional GDP growth, in line with the work by BHM. BHM assume that their model implicitly assumes adaptation, meaning that the results reflect a combination of the impacts of weather variability and the longer-term climate conditions to which economies have already adapted.

2.1. The climate-economy relationship

But how does climate affect the economy? This question seeks not only to understand why economies evolve differently and how present economies manage climate change, but also how they might respond to future climatic changes and what the impact of this would be (Hsiang, 2016). But this link between climate and the economy is complex and influenced by multiple channels, including exposure, adaptive capacity, and institutional context (Hsiang, 2016). Yet, it is of great importance since climate change may alter the course of economic development (Burke, Hsiang, et al., 2015; Hsiang, 2016; Tol, 2018).

Given this potential, quantifying this relationship is of great importance for informing policy design (Kolstad et al., 2020; Chang et al., 2023; Kotz et al., 2024). A clearer understanding of how climate change impacts economic growth helps to evaluate the trade-offs between the costs of mitigating and the benefits of avoiding climate-related damages (Kotz et al., 2024; Chang et al., 2023).

But how can this link between climate and the economy be quantified? One field that tries to answer this question is the field of climate econometrics (Hsiang, 2016). Climate econometrics focuses on how climate change impacts economic outcomes by means of statistical models that use historical climate and economic data (Kolstad et al., 2020; Chang et al., 2023). These models often rely on panel data methods, which allow researchers to study the same areas over time and control for differences between areas that do not change, such as soil quality, institutional settings, or infrastructure, but also general trends over time experienced by all locations such as global recessions and wars (Kolstad et al., 2020; Chang et al., 2023). This approach allows for more valid identification of causal effects, as it controls for location-specific characteristics and global time trends (Kolstad et al., 2020). These models often include non-linear relationships, meaning they can capture the complex relationships between climate variables and economic activity (Burke, Hsiang, et al., 2015; Blanc et al., 2017; Mérel et al., 2021). Climate econometrics is a growing field of research (Chang et al., 2023). Econometric models are common practice in quantifying the relationship between climate and economy (Dell et al., 2014; Hsiang, 2016). These models help isolate the effects of climate from other variables that also affect the economy, but are not related to climate (Chang et al., 2023; Kotz et al., 2024). This is making it possible to detect patterns across time and space (Chang et al., 2023). In doing so, climate econometric models provide empirical evidence that can inform effective action and the design of future climate policies (Chen et al., 2019; Chang et al., 2023).

2.2. Evolution of the field

The field of climate economics has evolved rapidly over the past two decades, moving from theoretical models to data-driven ones (Hsiang, 2016). Early contributions focused on estimating the long-term economic costs of climate change through integrated assessment models (Stern, 2008), while more recent work has used historical weather and climate data to identify the effects on economic outcomes (Dell et al., 2012; Dell et al., 2014; Hsiang, 2016). The growing availability of high-resolution climate and economic data has made it possible to analyse observed relationships rather than relying only on projections or assumptions (Carleton et al., 2016; Burke, Hsiang, et al., 2015; Blanc et al., 2017). This has led to better informed insights about how specific climate variables, such as temperature and precipitation, affect growth and productivity across different contexts (Kotz et al., 2022). Two empirical strategies, cross-sectional and panel data models, are most often used. Cross-sectional analysis involves comparing different units of observation at a single point in time. A recognised weakness of cross-sectional analysis is that these models cannot include unobserved factors correlated with climate (these are factors that do not change over time, such as institutions and soil quality), which can cause omitted variable bias (Blanc et al., 2017; Carleton et al., 2016). Cross-sectional models are also static, which can lead to them over- or underestimating the damages of climate change (Chang et al., 2023; Kolstad et al., 2020).

That is why an increasing number of studies use panel data, which consist of repeated observations for the same units over time (Chang et al., 2023). This structure allows researchers to track changes within each unit and to control for unobserved characteristics that do not vary over time, also called fixed effects (FE). Panel data models, especially those that include FE, are widely applied to estimate the impact of climate and other environmental factors on economic outcomes (Mérel et al., 2021). Despite this progress, several challenges remain. One challenge is that establishing causality is difficult, as climate and economic conditions may be influenced by shared underlying factors (Dell et al., 2014; Kolstad et al., 2020). Another challenge is adaptation. Over time, economic agents adapt to changing climate conditions, which can make the effects of climate harder to observe (Mérel et al., 2021). In addition, weather shocks can obscure shifts in climate, so it is important to separate year-to-year weather anomalies from long-term trends (Blanc et al., 2017). Therefore, empirical research into the effects of climate change must distinguish between weather shocks and climate shifts to avoid inaccurate estimates of long-run damages (Carleton et al., 2016). Finally, many studies are limited by the quality of available data. Gaps in coverage, missing years, or data only at the national level can make it harder to study the effects of climate on the economy (Rosés et al., 2021; Kotz et al., 2022).

This chapter will now turn to look at the channels through which temperature, precipitation and SLR shape economic outcomes, as well as their compounding effects over a period of 115 years. The dataset by the Climatic Research Unit, University of East Anglia (n.d.) and by Rosés et al. (2021) can supply this extended time series at the regional, NUTS-2 level, for 170 European regions. Rosés et al. (2021) note that using only national averages of climatic and economic variables for Europe can conceal substantial within–country differences, sometimes larger than those between countries. Because of this statement, the importance of a regional analysis is investigated. To this end, an influential climate-econometric model, the model by BHM, is re-analysed and extended following the notation of Clemens (2017). This re-analysis and extension are designed to address the three research questions, each corresponding to a specific gap in the existing literature.

2.3. Channels of climate impact

An increasing number of studies show that climate variables such as temperature, precipitation, and SLR influence the economy through complex channels (Duchenne-Moutien et al., 2021). This section examines the channels through which each of these variables affects economic output.

2.3.1. Temperature and economic outcome

The earlier work by Dell et al. (2008) and Dell et al. (2009) contributed to establishing the link between temperature and economic growth. Building upon this work, BHM show a non-linear relationship between temperature and economic outcomes, where productivity increases up to a temperature threshold, globally around 13°C, before declining at higher temperatures (Burke, Hsiang, et al., 2015). BHM found that this pattern appears to be consistent across countries and sectors, including both agricultural and non-agricultural industries, and has remained stable since 1960. Though low-income, heat-exposed areas are disproportionately affected due to weaker infrastructure and limited adaptive capacity, such as limited resources (Acevedo et al., 2020; Xia et al., 2018; Burke, Hsiang, et al., 2015). The capacity to implement large–scale protective infrastructure, such as dikes against SLR, is directly linked to a country's GDP and wealth (Hinkel et al., 2014). These findings suggest that these areas face higher vulnerability to climate change. Importantly, the observed reductions in economic output come from lower growth rates rather than just temporary declines in output levels (Burke, Hsiang, et al., 2015). The channels through which higher temperatures reduce economic outcomes include reduced labour productivity due to heat stress and lower efficiency in sectors sensitive to temperature, especially in areas with limited access to cooling or adaptation measures (Acevedo et al., 2020; Xia et al., 2018). After the threshold temperature is reached (as mentioned defined by BHM at 13°C), output losses range from 2-3% per 1°C increase in heat-sensitive sectors like manufacturing and agriculture (Somanathan et al., 2021; Burke, Hsiang, et al., 2015). Heat exposure also leads to health impacts that reduce working time and long-term productivity. Furthermore, temperature increases can have delayed effects (Kahn et al., 2021). This means that high temperatures in one year may reduce output in subsequent years (Xia et al., 2018; Chen et al., 2019). Although some adaptation to temperature is possible, such as shifts in working hours, land- and technology-use, these responses are often slow and uneven across areas (Mérel et al., 2021). But temperature is not the only climate variable that shapes economic outcomes; precipitation also influences economies (Tol, 2021).

2.3.2. Precipitation and economic outcome

Precipitation is another climate variable examined to understand through which channels it affects economic outcomes. Recent research shows that changes in precipitation patterns, including both total rainfall and its distribution across time, influence economic growth (Kotz et al., 2022). While moderate rainfall variability can be absorbed without any disruption, both too much and too little precipitation tend to reduce economic output (Kotz et al., 2022). This relationship is non-linear and highly context-dependent (Kotz et al., 2022; Khan et al., 2022). For instance, increased rainfall can enhance agricultural output in some low-income countries, particularly where irrigation is limited. Precipitation shocks can also indirectly affect economic output through changes in land use (Malhi et al., 2021). In response to persistent drought, farmers may expand cultivated land to offset productivity losses, leading to long-term changes in land productivity and environmental degradation (Zaveri et al., 2020; Malhi et al., 2021). In more industrial settings, changes in precipitation can interrupt production cycles, degrade inputs, and limit labour efficiency. Additionally, Malhi et al. (2021) stress that precipitation extremes, such as heavy rainfall or droughts, are expected to intensify in the near future due to climate change, with area-specific effects on water availability, soil erosion, and crop productivity. These changes are particularly harmful to agriculture in developing countries, where the capacity to adapt is limited and the dependence on rain is great (Khan et al., 2022). The recent findings by BHM state that precipitation has a non-linear impact on economic growth. Their analysis includes a control variable for precipitation since changes in local annual temperatures tend to be correlated with changes in precipitation (Burke and Emerick, 2016). BHM reveal that economic growth is influenced by annual average rainfall, though the effect is weaker and less consistent than that of temperature. Importantly, their model accounts for precipitation as a non-linear (quadratic) term, capturing the possibility that both very low and very high levels of rainfall may impact growth. While temperature emerges as the climatic driver of economic performance, variations in rainfall still have measurable effects, particularly in regions dependent on agriculture or vulnerable to flooding (Cortés Arbués et al., 2024; Kotz et al., 2022). In that sense, precipitation can be used as an indicator for flooding since heavy precipitation is often linked to river floods and water–related damages (Kotz et al., 2022).

2.3.3. SLR and economic outcome

SLR represents a slow but important dimension of climate change, particularly for coastal regions (Chatzivasileiadis et al., 2023). As sea levels continue to rise, the risk of coastal flooding, land loss, and damage to infrastructure grows (Chatzivasileiadis et al., 2023). Unlike temperature and precipitation, which affect a wide range of areas more immediately, the effects of SLR are slower to develop but can be hard felt in specific locations (Nováková et al., 2018). Nováková et al. (2018) note that most empirical work has focused on temperature and precipitation, with limited direct evidence on the economic effects of SLR. Empirical studies on the relationship between SLR and economic growth remain varied. The retrospective study by Nováková et al. (2018), focused on the historical effects of SLR in the United States, reported no stable or statistically significant impact on economic growth. The authors propose that the effects of SLR may be relatively limited in advanced economies such as the United States, in contrast to potentially greater impacts in less developed regions. However, examining this hypothesis remains difficult, largely because of constraints in the availability and quality of comparable data (Nováková et al., 2018). From a methodological perspective, assessing the impact of SLR on overall economic output is difficult, as historical variation in SLR has been relatively modest. Nonetheless, several studies caution that excluding SLR from broader estimates of climate-related economic damages may lead to an underestimation of costs, especially in lower-income and coastal countries, where vulnerability is often greater (Kotz et al., 2024). Both the paper by Tol (2018) and Tol (2021) note that current models often do not capture the full range of potential damages of SLR that lie outside the current range of historical data. In his paper, he specifically states that unprecedented SLR is an example of such an event. This indicates that even extensive historical data may be insufficient to reliably project the full scale of future damages for phenomena like SLR, precisely because the projected future changes are different from anything observed historically. It is important to look at SLR over a long time series because of its inherent nature to develop over centuries and its cumulative impacts. This long-term perspective is also needed to inform policy design and suitable adaptation strategies (Kirezci et al., 2020; Cortés Arbués et al., 2024).

A study by Hinkel et al. (2014) shows that, without investment in coastal protection, SLR could lead to direct global annual flood costs amounting to over 5% of GDP by the end of the century. Their findings also indicate that even relatively modest rises in sea level, when combined with growing populations and industries in coastal areas, would increase the number of people and value of property and industry at risk. A related study by Kirezci et al. (2020) shows that the exposure of land, population, and economic output to coastal flooding is expected to increase throughout the century. Rahimi et al. (2020) researches the compounding effects of SLR and heavy coastal precipitation. Their study finds that when these factors occur together, they can overwhelm existing flood protection infrastructure and increase the damage from flooding. Such compound events have the potential to disrupt infrastructure, which may in turn increase economic vulnerability and reduce economic growth in affected regions. A more recent study by Chatzivasileiadis et al. (2023) found empirical evidence from European regions that shows that SLR has already reduced regional GDP growth in coastal regions in the period from

1900 to 2020. This study found that SLR increases the risk of coastal flooding and land loss, which creates economic damages and constrains economic activity in these areas. As Rahimi et al. (2020) state that SLR in combination with heavy precipitation can have compounding effects. The work by Khan et al. (2022) does not isolate the specific economic contribution of SLR from temperature and precipitation. Instead, they created a natural disaster variable that is used to analyse the effects of mean temperature, precipitation, sunshine duration, and mean sea level pressure on GDP. The results show that the natural disaster variable, which includes SLR, has a significant negative effect on economic growth.

Thus, a clear gap in the existing literature is the absence of studies that examine the compounding effect of temperature, precipitation, and SLR on macroeconomic outcomes such as GDP growth. While temperature and precipitation are often analysed together, and SLR is typically studied in isolation, mainly in relation to coastal damage. There is little evidence of models that capture their potential compounding effects. Given suggestions in the literature that SLR, particularly when combined with heavy rainfall, may amplify economic impacts, it is important to investigate whether such compounding effects influence aggregate economic performance.

2.4. A regional analysis on an extended time series

Recognising that climate represents the distribution of individual weather events on which economic agents base strategic choices, researchers seek to understand how these longer-term climatic patterns affect economic outcomes over time (Mérel et al., 2021; Hsiang, 2016). A central question in the field of climate economics is how quickly economic agents adjust to changes in their environment (Burke and Emerick, 2016; Mérel et al., 2021). Adaptation refers to actions or investments undertaken to reduce the influence of climate change (Hsiang, 2016). Economic agents make adaptations based on climate normals. Climate normals are defined as average weather conditions observed over a standard period of around 30 years, as defined by the World Meteorological Organisation (Burke, Hsiang, et al., 2015). To understand how these longer-term climatic patterns affect economic activity over time requires extended time series data that capture sufficient variation in climate over time (Chang et al., 2023). It is important to understand how economies respond to climate change over time, especially when considering how adaptation may meaningfully reduce future climate damages. Many studies on the effects of climate change use short-term changes in weather to estimate economic impacts. These studies can provide useful information, but they do not show how economies respond to long-term climate trends. If economic agents adapt slowly to new climate conditions, short-term responses may give a misleading picture (Kolstad et al., 2020). For instance, behavioural adjustments, infrastructure investment, or shifts in production methods often take time and may only be observable over longer periods. As a result, damage estimates based only on short time frames might either overstate or understate the eventual impact of long-term climate change damages. Simply extrapolating short-run effects into the future may overestimate the long-term consequences of climate change if meaningful adaptation takes place over time (Chang et al., 2023). Burke and Emerick (2016) show the importance of understanding climate change by examining the effects of long-term exposure to temperature and precipitation changes on U.S. agricultural productivity. Their findings suggest that exposure to extreme heat reduces crop yields, particularly for corn and soy, and that these effects remain, even over a 20-year period. They also reveal how repeated heat extremes can slow growth rates and lead to larger cumulative losses than one extreme weather shocks suggest. Distinguishing between temporary level effects (a single bad harvest due to a hot year) and lasting growth effects (lower yields from that year on) is only possible when data span many years (Chang et al., 2023; Burke, Hsiang, et al., 2015). The paper by Dell et al. (2009) shows that if adaptation is meaningful, it can reduce climate change damages of temperature shocks by 50%. Dell et al. (2009) discovered this after they compared a cross-sectional and panel data model. In this way, they compared a model that examines variation across many units at a single point in time to a model that follows the same units over several periods. By using extended time series to study the relationship between climate variables and economic outcomes, the potential effects of adaptation could be better understood.

The literature also notes the importance of carrying out analyses at the regional level (Khan et al., 2022; Kotz et al., 2024). Examining the relationship between climatic and economic variables at the regional, NUTS-2 level offers several advantages, particularly to capture regional heterogeneity in both exposure and adaptation to climate variability (Rosés et al., 2021). Regional level data allows for the investigation of place-specific adaptation, which is influenced by local governance, infrastructure, and sectoral composition (Kalkuhl et al., 2020; Rosés et al., 2021). This spatial resolution is especially important in Europe, where climate exposure, economic structure, and institutional capacity vary greatly across regions (Rosés et al., 2021). The paper by Rosés et al. (2021) explains that Europe is fundamentally different from large countries with relatively centralised economies such as the United States because it consists of a mosaic of small, historically distinct national economies, each with its own institutions, economic structures, and development trajectories. They therefore argue that treating Europe as a collection of national units often obscures substantial within-country (regional) variation, which is sometimes larger than the differences observed across countries (Rosés et al., 2021).

But here is a gap in the climate econometrics literature that lies in the limited use of long-term, multi-regional analyses at the NUTS-2 level (Kalkuhl et al., 2020; Cortés Arbués et al., 2024; Diffenbaugh et al., 2019). Most existing studies focus on short-term impacts and/or country-specific data, providing valuable but incomplete insights into the economic consequences of climate change (Kotz et al., 2024). As BHM emphasise, the relationship between climate and economic output is highly non-linear, implying that both moderate and extreme climatic changes must be studied over longer horizons to understand the trajectory of economic impacts. Without extended time series, it is difficult to assess how economies cope with or adapt to or fail to adapt to long-term climate changes. The absence of long-term analyses also slows down the ability to account for the cumulative effects of climatic variables like temperature, precipitation and SLR on economic outcomes, which unfold over decades, influencing adaptive capacities, technological advancements, and policy responses (Chang et al., 2023; Kahn et al., 2021; Mérel et al., 2021). As Stern (2008) points out, climate change is a cumulative phenomenon with long lags between climate change causes and the visible effects on the economy. Furthermore, economic sensitivity to climate is not static. It evolves over time as economies undergo structural transformations such as industrialisation, urbanisation, and sectoral shifts (Chang et al., 2023). These shifts can vary greatly over space, and thus the literature underscores the need for detailed regional disaggregation (Khan et al., 2022). Capturing these dynamics requires extended time series with sufficient historical depth and detailed regional data.

2.5. Summary of knowledge gaps

So altogether, based on the climate econometric literature, three research gaps could be identified. Firstly, many studies rely on datasets covering only a relatively short period of time (Diffenbaugh et al., 2019; Kalkuhl et al., 2020; Cortés Arbués et al., 2024). Without time series covering a longer period of time, it is difficult to assess how economies cope with long–term climate changes (Chang et al., 2023; Kahn et al., 2021; Mérel et al., 2021). Secondly, most studies rely on single–country or one–region–specific data to estimate damage variations (Cortés Arbués et al., 2024). But climate impacts are highly heterogeneous, varying across geographic, economic, and institutional contexts. Using data at the regional, NUTS-2, level is necessary to capture regional heterogeneity, which can vary greatly across regions, especially in Europe (Rosés et al., 2021). Lastly, the compounding effects of temperature, precipitation and SLR are rarely studied together. Although temperature and rainfall are often analysed together, SLR is not included (Chang et al., 2023). This gap is especially important for regions vulnerable to frequent or overlapping climate risks.

This research aims to contribute to closing these gaps by evaluating the sensitivity of the model by BHM by conducting a robustness test, as posted by Clemens (2017). A robustness test can have two alternative forms depending on the usage of the same or different data compared to the original study and/or altering the model specification. For this research, the model specification by BHM will be altered, and different data will be used as input. This means that this thesis performs both a re–analysis and an extension test (Clemens, 2017). Note that rebuilding and altering the model specification of the original model does not mean that the original model is being replicated. A replication test, according to Clemens (2017), is a test that evaluates whether an original studys findings can be independently reproduced using the same methodology, so without altering the programming code of the original study. A replication test ensures that the statistical analysis yields the same results as the original study by either reanalysing the same dataset or resampling from the same population. Since this is not the intention of this research, identical results are not expected from the robustness test as they did in the original study by BHM. This research examines whether the estimated relationships from the original study by BHM are robust to modifications to model specification and data. Reanalysis and extension tests can provide useful insights into the stability and potential generalisability of the findings (Clemens, 2017). Through a reanalysis and extension of the model by BHM, applied to climatic and economic data from 170 NUTS-2 regions in Europe between 1900 and 2015 provided by the Climatic Research Unit, University of East Anglia (n.d.) and Rosés et al. (2021) datasets, this study investigates whether the climate-related economic effects identified in the original work persist across different regional contexts and over an extended time frame. If the results are in line with the original findings, this would increase the confidence in the robustness of the estimated relationships. Large differences could suggest that responses to climate shocks are context-specific, indicating that the conclusions of the original studies may not be directly transferable to other geographic and temporal contexts (Clemens, 2017).

The identification of the knowledge gaps in the existing climate econometrics literature, as well as the formulation of the intention of this research, has led to the formulation of the main research question and sub-questions as mentioned in Chapter 1, see section 1.4.

3
The Climate-Economy Model

The model developed by BHM investigates how temperature influences economic productivity at the global level, by aggregating country–level data. It focuses on estimating the effect of annual average temperature on the growth rate of GDP per capita. Rather than modelling separate economic mechanisms such as crop yields or labour productivity, the model uses a reduced-form approach to capture the aggregate effect of temperature fluctuations. The starting point of this model is that the relationship between temperature and economic growth is non-linear, meaning that both cold and hot extremes may reduce productivity, while moderate temperatures are more favourable. This approach helps to align previous findings in the literature, where micro–level data (such as firm-level or plant-level data) showed strong temperature effects, while macro-level studies reported limited or linear impacts. By applying a unified framework to global panel data, the model by BHM identifies a consistent, concave relationship between annual temperature and economic output, with peak productivity occurring at approximately 13 °C globally.

The objective of the model by BHM is to use a panel dataset to estimate a global non-linear relationship between annual average temperature and per–capita economic growth. The model tests whether temperature has a statistically significant influence on growth after controlling for time-invariant country characteristics, global shocks, and country-specific time trends. By specifying a flexible quadratic functional form, as will be shown in Equation 3.9, the model identifies whether this relationship is concave, indicating an optimal temperature beyond which additional warming reduces growth. Instead of assuming that each additional degree of warming has the same impact everywhere, the model allows the marginal effect to vary depending on the starting temperature level, which, in their model, is country–specific. This design makes it possible to detect threshold effects and to estimate how growth responses differ between countries with colder or warmer baseline climates. In doing so, the model provides insight into which countries may benefit from moderate warming and which are more vulnerable to heat–related economic losses.

This chapter provides a detailed overview of the BHM model. It outlines the theoretical motivation, data structure, estimation strategy, and main empirical results. The aim is to clarify how the model links variation in climate variables to economic growth and why it serves as a suitable reference for the empirical strategy developed in this thesis. The following sections explain the production logic, in section 3.1, the regression structure in subsection 3.2, and insights into fitted temperature-growth relationship in section 3.4, providing a foundation for applying a similar approach to the regional European data introduced in the next Chapter, Chapter 4.

3.1. Derivation of the empirical regression model

The framework of the model by BHM begins by focusing on a small part of the economy at a given location and moment. This approach typically uses a mathematical model to describe how output is generated. A common tool for this is a Cobb-Douglas production function (Tol, 2021), as applied in the original model by BHM. Consider a small economic unit, like a farm, a factory, or a small region, located at a specific place ℓ during a certain time t and operating within a particular industry i (Burke, Hsiang, et al., 2015). Total output $Y_{i\ell t}$ of the economy is often modelled as a function of its productive resources (such as labour and capital) and the environmental conditions (such as temperature and precipitation) it experiences and given by the following Equation 3.1 in industry i, location ℓ and time t:

$$Y_{i\ell t}(T_{\ell t}) \;=\; p_i\big(A_i^K(T_{\ell t})\,K_{i\ell t}(T_{\ell t})\big)^{\alpha}\,\big(A_i^L(T_{\ell t})\,L_{i\ell t}(T_{\ell t})\big)^{1-\alpha} \tag{3.1}$$

Here $A_i^K(T)$ and $A_i^L(T)$ are the productivities of capital and labour that depend on instantaneous temperature T. $K_{i\ell t}$ and $L_{i\ell t}$ are the quantities of capital and labour employed at location ℓ. p_i is the price of a unit of output in industry i. In a competitive market, firms choose capital and labour so that the ratio $K_{i\ell t}/L_{i\ell t} = \alpha/(1-\alpha)$, in this way the function has constant returns to scale. Because of the constant returns to scale, any proportional change in both inputs leads to the same proportional change in output. For example, doubling capital and labour together doubles output, and halving both cuts output in half. To simplify notation and bundle prices with production, BHM then define:

$$U_{i\ell t} \;=\; p_i\, K_{i\ell t}^{\alpha}\, L_{i\ell t}^{1-\alpha},$$

Which represents the monetary value of total resources allocated to industry i at location ℓ at time t. Substituting into Equation 3.1 and using $K/L = \alpha/(1-\alpha)$, the Equation 3.1 of total output simplifies to:

$$Y_{i\ell t}\big(T_{\ell t}\big) \;=\; \underbrace{\big(A_i^K(T_{\ell t})\big)^{\alpha}\,\big(A_i^L(T_{\ell t})\big)^{1-\alpha}}_{f_i\left(T_{\ell t}\right)}\; p_i\; K_{i\ell t}^{\alpha}\; L_{i\ell t}^{1-\alpha} \;=\; f_i\big(T_{\ell t}\big)\, U_{i\ell t} \tag{3.2}$$

This is the unitlevel productivity function and describes how temperature directly influences productivity at the unit-level. Under constant returns to scale, all temperature effects on capital and labour combine into this single function $f_i(T)$. Equation 3.2 shows that instantaneous temperature affects output through $f_i(T)$, while $U_{i\ell t}$ captures the monetary value of the total resources allocated to production at location ℓ. Together, they allow temperature-driven productivity changes to be isolated from input levels, providing a clear link between temperature and economic output.

3.1.1. Production as a function of average temperature

Aggregating micro-level outputs over all locations ℓ and time t within a year τ, and summing across all industries i, yields:

$$Y_{L\tau} \;=\; \sum_i \int_{t\in\tau} \int_{\ell\in L} f_i\big(T_{\ell t}\big)\, U_{i\ell t}\; d\ell\, dt$$

BHM define the annual average temperature in country L during year τ as $\overline{T}_{L\tau}$. Let $g_i(T-\overline{T}_{L\tau})$ denote the distribution of instantaneous temperature (temperature at a specific moment) deviations around $\overline{T}_{L\tau}$. Under the assumptions that industries and locations contribute additively (the output produced by each industry at each location simply adds up to the total output of a country), each distribution $g_i(\cdot)$ retains its shape (shifting only with its mean), and capital and labour do not relocate quickly in response to temperature changes, total output can be written as:

$$Y_{L\tau}\big(\overline{T}_{L\tau}\big) \;=\; \sum_i \int_{-\infty}^{+\infty} f_i(T)\; g_i\big(T-\overline{T}_{L\tau}\big)\; dT \tag{3.3}$$

Here $f_i(T)$ represents the unit (micro–level) productivity response to temperature T, and $g_i\big(T-\overline{T}_{L\tau}\big)$ is the probability of experiencing temperature T when the mean is $\overline{T}_{L\tau}$. $g_i\big(T-\overline{T}_{L\tau}\big)$ gives the likelihood of observing temperature T in a year when the average temperature is $\overline{T}_{L\tau}$. The integral aggregates micro-level productivity over all T, and the summation over i aggregates across all industries.

3.1.2. Piecewise linear response

At the micro–level, each small production unit in industry i responds to changes in temperature. Following BHM, the productivity function of these units, denoted as $f_i(T)$, is assumed to follow a piecewise linear relationship with temperature:

$$f_i(T) \;=\; \begin{cases} c_1 + b_1 T, & T < \tilde{T} \\ c_2 + b_2 T, & T \geq \tilde{T} \end{cases} \tag{3.4}$$

Here, T is the instantaneous temperature, $\tilde{T}$ is the temperature at which productivity peaks, b_1 and b_2 are the slopes of the productivity response below and above $\tilde{T}$, and c_1 and c_2 are intercepts chosen such that the function is continuous at $\tilde{T}$. This implies that $c_1 + b_1\tilde{T} = c_2 + b_2\tilde{T}$.

BHM then go on to consider a full year with an average temperature $\overline{T}$. Throughout this year, each production unit experiences a range of temperatures. The function $g_i(T - \overline{T})$ represents the fraction of time that units in industry i spend at temperature T, relative to the annual average temperature $\overline{T}$. This distribution shows how exposure to different temperatures is spread over time and space. Assuming that this distribution is symmetric around $\overline{T}$, the change in output per unit of productive input in response to a small change in the annual average temperature is given by:

$$\frac{\partial}{\partial \overline{T}}\big(Y_i/M_i\big) \;=\; b_1\, m_{i1}(\overline{T}) \;+\; b_2\, m_{i2}(\overline{T}) \tag{3.5}$$

In this Equation 3.5, Y_i is the total output of industry i, M_i is the total quantity of productive resources (combining capital and labour), and $\frac{Y_i}{M_i}$ represents output per unit of productive input. The terms $m_{i1}(\overline{T})$ and $m_{i2}(\overline{T})$ are the shares of time that units spend below and above the threshold temperature $\tilde{T}$, respectively. By definition, these weights sum to one: $m_{i1} + m_{i2} = 1$. If most time is spent at temperatures below $\tilde{T}$, then m_{i1} is large and the positive slope b_1 dominates the response, so an increase in $\overline{T}$ raises output. Opposite to this, if most time is spent above $\tilde{T}$, then m_{i2} is large and the negative slope b_2 dominates, which points towards a reduction in output. The turning point temperature $\overline{T}^*$, at which the two effects balance and productivity peaks, will then satisfy:

$$b_1\, m_{i1}(\overline{T}^*) + b_2\, m_{i2}(\overline{T}^*) = 0$$

This relationship creates a concave shape between average temperature and productivity per unit of input. Although the individual productivity function $f_i(T)$ has a kink at $\tilde{T}$, the aggregation across time and units (described by the distribution $g_i(T - \overline{T})$) smooths this effect. As a result, the macro-level productivity function $\frac{Y_i}{M_i}$ is continuous and differentiable, even though the micro–level response is not. This smoothing becomes more pronounced when the distribution $g_i(T - \overline{T})$ is wider, either due to more variation in climate conditions or longer observational periods.

While the piecewise linear model is not directly used in later stages of the analysis by BHM, it serves as a conceptual foundation to justify the non-linear relationship between temperature and output observed at the macro-level. The functions $f_i(T)$ and $g_i(\cdot)$ describe how individual production units respond to temperature. However, these underlying mechanisms cannot be observed directly in macroeconomic datasets. Therefore, the authors move away from this micro–level structure and instead adopt a reduced-form empirical approach grounded in the Solow growth model, which will be discussed in the following section 3.2. The sharp kink in the micro–level productivity function motivates the use of a flexible functional form in the estimation, but the specific piecewise structure is not used further in the analysis. Rather than modelling the distribution of exposure or micro–level heterogeneity explicitly, the authors estimate how average annual temperature influences GDP growth, allowing the data to reveal the aggregate effects of climate variation.

3.1.3. Model foundation and growth transformation

BHM estimate a direct and observable relationship between GDP per capita growth and average temperature, using a Solow-style framework, a standard economic model where output depends on capital, labour, and productivity. In this framework, total output in period t can be written as:

$$Y_t \;=\; \psi\big(\overline{T}_t\big)\, \gamma\, M_t \tag{3.6}$$

In this Equation 3.6, $\psi(\overline{T}_t)$ captures the effect of average temperature on productivity. The parameter γ is a constant which represents structural factors such as technology and efficiency that do not vary with temperature. M_t represents the capital stock in year t, which are machines, infrastructure, and other physical equipment used in production.

BHM let capital evolve over time using to the standard law of motion:

$$\Delta M_t = s\, Y_t - \delta\, M_t$$

Here, s is the savings rate, the part of output that is invested in new capital and δ is the depreciation rate, representing the share of capital that loses its value each year or is no longer in use. Capital increases when investment is higher then depreciation and decreases otherwise. Substituting this expression into Equation 3.6, output in period t can be expressed in terms of last period's capital and temperature as:

$$Y_t \;=\; \psi(\overline{T}_t)\,\gamma\Big(M_{t-1} + s\,\psi(\overline{T}_{t-1})\,\gamma\,M_{t-1} - \delta\,M_{t-1}\Big) \tag{3.7}$$

This reformulation shows how temperature in both the current year, $\overline{T}_t$, and the previous year, $\overline{T}_{t-1}$, influences output through two channels: a direct impact on productivity, via $\psi(\overline{T}_t)$, and an indirect effect on capital accumulation, as past productivity determines the savings and investment that shape next years capital stock. While this form is not estimated directly in the model by BHM, it provides economic intuition for the reduced-form approach used for the main regression model. It clarifies how short-term fluctuations in temperature may have persistent effects on output. In the empirical approach by BHM, observed temperature variations are linked to output growth, without explicitly estimating each structural component of the production process.

An important econometric challenge is that output levels typically follow a unit root process, meaning their average and variability can drift over time. If one were to regress the output level Y_t on temperature, any long-run trend in income would tend to coincide with climate patterns by chance. Such a coincidence creates a spurious correlation: two variables that both trend can appear linked even when there is no real effect. To avoid this, BHM take the natural logarithm of output and then apply first differencing. This produces:

$$\Delta\ln(Y_t) \;=\; \ln(Y_t) \;-\; \ln(Y_{t-1})$$

This transformation measures the year-on-year growth rate of GDP per capita. It has two advantages. First, it isolates short-run changes in output, helping to identify how annual deviations in temperature influence economic growth rather than long-run income levels. Second, it removes trending behaviour in GDP, which differs systematically across countries, and avoids picking up non-causal associations driven by these trends. The result is a variable, $\Delta\ln Y_{it}$, that is (approximately) stationary, allowing the model to isolate the impact of temperature variability on economic growth.

3.2. Reduced form regression model

Using this transformed variable as the dependent variable from the previous subsection (3.1.3), BHM estimate the following reduced-form panel regression model:

$$\Delta\ln Y_{i,t} \;=\; h(T_{i,t}) \;+\; \lambda_1\,P_{i,t} \;+\; \lambda_2\,P_{i,t}^2 \;+\; \mu_i \;+\; \nu_t \;+\; \theta_{i1}\,t \;+\; \theta_{i2}\,t^2 \;+\; \varepsilon_{i,t} \tag{3.8}$$

This regression model links the annual growth rate of real GDP per capita, $\Delta\ln Y_{i,t}$, to climate and control variables, while using panel data. The term $h(T_{i,t})$ captures the potentially non-linear effect of annual average temperature on economic growth, and is specified as:

$$h(T_{i,t}) = \beta_1 T_{i,t} + \beta_2 T_{i,t}^2 \tag{3.9}$$

The choice for this quadratic functional form is grounded in the micro–economic framework developed by BHM and outlined in Section 3.1. In that framework, productivity at the unit level responds to temperature in a piecewise linear manner (see Equation 3.4). When this response is aggregated over time, locations, and industries, the result is a smooth, concave relationship between average temperature and output per unit of input, as described in Equation 3.5. Although the piecewise form is not estimated directly, it motivates the use of a more flexible and continuous specification at the macro-level. The quadratic approximation in Equation 3.9 captures the empirical pattern observed in the data: economic performance improves with rising temperatures up to a certain point, after which further warming reduces growth. This concave shape is consistent with both theoretical expectations and empirical regularities, making the quadratic approximation a clear and practical way to capture non-linear climategrowth relationships in the reduced-form model BHM.

In Equation 3.8, i denotes a country and t a year. The dependent variable, $\Delta\ln Y_{i,t}$, represents the annual growth rate of real GDP per capita. Temperature $T_{i,t}$, measured in $\circ$C, enters through the non-linear function $h(\cdot)$. Precipitation $P_{i,t}$, measured in millimetres, enters both linearly and quadratically to account for potential non-linear effects. Country FE μ_i control for time-invariant differences across countries, such as geography or institutions, while year FE ν_t capture global shocks common to all countries in a given year, such as financial crises or commodity price fluctuations. To control for gradual structural changes within countries, the model includes country-specific time trends, both linear and quadratic ($\theta_{i1}t + \theta_{i2}t^2$), which may reflect factors like demographic change or evolving policy environments. The error term $\varepsilon_{i,t}$ captures unexplained variation in growth. A cluster-robust variance estimator is used, allowing $\varepsilon_{i,t}$ to be serially correlated within countries over time, while assuming independence across countries. This ensures consistent inference even in the presence of within-country autocorrelation. Estimation details are further discussed in Subsection 3.2.2.

3.2.1. Panel data

The research strategy of the model by BHM relies on panel data, which is widely used in climate econometrics to examine the complex relationship between climate variables and economic outcomes (Kolstad et al., 2020). Panel data combines cross-sectional and time-series dimensions by tracking the same observational units, such as countries or regions, over multiple years (Blanc et al., 2017; Chang et al., 2023; Mérel et al., 2021). This leads to a higher number of observations and therefore makes it possible for researchers to account for both differences across units and changes within them over time, which makes statistical estimates more precise (Baltagi et al., 2011; Blanc et al., 2017; Chang et al., 2023; Kotz et al., 2024). Panel data improves the estimation efficiency because it increases the variation in the data since it combines observations between units with observations over time (Baltagi et al., 2011). Because of this, the number of degrees of freedom is increased since $N \times T$ observations (where N is the number of units and T the number of periods) (Blanc et al., 2017). Because of this, the precision of parameter estimates is increased. Instead of capturing a static observation of a unit at one moment in time, it enables researchers to analyse how economic outcomes evolve in response to changing climate conditions within a specific unit, while also comparing these effects across units. Because panel data allows for evolution over time, it increases the statistical precision and supports more robust results of the model at hand (Baltagi et al., 2011).

To quantify relationships between variables using panel data, regression models are used. These models generate statistical estimates, which are numerical values that represent how an outcome, such as GDP growth, changes when a predictor, such as temperature, varies by one unit, holding other factors constant. Through the model, these estimates are derived from the data by minimising the difference between the models predictions and actual outcomes (Chang et al., 2023). The models coefficients provide insight into the strength and direction of relationships between climate variables and economic growth (Kotz et al., 2022).

A strength of a panel data regression model is its ability to control for unobserved factors that do not vary over time, such as geography, historical and geopolitical factors, institutions, or soil quality (Chang et al., 2023; Kotz et al., 2024; Blanc et al., 2017). These unchanging characteristics, known as unobserved time-invariant heterogeneity, may influence both climate conditions and economic growth. If these are not included, it can introduce omitted variable bias, leading to incorrect estimates of the true climateeconomy relationship (Blanc et al., 2017). To prevent omitted variable bias, BHM include FE in their model. Omitted variable bias can occur when an unobserved factor both influences the dependent variable and correlates with one or more explanatory variables. In such cases, the model may wrongly attribute part of the omitted variable effect to the included variables, which in turn leads to wrong inference about the true climateeconomy relationship. FE control for such unobserved, time-invariant differences between geographic units. By controlling for this, FE help to isolate the effect of climate variation on economic growth (Blanc et al., 2017). In this way, panel data regression models using FE have an advantage over cross-sectional models. By using variation over time within each unit, regressions using panel data with FE are able to control for a wider range of relevant factors, both observed and unobserved, that could influence economic outcomes (Kolstad et al., 2020).

3.2.2. Model estimation

BHM estimate their model using Ordinary Least Squares (OLS), a standard method in regression analysis. OLS estimates the relationship between variables by finding the coefficients that minimise the sum of squared differences between the observed values of the dependent variable and the values predicted by the model (Baltagi et al., 2011).

For OLS estimates to be reliable, certain assumptions about the error term (the unexplained part of the outcome) need to hold. One assumption of those is that the variance of the error term has to remain constant across all observations. This assumption is called homoscedasticity (Baltagi et al., 2011). If the variance varies, it is called heteroscedasticity. This violation of homoscedasticity can be due to measurement errors, missing variables, or randomness in the data, and can lead to unreliable standard errors and misleading results. A second assumption is that errors across observations should not be related to each other. This assumption is called no serial correlation (Baltagi et al., 2011). In the case of panel data, this means that the error in one year should not systematically affect the error in the next. However, in panel data regression models such as the one used by BHM, observations are grouped by country and repeated over time. This often results in error terms being correlated within each group, in this case per country, and therefore having unequal variances, violating both assumptions above.

To address this issue and ensure valid results, BHM apply a cluster-robust variance estimator. This method allows for arbitrary patterns of heteroscedasticity and autocorrelation within each cluster (which is by country in the BHM model), while assuming independence across countries (Baltagi et al., 2011). This improves the accuracy of standard errors, confidence intervals, and statistical tests in the presence of clustered data (Bester et al., 2011). Although the regression coefficients are still estimated using OLS, the corrected standard errors make the results more valid and suitable for panel data analysis. If there is no cluster-robust variance estimator used, the standard errors of the model will be smaller, while they are in reality larger. This can lead to overstated statistical significance of the estimated coefficients, meaning that they may appear more precise or reliable than they actually are. This will make the model less valid.

The model estimated by BHM is a reduced-form regression model (given by Equation 3.8). This means it directly links observed economic outcomes to climate variables, without explicitly modelling the underlying mechanisms. The advantage of this approach is that it captures the combined effect of all causal channels, whether known or unknown, without requiring detailed assumptions about the structure of the economy or the behaviour of economic agents (Kolstad et al., 2020). However, this simplicity has a cost. Because the model does not explicitly model all the individual mechanisms, it cannot identify how or why climate affects economic growth. In addition, reduced-form estimates are typically valid only within the historical range of the data. If future climate conditions differ substantially, such as through extreme temperature shifts or unprecedented SLR, the model may not predict outcomes accurately. In this sense, reduced-form models provide a useful summary of past relationships, but their predictive power outside observed conditions is limited (Carleton et al., 2016; Kotz et al., 2022).

3.3. Temperature-growth relationship

The main aim of the model by BHM is to investigate whether temperature variation affects economic growth, and if this relationship is non-linear. Using the findings of earlier micro–level studies, the authors note that modest warming can increase productivity in colder areas, while more extreme warming can reduce it. The model aims to explore whether similar threshold effects exist at the global level, and if so, to identify the temperature range in which economic performance peaks before declining on each side.

3.3.1. Non-linear effect

A non-linear relationship means that the effect on economic growth of a 1 °C temperature increase is not the same across all temperature levels. Thus, the economic response to a rise in temperature depends on the initial temperature of a country. The results of BHM suggest that GDP growth is highest, globally, at an average annual temperature of approximately 13 °C. This relationship is captured using a quadratic temperature term in the model (Equation 3.9).

To identify the temperature at which growth peaks, BHM take the derivative of the impact function, Equation 3.9 with respect to temperature:

$$\frac{\partial h}{\partial T} = 0 \quad \Longleftrightarrow \quad T = T^*$$

This derivative measures how GDP per capita growth responds to a change in temperature. The turning point, T^*, occurs when this derivative equals zero. This is the temperature at which the effect of temperature on GDP per capita growth changes direction. If the estimated relationship is concave (the quadratic coefficient is negative), T^* represents a maximum and growth declines beyond this point. If the relationship is convex (the quadratic coefficient is positive), T^* marks a minimum, and growth increases with further warming. BHM observe a concave relation between temperature and economic growth, making the derivative of the impact function T^* a maximum.

3.3.2. Explanatory variables and dependent variable

The main explanatory variables in the model are temperature and precipitation, even though precipitation is included to control for temperature. This is a common practice since precipitation can confound the effects of temperature (Burke, Hsiang, et al., 2015). This arises since temperature and precipitation tend to correlate over time (Burke, Hsiang, et al., 2015; Blanc et al., 2017). In the model by BHM temperature T_{it} is defined as the annual average temperature for country i in year t, measured in °C. To calculate the annual average temperature, BHM use 0.5° gridded monthly temperature data weighted by the population. Each grid cell, denoted by $\mathcal{C}_i$, covers approximately 50 by 50 kilometres near the equator. The population-weighted average temperature is given by:

$$T_{i,t} = \frac{\sum_{c \in \mathcal{C}_i} \mathrm{POP}_c \, T_{c,t}}{\sum_{c \in \mathcal{C}_i} \mathrm{POP}_c}$$

Using population weights ensures that the calculated temperature reflects the climate conditions experienced by most people and the areas of greatest economic activity (Hsiang, 2016; Burke, Hsiang, et al., 2015). Precipitation P_{it}, measured in millimetres and is also population–weighted.

To isolate the true effects of temperature and precipitation, the model, as mentioned in Equation 3.8 also includes several control terms. Country FE μ_i account for time-invariant characteristics such as geography or institutional context, while year FE ν_t control for global shocks affecting all countries in a given year. The regression also includes country-specific linear and quadratic time trends, $\theta_{i1}t + \theta_{i2}t^2$, which capture slow-moving changes within countries that are unrelated to climate, such as demographic shifts.

The dependent variable is the annual growth rate of real GDP per capita, calculated as the first difference of its natural logarithm. By differencing the log of GDP per capita, the model captures short-term fluctuations in growth. This has an advantage since it allows for making comparisons across countries regardless of their initial level of GDP (Burke, Hsiang, et al., 2015). But this transformation is also necessary for valid inference, especially when working with time series or panel data that may contain unit roots, as mentioned in subsection 3.2 (Burke, Hsiang, et al., 2015).

As discussed in Chapter 2, a central debate in climate econometrics is whether temperature affects the level of economic output or its growth rate. Earlier models often assumed level effects, which means that they assumed that damages did not accumulate over time. In contrast, BHM focus on growth effects, assuming that damages from climate shocks can persist and compound.

3.4. Insight from the model

The regression results from the model by BHM are the estimated coefficients $\hat{\beta}_1$ and $\hat{\beta}_2$, which define the global temperature–growth response function $\hat{h}(T)$. This function describes how a 1°C change in annual average temperature affects GDP growth, depending on the starting temperature level. When the fitted function $\hat{h}(T)$ is plotted across the observed temperature range in the panel data, it forms a concave curve. Using $\hat{h}(T)$, the marginal effect of additional warming on GDP growth can be computed for each country, based on its historical average temperature. The hat notation means that these are sample–based estimates, not the true underlying values. While the estimates provide a best-fit summary of observed data, they remain approximations and should be interpreted with this in mind.

The key empirical finding of BHM is the concave shape of the temperature-growth relationship. In their benchmark model, the fitted function is:

$$\hat{h}(T) = 0.0178\,T - 0.0007\,T^2,$$

which implies a turning point at:

$$\widehat{T}^* = -\frac{\hat{\beta}_1}{2\,\hat{\beta}_2} \approx 13°\text{C}.$$

For temperatures below the threshold of ($T < 13$°C), the marginal effect is positive, indicating that moderate warming can increase economic growth. But for temperatures above this threshold ($T >$ 13°C), an increase in temperature reduces growth. Thus BHM conclude that colder countries, such as many in Europe and North America, may experience small positive effects from moderate warming, while countries in warmer climates, especially in tropical regions, are projected to suffer stronger negative impacts from further temperature increases.

3.5. Relevance of the BHM model

Understanding the structure and logic of the model by BHM is important for this research, as it provides the theoretical and empirical foundation for estimating climate–economy relationships using historical data. The model introduced a novel reduced–form approach that identifies a concave relationship between annual temperature and GDP per capita growth, using global panel data. This thesis re–analyses and extends the BHM model to test its robustness with a longer time series and an additional climate variable at the regional–level in Europe.

In the following chapter, the model will be applied using alternative data sources. Specifically, the Rosés et al. (2021) V6 dataset will be used to obtain historical GDP per capita figures at the regional (NUTS-2) level in Europe. Climate variables, temperature, precipitation, and SLR, will be incorporated to estimate how past climatic variation is related to long-run economic growth across European regions. The methodological choices made by BHM, including the use of panel data, and FE, serve as a blueprint for constructing the extended empirical framework of this study.

4
Data preparation

This chapter provides an overview of the datasets used to answer the central research question: How have temperature, precipitation, and SLR, affected regional GDP per–capita growth in Europe between 1900 and 2015? The chapter describes the source, structure, and transformation of the data prior to regression analysis. It begins by introducing the main data sources and preparation steps in section 4.1. In this section, the inclusion and exclusion criteria applied to construct a balanced panel of regions are discussed. In Subsection 4.2.1 the exploratory data analysis will be discussed and lastly in Subsection 4.3 the control variables used in this research.

4.1. Data Sources and preparation

The empirical analysis is based on a panel dataset of European NUTS-2 regions, covering the period from 1900 to 2015. NUTS–2 regions are defined under the European Unions Nomenclature of Territorial Units for Statistics and represent sub–national administrative areas (regions). The dataset consists of benchmark-year GDP, population, and climate observations for 170 regions across 16 countries. The climate variables in this panel dataset include averages of annual temperature and precipitation, as well as the difference in relative SLR since 1900.

To ensure data consistency and validity, the data is restricted based on the following inclusion criteria. First, regions must have at least 12 benchmark-year GDP observations between 1900 and 2015 to allow for the computation of per capita growth. Second, annual average temperature and precipitation values must be available for each benchmark year that has a regional GDP observation. Third, population data must be available for benchmark years in which regional GDP is available, as it is used to construct per capita GDP. Since the regional GDP and population data are complete for the benchmark years, no regions needed to be excluded for this reason. Temperature and precipitation data are also available for all benchmark years, except for two of the 173 regions, which were excluded due to missing values. One additional region, Flevoland (NL23), was excluded because it only has five benchmark–year GDP observations. Since all other regions have 12 observations, Flevoland was dropped. After these exclusions, the final dataset contains 170 regions. Finally, no exclusion is imposed based on SLR coverage, but a thorough interpolation was needed. After interpolation, 82 coastal regions and 88 non–coastal regions having 12 SLR observations for benchmark years were established.

4.1.1. Economic data

This research uses the V6 dataset by Rosés et al. (2021), which provides the GDP estimates available for twelve benchmark years within this time frame, with each year containing one observation per region. These benchmark years are visualised in Figure 4.1.

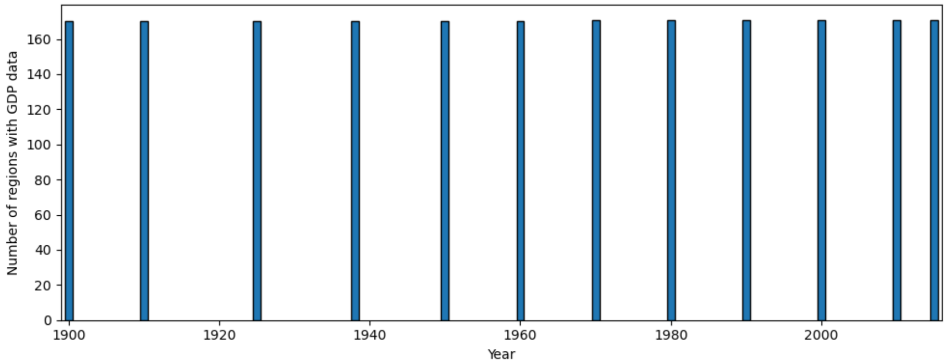

Figure 4.1: Count of GDP observations per benchmark year

The main variable of interest is regional GDP, expressed in 2011 international dollars and adjusted for purchasing power parity. This adjustment ensures that a dollar of output represents a similar volume of goods and services in every region, allowing for meaningful comparisons across space and time. Additional variables include annual population (in thousands) and land area (in square kilometres). Because GDP is not recorded annually but only for certain benchmark years, there are long intervals between some observations. Filling in these gaps through interpolation would risk introducing bias. Therefore, only the recorded values from the benchmark years are used. When GDP values are available for two consecutive benchmark years, the growth rate is calculated directly from these observations.

To prepare the data, regional GDP is first expressed in per capita terms by dividing total output by the regional population in each given benchmark year. This provides a measure of average economic output per person. Figure 4.2 illustrates how GDP per capita has changed across regions between 1900 and 2015. The early part of the century shows modest growth, with a stronger rise beginning after 1950. According to Rosés et al. (2021), this increase in GDP per–capita can be linked to post-war reconstruction or economic integration. The 1950s and 1960s marked a period of strong convergence in Europe, driven by industrial expansion, structural transformation, and improved infrastructure. Over time, differences between regions have become much larger than they were, especially from the 1960s onwards.

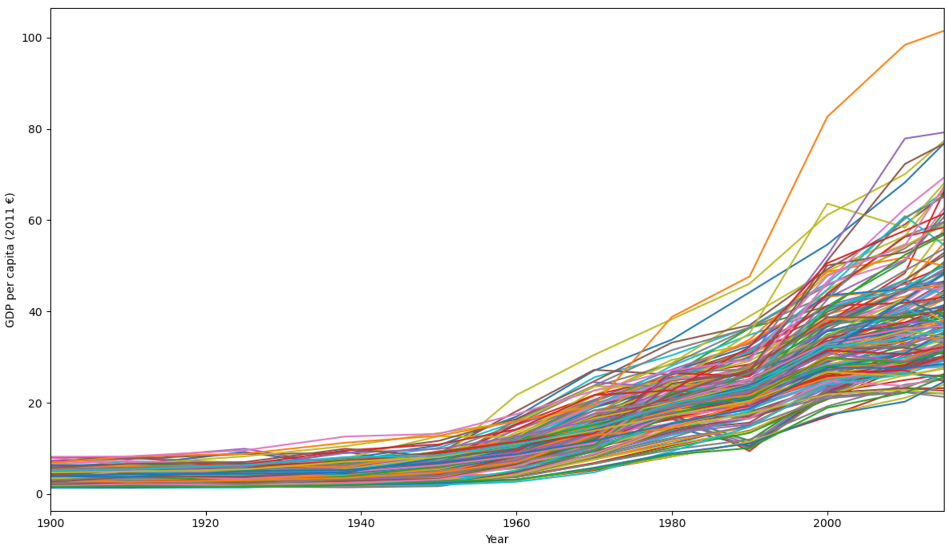

Figure 4.2: Evolution of GDP per capita over time by (NUTS–2) region

To complement the regional trajectories shown in Figure 4.2, Figure 4.3 presents the average GDP per capita across all NUTS–2 regions over time. This figure shows the broader trend of economic growth throughout the twentieth century in the 170 regions of Europe, with a more pronounced increase after the Second World War. The steady upward path after 1950 reflects the combined effect of economic recovery and technological progress across Europe, as described by Rosés et al. (2021). This average growth pattern underscores the long-run increase of regional incomes.

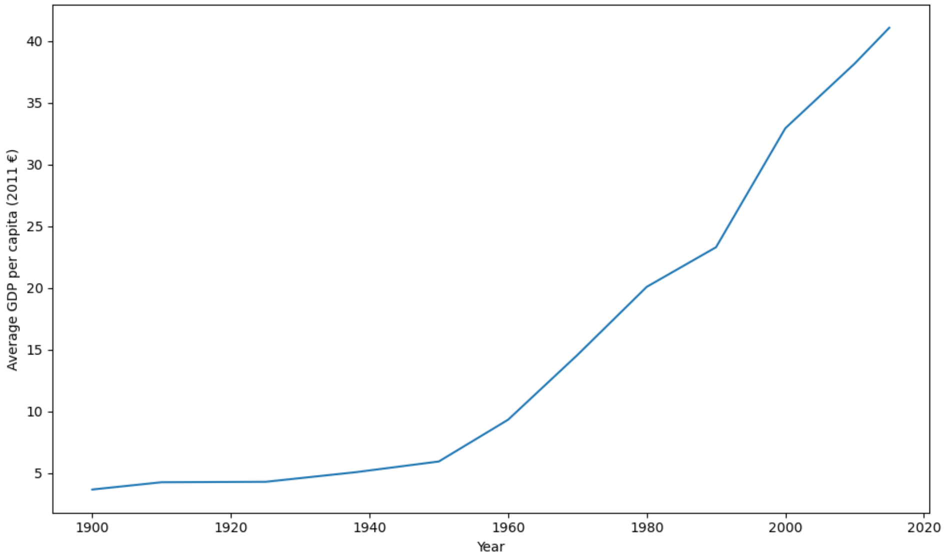

Figure 4.3: Average of GDP per capita over time by (NUTS–2) region

Following the approach by BHM, the main dependent variable in the model is defined as the change in the natural logarithm of GDP per capita between two consecutive benchmark years. Observations that lack sufficient information to calculate this value are excluded, so the sample begins in 1910. After this step, the dataset becomes a balanced panel of 170 regions, covering the period from 1910 to 2015. It is now ready to be merged with the climate data from Climatic Research Unit, University of East Anglia (n.d.) for the analysis described in Chapter 3.

Figure 4.4 displays the annual GDP per capita growth for each region from 1910 to 2015. Most regions show small but positive growth rates, typically between 0.1 and 0.25. A strong decline is visible around 1920, which coincides with the economic aftermath of the First World War. A notable recovery follows in the 1950s and 1960s, with many regions experiencing growth rates above 0.5. The oil crises of the 1970s and 1980s appear as a marked slowdown. Growth picks up again from the late 1990s into the early 2000s, followed by a downturn after the global financial crisis of 2008 (Rosés et al., 2021). Figure 4.4 also shows the considerable volatility of regional growth over time. Fluctuations are particularly strong during periods of global or continental shocks such as wars, depressions, and financial crises. According to Rosés et al. (2021), this volatility reflects both the sensitivity of regional economies to macroeconomic conditions and structural differences in how they respond to these kinds of shocks. Such heterogeneity underlines the importance of examining regional dynamics over time rather than relying solely on national aggregates.

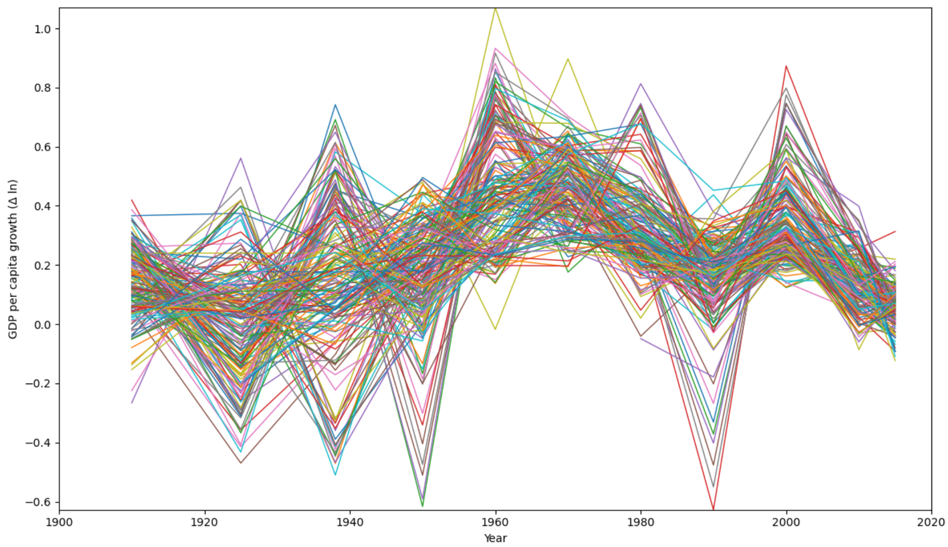

Figure 4.4: Evolution of GDP per capita growth over time by (NUTS–2) region

Figure 4.5 presents the average growth rate across all NUTS–2 regions. In the early decades of the twentieth century, growth remained modest, reaching close to zero around the time of the first world war. During the 1930s, average growth started to recover and increased sharply after 1950, peaking in the 1960s. This expansion slowed through the 1970s and 1980s, before another period of faster growth began around the 2000s. Following the 2008 crisis, growth declined again, falling to approximately 0.07 by 2015. The pattern reflects a strong post-war boom and a lack of sustained recovery in the aftermath of recent economic shocks (Rosés et al., 2021). It is important to note that the y–axis in Figure 4.5 represents the GDP per–capita growth in benchmark years. This means that the growth rates reflect the growth rates between those benchmark years, not annual growth rates as in the paper by BHM.

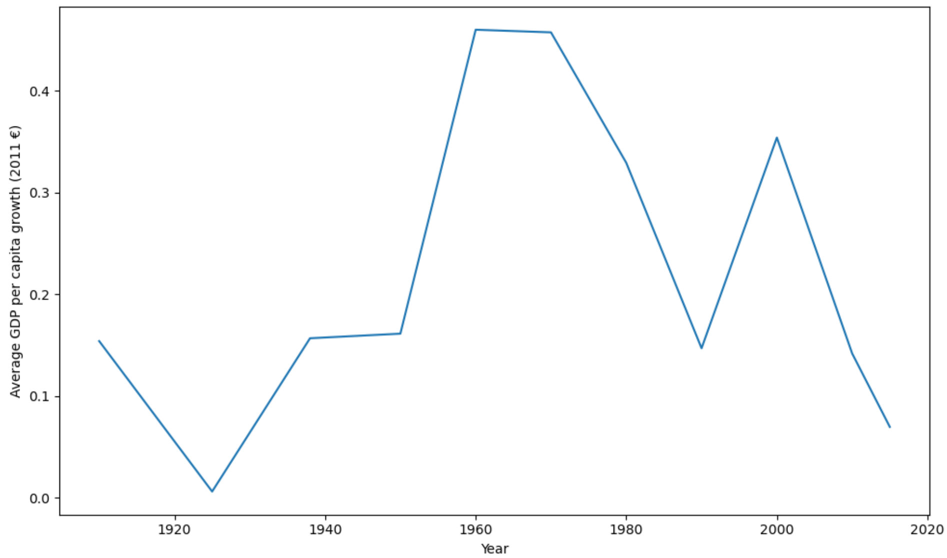

Figure 4.5: Average GDP per capita growth over time across all (NUTS–2) regions

4.1.2. Climate data

The climate data comes from the Climatic Research Unit, University of East Anglia (n.d.) dataset maintained by the University of East Anglia (Climatic Research Unit, University of East Anglia, n.d.). This dataset provides gridded annual observations for temperature and precipitation at a resolution of 0.5° × 0.5°, meaning each cell covers roughly 50 km × 50 km near the equator, with slightly smaller dimensions at higher latitudes. For this research, the CRU temperature and precipitation values have already been aggregated to NUTS-2 regions using these 0.5° × 0.5° grid cells. The temporal coverage begins in 1900 and extends through 2021. Since the Rosés et al. (2021) only provides regional GDP data up to 2015 the climatic data is cut at this year. Also, for two of the 173 regions, there was no temperature and precipitation data available, which resulted in the dropping of these two regions. Leaving 171 remaining regions. One region, Flevoland (NL23), only had 5 observations since the land was reclaimed from the sea during the 20th century after which the province itself was only officially established in 1986 (Rosés et al., 2021). Since all the other regions do have 12 observations, Flevoland was dropped from the dataset. Leaving 170 remaining regions. The main climate variables used in the analysis are the mean annual temperature in °C for each region, and the mean annual precipitation in millimetres. In addition, annual relative sealevel data (in millimetres) are obtained from the Permanent Service for Mean Sea Level Revised Local Reference dataset for each coastal region (Holgate et al., 2013).

Temperature

Figure 4.6 shows the evolution of the average 30-year temperature across all regions from 1900 to 2015. The construction of this rolling mean is explained in Appendix D, subsection D.1. During the first half of the twentieth century, a steady rise in temperature is observed, increasing from approximately 8.3°C in 1900 to just under 9.0°C by 1960. This warming trend then slows, with a slight decline between 1960 and 1980, often referred to as the mid-20th century cooling (Malhi et al., 2021; Dell et al., 2008). From 1980 onward, however, the trend shifts, more sharply than before, upward, reaching nearly 9.8°C by 2015. This recent acceleration is consistent with global patterns linked to anthropogenic climate change (Kahn et al., 2021; Malhi et al., 2021; Stern, 2008).

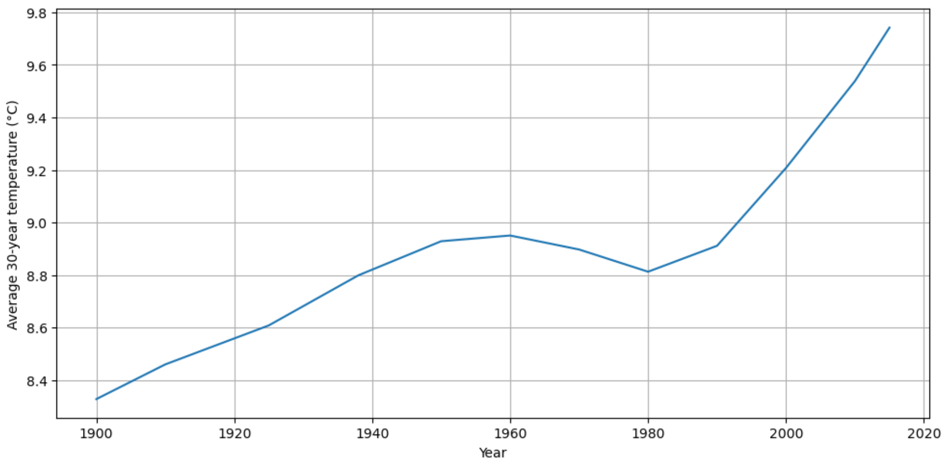

Figure 4.6: Average 30-year rolling mean temperature across 170 regions

Figure 4.7 shows the annual average temperature for each NUTS–2 region in europe from 1900 to 2015. Each coloured line represents one region, showing both the within-region trend and the differences between regions across the full period. There is a clear upward trend, reflecting long-term warming consistent with climate studies (Twardosz et al., 2021). A notable feature is a sharp dip around the year 2010. This dip is unexpected and not supported by the literature, as temperatures in Europe and globally continued to rise during this period (Twardosz et al., 2021). because the dip is visible across many regions at the same time, it is likely caused by a structural break or coverage difference in the panel data. Since this pattern is present in many regions, no corrections are made. Figure 4.6 shows the rolling average temperature over thirty years for the same period. This method smooths short-term fluctuations while preserving long-term trends. the rolling average shows a consistent increase without the sharp dip near 2010. This supports the interpretation that the dip seen in figure 4.7 is not a climate signal but a data artefact related to panel structure or regional coverage in the year 2010.

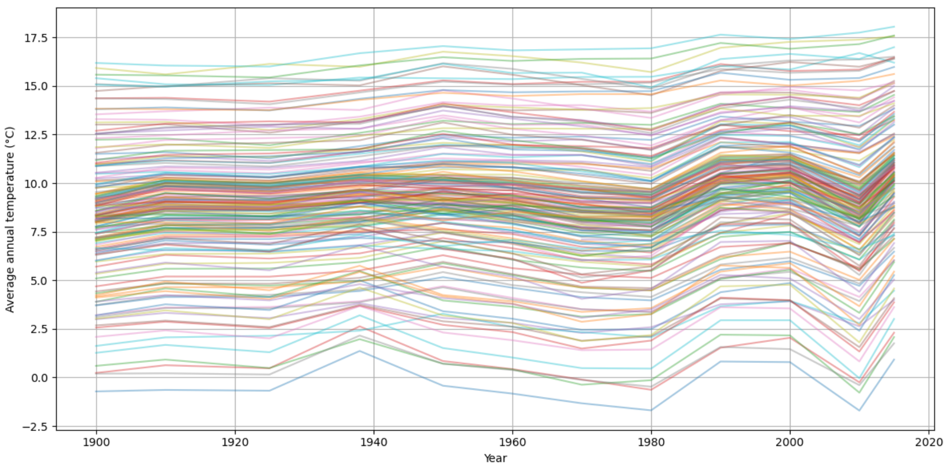

Figure 4.7: Annual average temperature over time across 170 regions

Taking the average over the 170 regions gives Figure 4.8. The figure shows an upward trend in temperature, which aligns with the expected long-term warming pattern (Twardosz et al., 2021). The sharp dip around the year 2010, followed by a strong increase towards 2015, which was also visible in Figure 4.7, is also visible here. The overall pattern confirms the long-term warming in Europe, with the average temperature increasing from about 8.3 °c in the early 1900s to over 10 °c by 2015.

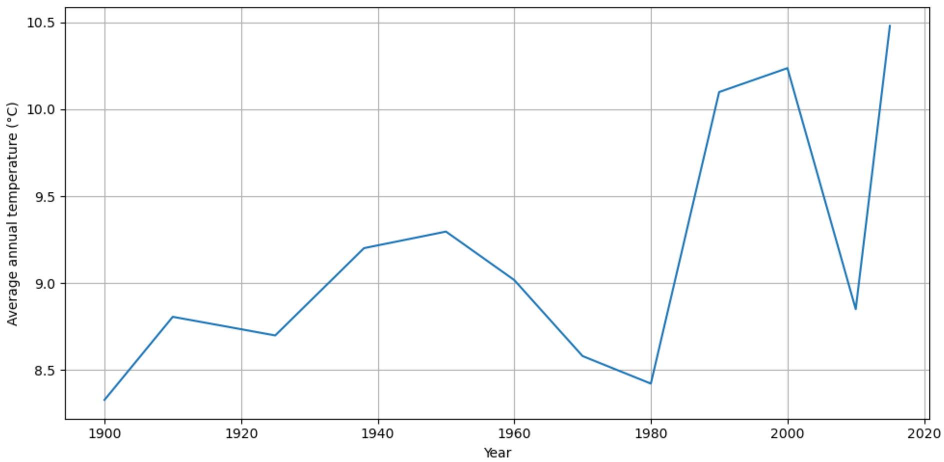

Figure 4.8: Average annual temperature over time

Compared to Figure 4.7 and Figure 4.8, Figure 4.6 presents a clearer view of long-term temperature trends by using a 30-year rolling mean. While the annual and regional figures show substantial year-to-year variability and a sharp dip around 2010, the rolling mean smooths these fluctuations and confirms a steady upward trend. The figure shows an increase from about 8.3 °C in 1900 to nearly 9.8 °C by 2015, with a brief plateau between 1960 and 1980, consistent with the mid-20th century cooling described in the literature (Malhi et al., 2021; Dell et al., 2008). This recent acceleration after 1980 aligns with global warming patterns linked to anthropogenic climate change (Kahn et al., 2021; Malhi et al., 2021; Stern, 2008). The absence of a sharp dip around 2010 in the rolling mean supports the interpretation that the dip seen in the annual figures is likely a data artefact related to panel structure or coverage in that year rather than a climate signal.

Precipitation

Figure 4.9 illustrates the development of average 30-year precipitation from 1900 to 2015. Between 1900 and 1935, precipitation levels increased from about 645 mm to over 670 mm. A decline followed in the 1940s, dropping to around 662 mm by 1950. From that point onwards, precipitation has increased, with only small fluctuations. After 1980, levels remained relatively stable until a sharp rise occurred after 2000, peaking near 687 mm around 2010. A slight drop is observed in the final data point for 2015. Overall, the long-term trend points to a moderate increase in precipitation across the 170 regions across the observed period. The literature offers some support for a long-term increase in global precipitation between 1900 and 2015 but does not confirm the specific fluctuations or numerical values described for Europe. While studies such as Malhi et al. (2021) note a modest positive trend over the full period, others, like Dell et al. (2009), report a decline in precipitation from the 1950s onward, suggesting more complex or regionally variable dynamics. Overall, while the literature offers some support for a long-term increase in global precipitation between 1900 and 2015, studies focusing on Europe reveal more nuanced trends with significant seasonal and decade-long variability (Pauling et al., 2006).

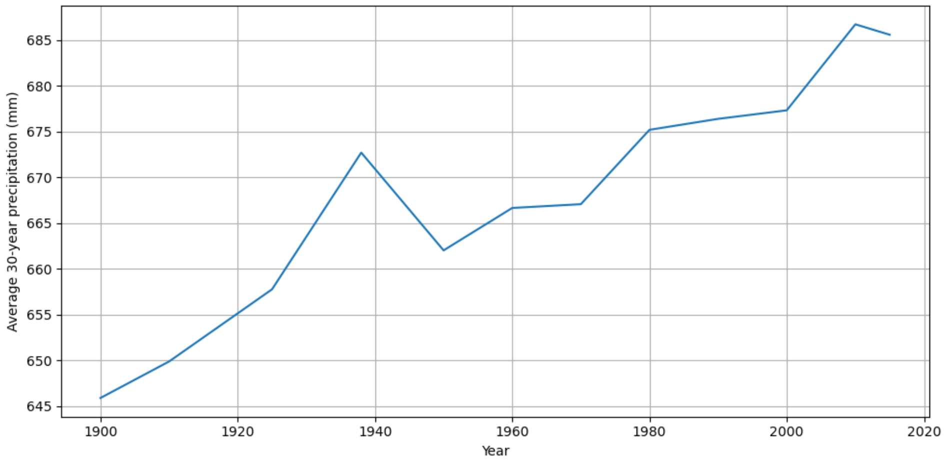

Figure 4.9: Average 30-year rolling mean precipitation across 170 regions

Figure 4.10 shows the average annual precipitation for each NUTS–2 region in Europe from 1900 to 2015. Each coloured line represents one region, showing differences in both the level and the variation of precipitation across regions. the figure shows a high variation between regions, with some regions having average precipitation below 400 mm per year while others exceed 1400 mm per year. Many regions experience a sharp peak around the year 1960 and a smaller peak around the year 2010, which is seen across many regions.

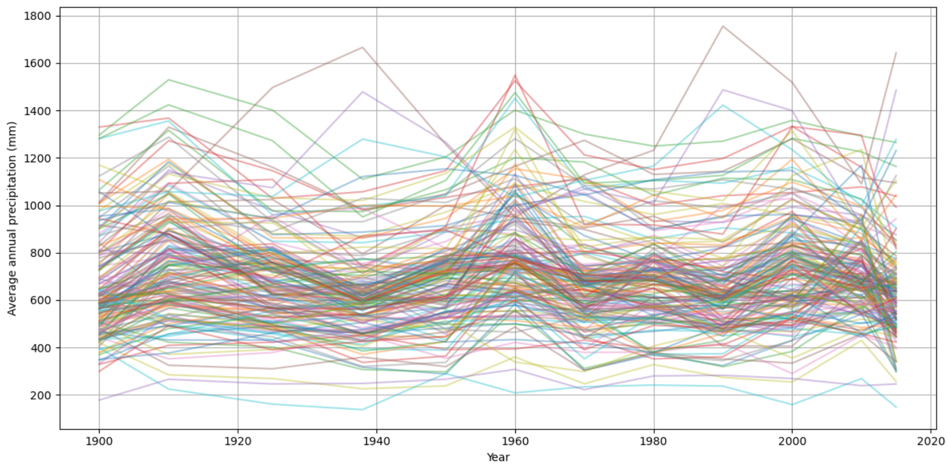

Figure 4.10: Annual average precipitation over time across 170 regions

The average annual precipitation is given by Figure 4.11. Compared to Figure 4.10, the pattern in Figure 4.11 is smoother and shows the average across all regions instead of individual lines. The sharp peak around 1960 and the smaller peak around 2010 remain visible, indicating that these features are present in many regions and affect the overall average. However, the figure does not show a clear increasing or decreasing trend over the period, suggesting that while temperature shows a systematic increase over time, precipitation remains highly variable without a clear directional change.

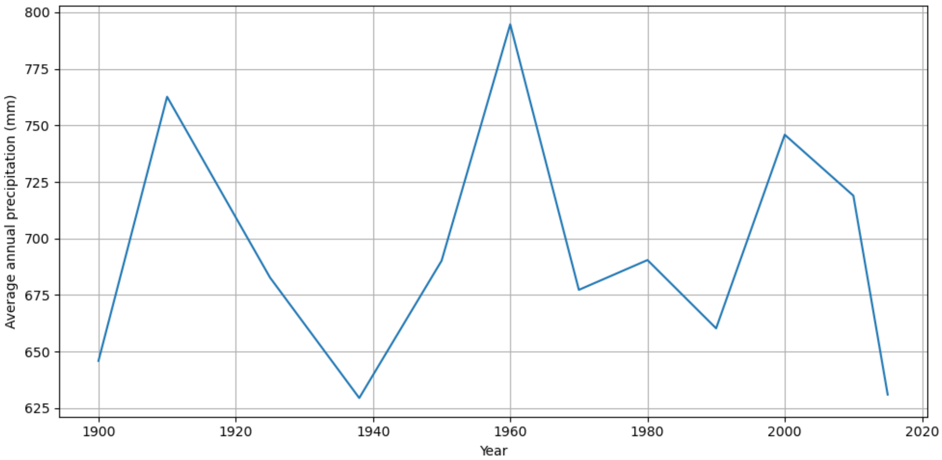

Figure 4.11: Average annual precipitation over time

Compared to Figure 4.10 and Figure 4.11, Figure 4.9 presents a clearer view of long-term trends by using a 30-year rolling mean. While the annual and regional figures show strong variability and sharp peaks around 1960 and 2010, the rolling mean smooths these fluctuations. It shows a slow increase in average precipitation from about 645 mm in 1900 to nearly 687 mm around 2010, followed by a slight decline. This suggests a moderate upward trend in precipitation over the period while confirming that year-to-year and decade-to-decade variation remains high. The rolling figure thus aligns with literature indicating a modest long-term increase in precipitation (Malhi et al., 2021).

SLR data preparation

The SLR series for each region is prepared in several steps. First, for the initial years 1900 and 1910, missing values are synchronised within each region. For all regions with more than one observed value in the full time series, if either 1900 or 1910 is missing but the other is available, the missing entry is filled by copying the observed value, depending on which value is observed. Second, for any remaining missing values, linear interpolation is applied within each region, but only for gaps of up to two consecutive years. For each missing year t with observed values at $t-2$ and $t+2$, the gap is filled as:

$$\mathrm{SLR}_{i,t} = \frac{\mathrm{SLR}_{i,t-2} + \mathrm{SLR}_{i,t+2}}{2}.$$

Longer gaps are left unfilled and addressed through spatial interpolation. Each NUTS–2 region i is matched to its six nearest neighbouring regions using the shapefile provided by Rosés et al. (2021) containing the geometry of all regions. For any missing year t, the value is replaced by the average of observed values in those neighbours:

$$\mathrm{SLR}_{i,t} = \frac{1}{N_i} \sum_{j \in \mathcal{N}(i)} \mathrm{SLR}_{j,t},$$

Where $\mathcal{N}(i)$ denotes the set of neighbouring regions with valid observations in year t, and $N_i \leq 6$ is the number of available neighbours. After these steps, each coastal region has a complete SLR series starting from 1900. Letting $t_{i,0} = 1900$ denote the first year of observation, the cumulative change in sea level relative to 1900 is defined as:

$$\Delta\mathrm{SLR}_{i,t} = \mathrm{SLR}_{i,t} - \mathrm{SLR}_{i,\,1900} \qquad \text{for } t \geq 1900.$$

For regions with no single observed SLR value, all entries are set to $\mathrm{SLR}_{i,t} = 0$, resulting in a cumulative change of $\Delta\mathrm{SLR}_{i,t} = 0$ for all t.

Figure 4.12 shows the evolution of relative SLR across 82 coastal regions over the period 1900 to 2015. Each line represents the trajectory of one region. Although the absolute levels differ considerably, several broad patterns can be observed. At the beginning of the observed period, sea level changes vary widely across regions. Some regions exhibit early increases, while others remain stable or decline. Between the 1940s and 1960s, the sea level of all the regions came closer together, by either increasing, decreasing or no SLR. This could reflect measurement inconsistencies, vertical land movement, or local subsidence (Holgate et al., 2013). From the 1970s onward, regional trajectories became more stable or began to rise more consistently. After 1990, most regions exhibit a moderate upward trend, indicating an acceleration in relative SLR in recent times. Despite regional variation, the synchronisation of trends in this later period aligns with broader global patterns of SLR (Holgate et al., 2013).

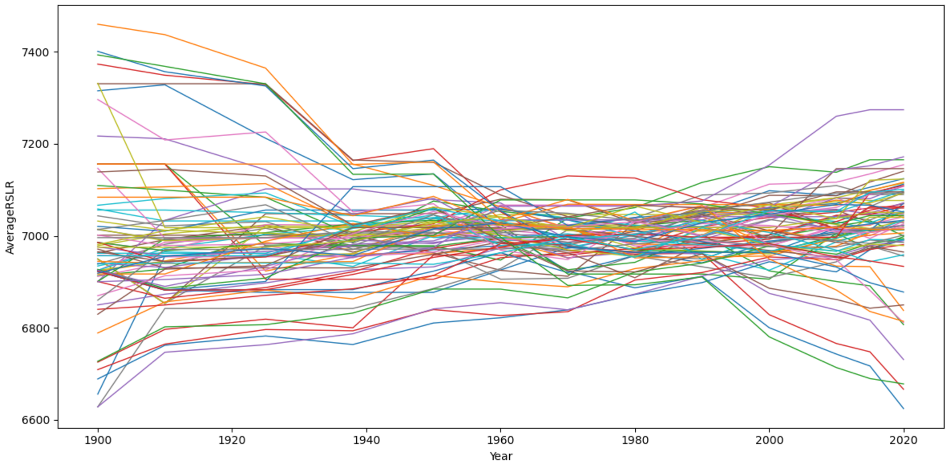

Figure 4.12: Evolution all regions with SLR data after interpolation

Figure 4.13 shows the average relative SLR across all regions after interpolation, measured from 1900 to 2020. The series begins just below 6990 mm in 1900 and displays a gradual increase in the early decades, reaching around 7000 mm by 1925. A noticeable drop occurs in the 1930s, followed by a sharp rise that peaks near 7010 mm in 1950. From that point, the average declines again during the 1960s and reaches its lowest point around 6970 mm in the early 1970s. Starting in the 1980s, a more consistent upward trend emerges, with small fluctuations around 2000, and a marked acceleration in the final two decades. By 2020, the average SLR surpasses 7017 mm, its highest level in the entire period. Overall, the figure shows a long-term rising trend in SLR across regions, especially after 1990, despite earlier periods of fluctuation. This pattern aligns with broader global trends of accelerating SLR due to climate change. The temporary mid-century decline may reflect short-term climate variability, limited data coverage, and changes in ocean heat content or land movements, as discussed Holgate et al. (2013).

Figure 4.13: Average over all regions with SLR data after interpolation

Figure 4.13 shows the change in relative SLR across 82 coastal regions between 1900 and 2020, measured as the difference from the 1900 baseline. The Figure 4.13 also shows that the SLR changes vary considerably across regions. Most regions have experienced a SLR of about 150, mm by 2020, while some regions have seen a decrease. These differences reflect local factors such as land uplift, subsidence, and ocean dynamics (Nováková et al., 2018; Holgate et al., 2013). What is also striking is that most regions have experienced a relatively stable SLR until 1950, after which the SLR experienced by regions started to diverge more. Also, a few regions show continued declines, which is likely due to post-glacial isostatic adjustment. Post-glacial isostatic adjustment is a long-term geological process involving the vertical movement of Earth's crust in response to the past loading and unloading of ice sheets during the last Ice Age (Holgate et al., 2013). As the massive ice sheets melted, the land beneath them, which had been depressed, began to slowly uplift (Holgate et al., 2013). This vertical land movement directly affects relative SLR, which is the local height of the sea relative to the land, as measured by tide gauges (Holgate et al., 2013).

To check if the highest and lowest values are realistic, the three regions with the lowest and highest SLR values in 2020 were examined. The largest decreases occur in northern Sweden and Finland, where land is shifting upward (Holgate et al., 2013). This is consistent with the literature (Chatzivasileiadis et al., 2023). The highest increases are observed in West Flanders (Belgium), Nord (France), and Yorkshire (UK), areas known for significant relative SLR. These findings also aligns with findings by Nicholls et al. (2021), who show that relative SLR is not only driven by climate–induced SLR but is also strongly shaped by subsidence. Particularly in populated delta regions and urban coastal areas, human-induced subsidence, caused by groundwater extraction and infrastructure load, can lead to local relative SLR rates up to four times the global average. Although such rapid subsidence is more prominent in Asian megacities, also in Europe, a large proportion of the population and economic activity are concentrated in coastal areas that face uneven exposure and vulnerability to SLR (Cortés Arbués et al., 2024; Nicholls et al., 2021).

4.2. Exploration of variables

The variables used in the regression $\overline{T}_{i,t}$, $\overline{P}_{i,t}$, and $\text{SLR}_{i,t}$ are included both as linear terms and in squared form to allow for potential non-linear effects. With all variables now defined for each region and year between 1900 and 2015, the climate panel is ready to be merged with the Rosés et al. (2021) GDP dataset and used in the extended specification of the BHM model. Each observation in the dataset is indexed by a NUTS–2 region identifier i and a year t. The Rosés et al. (2021) V6 dataset applies the 2010 version of the NUTS–2 classification, which remains consistent over the full time span from 1900 to 2015. This results in 170 unique region codes that serve as stable spatial identifiers. The temporal structure is defined by a set of benchmark years for which GDP is reported. Since the dependent variable is calculated as the first difference of the logarithm of per-capita GDP, the initial year (1900) cannot generate a growth value and is excluded from the regression sample. The final panel is balanced and covers all regions for each available benchmark period.

4.2.1. Exploratory data analysis

Prior to model estimation, an Exploratory Data Analysis (EDA) is undertaken to understand the characteristics of the data and detect any inconsistencies. First, descriptive statistics (mean, standard deviation, minimum and maximum) are calculated for the dependent and exploratory variables. Per-capita GDP growth $\Delta \ln \text{GDPpc}_{i,t}$, temperature $\overline{T}_{i,t}$, precipitation $\overline{P}_{i,t}$, and the difference in SLR $\Delta\text{SLR}_{i,t}$, as well as their squared terms. These tables can be found in Appendix A A. This step ensures that the variables fall within expected ranges and highlights any data quality issues. Second, pairwise correlations are computed between GDP growth and each climate variable and its square. Scatterplots are generated to visualise these relationships and assess possible non-linearities or multicollinearity. Multicollinearity refers to a situation in regression analysis where two or more independent variables are highly correlated, making it difficult to isolate the individual effect of each variable. This can inflate standard errors and reduce the reliability of coefficient estimates.

Figure 4.14 presents the correlation matrix of the variables used in the analysis, without their squared forms. Temperature and precipitation have a moderate negative correlation of -0.279, suggesting that higher temperatures tend to be associated with lower levels of precipitation. Temperature and SLR differences (dif_slr) show a moderate positive correlation of 0.370, indicating that higher temperatures are generally associated with higher SLR values. Precipitation and SLR are weakly positively correlated, with a coefficient of 0.038. The correlation between temperature and GDP per capita growth is weakly negative at -0.061. Precipitation and GDP per capita growth have a weak positive correlation of 0.076. Finally, SLR and GDP per capita growth are weakly negatively correlated, with a coefficient of -0.028. These low correlation values with GDP per capita growth suggest that none of the climate variables have a strong direct linear relationship with economic growth. This supports their inclusion in the regression model without concern for multicollinearity. The pairwise correlation matrix with inclusion of the squared terms, Figure A.20, is in Appendix A. Since high correlations between those variables are expected by definition. Including these can give a misleading impression of multicollinearity between explanatory variables when, in reality, such correlations are structural and not problematic for model identification if included deliberately for functional form (testing non-linearity).

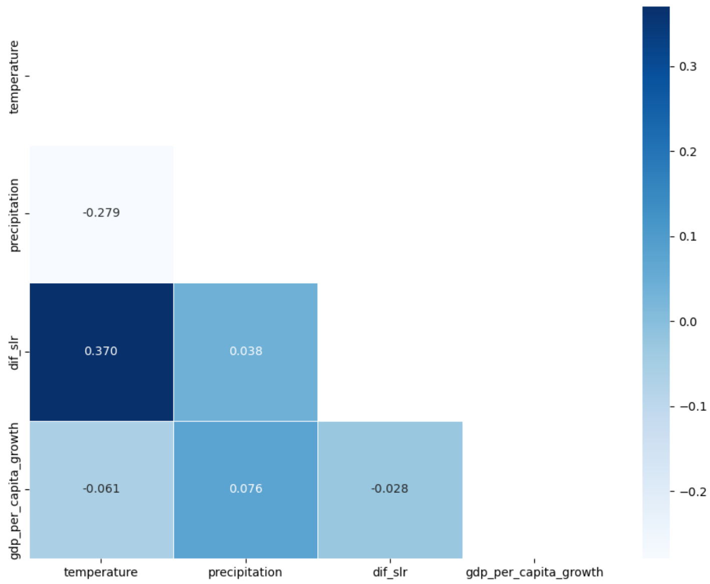

Figure 4.14: Lower triangle of the pairwise correlation matrix between GDP per capita growth and climate variables

Figure 4.15 shows scatter plots of GDP per capita growth against the squared annual average temperature for the benchmark years. Each panel represents one year, allowing patterns over time to be seen. The plots show high variation, with no strong or consistent relationship between temperature squared and GDP growth across the years. Some years, such as 1938 and 1950, show a slight negative pattern, while others, such as 1970 and 1980, suggest a weak positive association. Overall, the figure confirms that while temperature may affect economic growth in complex ways, the direct relationship is not clear and varies across time, supporting the need for econometric analysis to capture potential non-linear and interacting effects.

Figure 4.15: Scatter plots for benchmark years GDP growth and average annual temperature squared

Figure 4.16 shows scatter plots of GDP per capita growth against changes in SLR (dif_slr) for coastal regions across the benchmark years. Each panel represents a benchmark year, making it possible to see patterns over time. In the scatter plots for the years 2000 to 2015, the cloud of points appears to slope downward. This suggests that in these years, regions with higher values of SLR change tend to have lower GDP per capita growth rates. In other words, there is a weak negative association between SLR and economic growth during this period in the dataset. This could indicate that higher relative SLR may begin to have negative impacts on economic performance, for example, through increased flooding, higher protection costs, or disruption of coastal activities. However, the spread of the points shows high variation, and this pattern alone does not prove a causal relationship, underlining the need for econometric analysis to test whether this association holds when controlling for other factors.

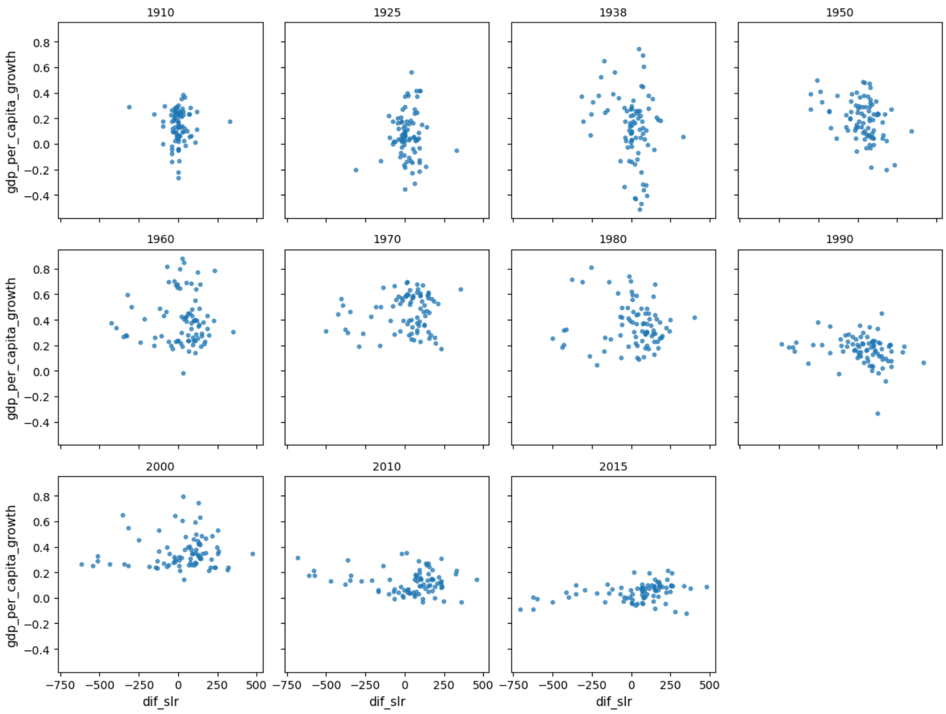

Figure 4.16: Scatter plots for benchmark years GDP growth and the difference in SLR for coastal regions only

4.3. Control variables

The data preparation steps outlined in this chapter ensure a consistent and balanced panel of European NUTS–2 regions, suitable for analysing the long-run effects of climate variables on economic growth. By carefully selecting regions with reliable GDP, population, and climate observations, the resulting dataset captures both temporal dynamics and spatial variation over the period 19002015. To isolate the effects of temperature, precipitation, and SLR on GDP per capita growth, the regression model includes several control variables. First, regional FE account for time-invariant characteristics, such as soil quality, geographic location, and institutional history, which might otherwise confound the analysis. Second, annual precipitation is added to control for short-term variability in water availability, which can impact agriculture and other sectors independently of longer-term climate trends. Third, country-specific linear and quadratic time trends are introduced to absorb gradual changes at the national level, such as demographic transitions, that may influence growth over time (Burke, Hsiang, et al., 2015). Finally, year FE control for global or continent-wide shockssuch as wars, economic crises, or technological advancesthat affect all regions in a given year. Together, these controls ensure that the estimated coefficients on the climate variables reflect region-specific responses, net of broader confounding influences (Burke, Hsiang, et al., 2015).

The next Chapter 5 presents the regression results, applying a reduced-form panel OLS model to the data constructed in this chapter. This analysis aims to estimate the long-run economic impacts of climate variation across European regions over the past century.

5
Results

This Chapter, together with Chapter 6, addresses the central research question: To what extent are the estimated economic impacts of climate, identified in the model by BHM, robust? Subquestion 1 and 2 guide the structure of this Chapter. Subquestion 1 examines whether the concave relationship between temperature and GDP per capita growth remains stable when extending the time frame by 65 years, from 1960-2010 to 1900-2015, using regional data of Europe instead of national level data of the whole world which excludes the the influence of earlier structural shifts such as industrialisation, world wars and post-war recovery (Rosés et al., 2021). By adding 65 more years of economic and climate data, this study assesses whether that concave pattern holds once long-run historical events are included. Subquestion 2 investigates whether an additional climate variable, SLR, alters the temperature-growth relationship as estimated in subquestion 1 or provides further explanatory value.

In the next Chapter 6, subquestion 3 will be discussed. Subquestion 3 will examine what effect certain regions, countries and/or time periods have on the results using a jackknife analysis. To better understand the sole explanatory power of temperature, a regression, using only temperature, will be conducted as well in Chapter 6.

The structure of this chapter is structured around two regression models. For the first subquestion, the model by BHM is extended with the dataset by Rosés et al. (2021) and the dataset by the Climatic Research Unit, University of East Anglia (n.d.). This model will serve as a baseline model to which the output of the models discussed in subquestions 2 and 3 will be compared. For the second subquestion, the model by BHM re-analysed and extended. Therefore, SLR is added to the baseline model to assess whether compound climate effects yield additional insight into regional economic outcomes. Each step is evaluated using model fit statistics, coefficient estimates, and diagnostic tests.

5.1. Subquestion 1: Re-analysis

This section investigates, through a re–analysis, whether the non-linear temperature-GDP relationship identified by BHM remains robust when using an extended time series and a different spatial scale and resolution. Rather than examining the whole world as done by BHM, this research focuses on 170 NUTS–2 regions in Europe. This finer spatial scale reveals how temperature affects growth in a generally cooler climate. Because European temperatures cluster below the global mean, most regions occupy the rising part of the temperature-growth curve. Thus, it is possible to explore how an increase in temperature influences growth in colder climate zones. This setup provides the basis for addressing subquestion 1: How robust is the model by BHM when extending the time frame to 115 years?

To get an understanding of the distribution of the average annual temperature, figure 5.1 highlights the northsouth temperature progression in Europe. Southern Europe, including Spain, Italy and Greece, shows the highest average temperatures, with many regions above 15 °C. Central Europe, including France and Germany, shows moderate temperatures around 10 to 15 °C. Northern Europe, including Scandinavia and the Baltic states, shows lower temperatures, often below 5 °C. This spatial pattern reflects Europe's climate distribution, where temperature decreases with increasing latitude.

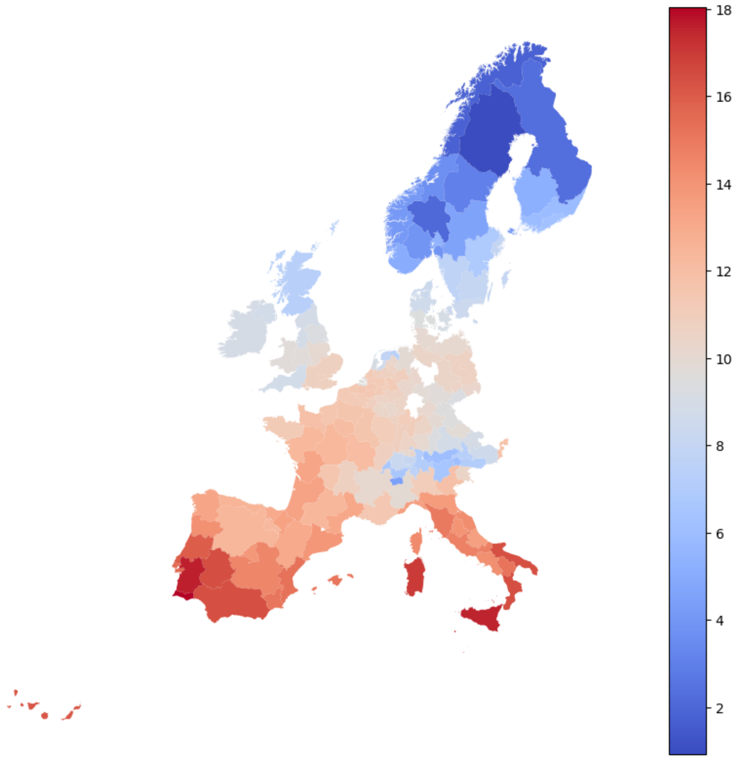

Figure 5.1: Average annual temperature in 2015 of the 170 regions analysed

5.1.1. Annual temperature-growth relationship

This section estimates the relationship between annual temperature, precipitation and GDP per capita growth in benchmark years. The following regression model, discussed in Chapter 3, Equation 3.8, is used on the regional level. Below, in table C.5, the four estimated coefficients from this regression are shown. The overall regression results are shown in Table C.4 in Appendix C. The model has an R^2 of 0.516, meaning that around 52% of the variation in GDP per capita growth is explained by the included variables. The adjusted R^2 is 0.329, which accounts for the number of explanatory terms and shows that a substantial part of the variation is still captured after correcting for model complexity. These values indicate that the model has moderate explanatory power and reinforce the relevance of including climate variables to understand regional economic growth.

Table 5.1: Estimated coefficients from OLS regression

Variable	Coef.	Std. Err.	t	P>|t|	[0.025	0.975]
temperature	0.0737	0.018	4.066	0.000	0.038	0.109
temperature_sq	-0.0031	0.001	-2.805	0.005	-0.005	-0.001
precipitation	-0.0003	0.000	-1.483	0.138	-0.001	0.000092
precipitation_sq	7.235e-08	9.08e-08	0.797	0.425	-1.06e-07	2.5e-07

As can be seen from table C.5, the coefficient of annual mean temperature is positive ($\hat{\beta}_1 = 0.0737$, $p < 0.001$), while the squared term is negative ($\hat{\beta}_2 = -0.0031$, $p = 0.005$), indicating a statistically significant concave relationship between temperature and economic growth. The estimated turning point is calculated as:

$$T^* = -\frac{\hat{\beta}_1}{2\hat{\beta}_2} = -\frac{0.0737}{2 \times (-0.0031)} \approx 11.8\,^\circ\mathrm{C} \tag{5.1}$$

This is close to the optimum of 13.0 °C found by BHM, suggesting strong consistency. Figure 5.2 visualises the quadratic function $h(T) = \hat{\beta}_1 T + \hat{\beta}_2 T^2$, showing the strength and direction of the effect average temperature has on GDP per capita growth. The precipitation terms are not statistically significant at conventional levels. The linear coefficient is small and negative ($\hat{\lambda}_1 = -0.0003$, $p = 0.133$), while the squared term is close to zero ($\hat{\lambda}_2 = 7.444 \times 10^{-8}$, $p = 0.413$). Even though the precipitation terms are not significant the joint significance test shows that the group of climate variables included in this model together have a statistically significant effect on economic growth. Table C.6 in Appendix C shows a F-statistic of 4.49 with a p-value of 0.0018, which indicates that the likelihood of observing such a result if these variables had no effect is very low. This suggests that variation in temperature and precipitation, including their non-linear effects, contributes meaningfully to explaining changes in economic growth across regions. Therefore, including these variables in the regression model is empirically justified.

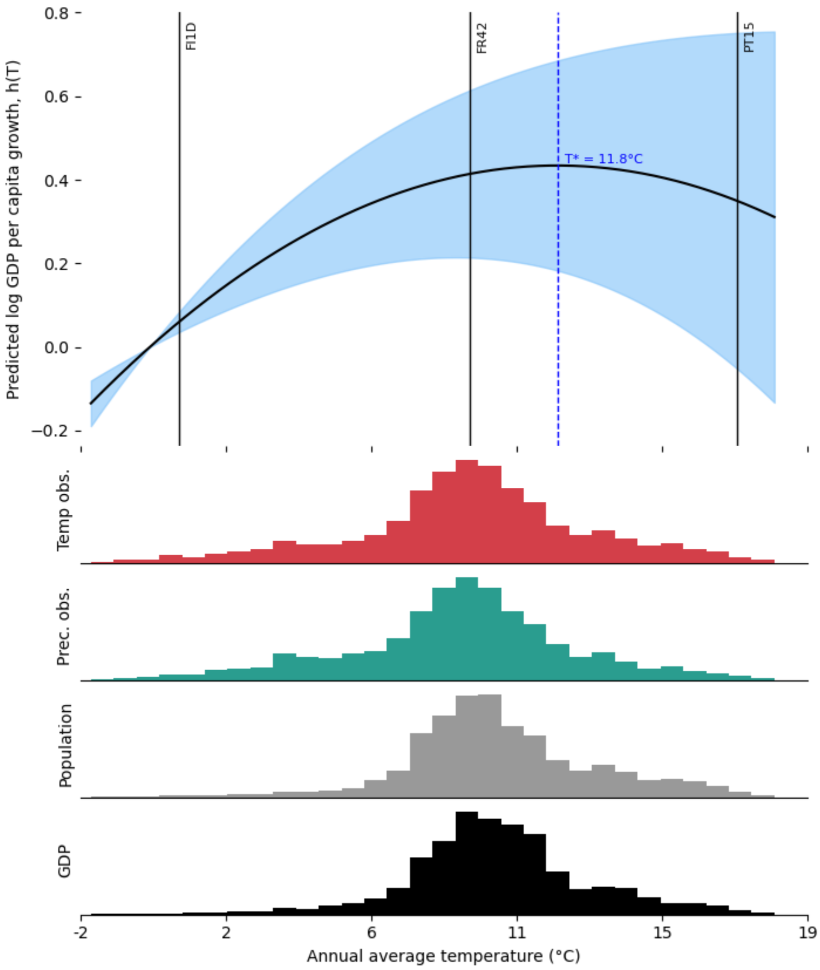

Figure 5.2: Annual average temperature and growth with 90% confidence band and distributions of temperature observation, precipitation observations, population and GDP

As can be seen from Figure 5.2, most observations are concentrated between 5 °C and 15 °C, suggesting that a large share of economic activity and population falls near the estimated optimum. The curve flattens near the peak and declines more sharply at higher temperatures. This indicates that warmer regions may be more vulnerable to temperature increases. Three representative regions are marked for reference: Pohjois- ja Itä-Suomi (FI1D) in Finland lies on the left-hand side of the curve, suggesting potential gains in regional GDP growth if moderate warming occurs. Alsace (FR42) in France is close to the peak, where current temperature conditions are close to optimal, according to this model. In contrast, Algarve (PT15) in Portugal lies beyond the turning point, indicating that the current average annual temperature is already decreasing regional GDP growth and that further warming may reduce growth even further in this already warmer region. These examples aim to illustrate how regional positioning on the temperature-growth curve can shape expected climate impacts. Figure 5.2 also shows four histograms. The first histogram shows that most temperature observations fall between 5 and 15 °C. The second panel presents precipitation-weighted temperature bins, indicating that total precipitation clusters within the same range as temperature. This suggests that temperature and precipitation conditions commonly align within the range where economic growth is highest. This raises the possibility that moving away from this climate range, due to changes in temperature or precipitation, could increase negative effects on regional economic growth. The third panel shows population-weighted observations, confirming that a large share of the population is exposed to temperatures near the estimated optimum. The fourth panel reveals that most GDP is also generated in this range. Together, these distributions underline that both climatic exposure and economic activity are concentrated near the temperature range where the model predicts the highest growth.

5.1.2. Robustness of the annual temperature-growth relationship over 115 years

The model developed by BHM estimates a non–linear relationship between temperature and economic growth, with a statistically significant positive linear temperature coefficient of 0.0127 and a negative squared term of −0.0005 (see column 1 of their Extended Data Table 1). Applying a comparable FE panel model to regional data in Europe over an extended time period (1900–2015), this study finds a similar concave relationship, but with substantially larger effect sizes: the linear temperature coefficient is 0.0737, and the squared term is −0.0031, as shown in Table C.5. These coefficients are roughly six times larger in absolute terms than those estimated by BHM. This suggests that regional economic output in Europe may be more sensitive to temperature variation than global averages imply. But also the greater variation at the regional level compared to the national level could help explain the larger estimated coefficients in this study are much larger than those of BHM. While the precipitation effects differ in sign and significance, the core temperature–growth relationship remains intact. A full overview of the coefficient comparison is presented in Table C.3 in Appendix C. The regression estimated in this section confirms a statistically significant concave effect of annual temperature on economic growth, with an estimated optimum of 11.8 °C. This relationship remains robust when extending the original time frame of BHM from roughly 50 years to 115 years. The similarity in shape, turning point, and significance, despite differences in magnitude, indicates that the model's findings are not sensitive to a time period extension or spatial context. Therefore, in response to Subquestion 1, the estimates of BHM are robust and hold over a much longer historical period and another spatial context, strengthening the claim that temperature has a non-linear and economically meaningful effect on growth.

According to Clemens (2017), differences in estimated coefficients across studies should not be interpreted as failed replications, especially when the underlying structure of the model is preserved. In this case, the observed differences are consistent with expectations given the more restricted geographical scope, longer historical horizon, and cooler average climate in the European panel. That the same non-linear functional form emerges in both settings supports the structural robustness of the BHM model, while highlighting regional variation in the sensitivity of economic growth to temperature.

To further assess the robustness of the estimated parameters, a jackknife analysis is performed in the next chapter (Chapter 6). By systematically re–estimating the model while excluding one region, one country, or one benchmark year at a time, this method evaluates the stability of the coefficient estimates. This method shows whether certain years or locations have a strong influence on the size or shape of the estimated effect. While the similarity in the functional form already indicates a degree of robustness, the jackknife analysis provides a more detailed assessment. If the coefficients change a lot when some years or regions are left out, it may mean that the model depends too much on those specific observations or on shifts in the data over time. The jackknife analysis, therefore, looks beyond just the curve and gives a broader view of how stable the model is. The outcomes of the jackknife analysis are discussed in Chapter 6, Section 6.1.

5.2. Subquestion 2: Re-analysis and extension

This section investigates, through a re-analysis and extension of the model by BHM, if there are compound effects of temperature, precipitation and SLR on economic outcomes. But also, to what extent these compounding effects alter the climate-economy relationship identified in the model by BHM? Therefore, this section evaluates how the inclusion of an additional climate variable affects the estimated temperature–growth relationship.

Figure 5.3 displays the spatial variation in the difference in SLR exposure across European NUTS 2 regions in 2015. The values represent the difference in average SLR compared to the baseline year 1900, capturing over a century of SLR. Regions shaded in red experienced an increase in SLR relative to 1900. In the dataset used in this study there are, in the benchmark year 2015, 57 regions experiencing an increase in SLR. The 25 regions in blue saw a decrease in SLR in 2015. Negative SLR can occur in regions where the land rises faster than the sea. This happens due to natural processes such as glacial rebound or tectonic uplift (Chatzivasileiadis et al., 2023). White regions indicate no recorded difference in SLR over this period. This is because they are non-coastal and thus not experiencing SLR directly. Figure 5.3 highlights clear geographic patterns: northern Europe, particularly Finland and parts of Sweden, shows relatively low or even negative differences, whereas northwestern and coastal regions of Western Europe experienced some of the largest increases. This spatial variation in SLR change is used as input for estimating the potential economic relevance of SLR in later parts of the analysis.

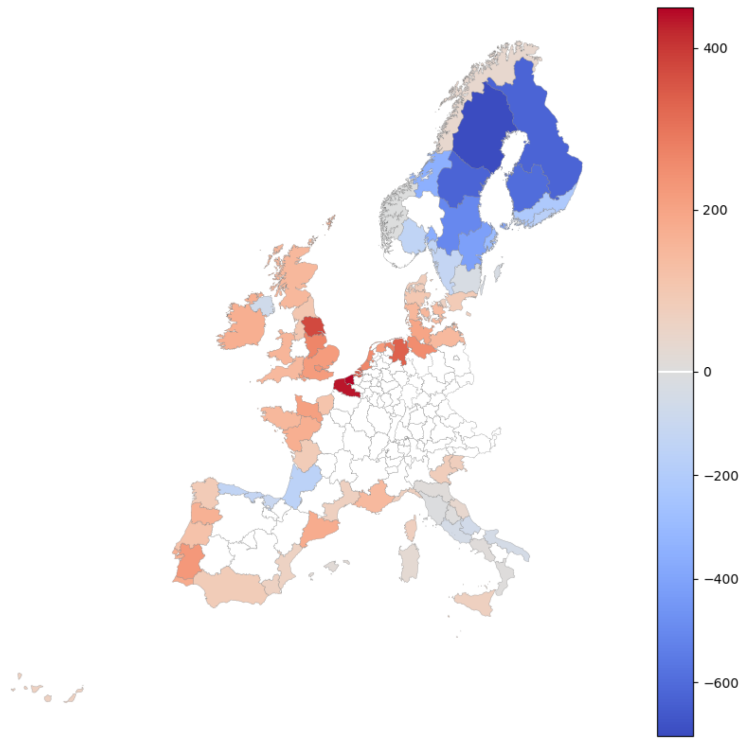

Figure 5.3: Estimated difference in SLR in 2015 by NUTS 2 region. Red regions experience positive SLR, while blue regions experience negative SLR. White, inland regions experience zero SLR.

5.2.1. Extending the model with SLR

This step extends the model by incorporating SLR in addition to temperature and precipitation. This step tests whether compound climate effects further influence the climate–growth relationship beyond those already captured by temperature and precipitation alone, as done in the section above (section 5.1).

Table C.10 shows that the temperature coefficients remain consistent with those estimated in Section 5.1. The linear term is positive ($\hat{\beta}_1 = 0.0727$, $p < 0.001$), and the squared term is negative ($\hat{\beta}_2 = -0.0030$, $p = 0.014$). Compared to the optimum found in the first regression in section 5.1 and visualised in 5.2 at ($T^* \approx 11.8\,°\mathrm{C}$), the estimated turning point in this model is similar to this regression, with SLR included, with an optimum at ($T^* \approx 12.0\,°\mathrm{C}$), suggesting that the inclusion of SLR does not substantially shift the optimal temperature for growth.

Table 5.2: Estimated coefficients for temperature, precipitation, and Sea level rise (clustered SE)

Variable	**coef**	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
Temperature	0.0727	0.021	3.547	0.000	0.033	0.113
Temperature^2	−0.0030	0.001	−2.455	0.014	−0.005	−0.001
Precipitation	−0.0003	0.000	−1.532	0.126	−0.001	8.24×10^{-5}
Precipitation^2	8.16×10^{-8}	9.03×10^{-8}	0.904	0.366	-9.53×10^{-8}	2.58×10^{-7}
Sea level rise	−0.0007	0.000	−2.326	0.020	−0.001	−0.000
Sea level rise^2	-1.199×10^{-6}	6.96×10^{-7}	−1.723	0.085	-2.56×10^{-6}	1.65×10^{-7}

As can be seen from Table C.10 precipitation terms remain statistically insignificant, with only a small change in their magnitude, supporting the conclusion that precipitation does not have a strong direct influence on growth in this dataset. In contrast, the the additional variable, SLR, provides additional explanatory content: the linear effect is negative and statistically significant (-0.0007, $p = 0.020$), while the squared term is very small but negative and only marginally significant (-1.199×10^{-6}, $p = 0.085$), suggesting a non-linear and increasingly negative effect of SLR at higher exposures. The joint significance test, in Table C.11 in Appendix C, returns an F-statistic of 5.08 with a p-value of 8.07×10^{-5}. This result indicates strong evidence that these climate variables jointly contribute to explaining variation in economic growth. The low p-value suggests that it is highly unlikely to observe such a result if none of the variables had an effect. Compared to the earlier test on only temperature and precipitation (Table C.6 in Appendix C), which yielded an F-statistic of 4.49 and a p-value of 0.0018, the inclusion of SLR leads to a slightly higher F-statistic and a more significant result. This indicates that adding SLR to the regression improves the overall explanatory power of the model. This implies additive effects, meaning that the estimated impact of each climate variable (temperature, precipitation, and SLR) is evaluated independently of the others. There are no interaction terms included, so the total effect on economic growth is the sum of the individual contributions of each variable.

The models explanatory strength is further confirmed by an R^2 of 0.518 and an adjusted R^2 of 0.330 (Table C.9 in Appendix C), indicating that about one-third of the variation in economic growth is explained after adjusting for the number of predictors. This represents a slight improvement compared to the model without SLR, which achieved an R^2 of 0.516 and an adjusted R^2 of 0.329. While the difference in model fit is modest, the statistical significance and sign of the SLR coefficients suggest that its inclusion is empirically relevant. Therefore, including temperature, precipitation, and SLR together is justified and strengthens the models ability to capture the economic effects of climate.

Figure 5.4 shows the estimated non-linear relationship between temperature and GDP per capita growth, now controlling for precipitation and SLR. The concave shape remains visible, with a peak at 12.0 °C, slightly higher than the estimated optimum in the model without SLR (11.8 °C). This implies that when the model separates the effect of SLR from temperature, the estimated tolerance to temperature increases a bit. The four histograms beneath the curve provide additional context. Their interpretation remains the same as in Section 5.1.

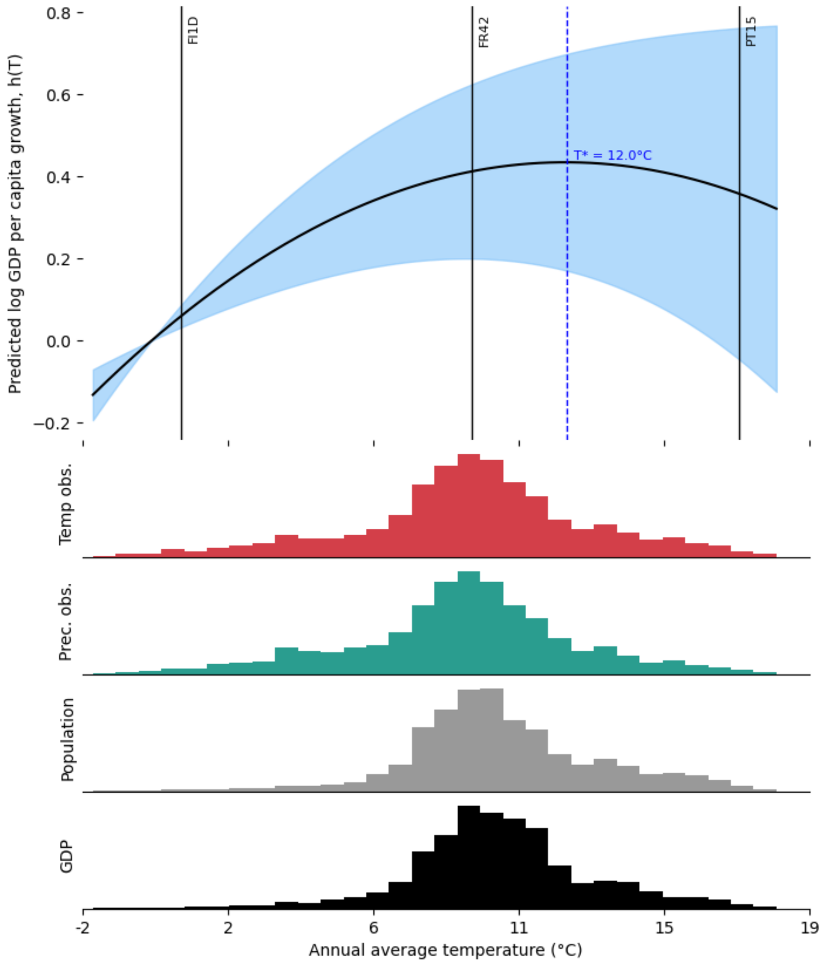

Figure 5.4: Estimated temperature–growth relationship of temperature, precipitation and SLR on GDP per capita growth.

The first subplot in Figure 5.5 (below) shows the estimated effect of annual average temperature on GDP per capita growth. This visualises the non-linear, quadratic temperature-growth relationship found in the model, aligning with the concave shape found by BHM.

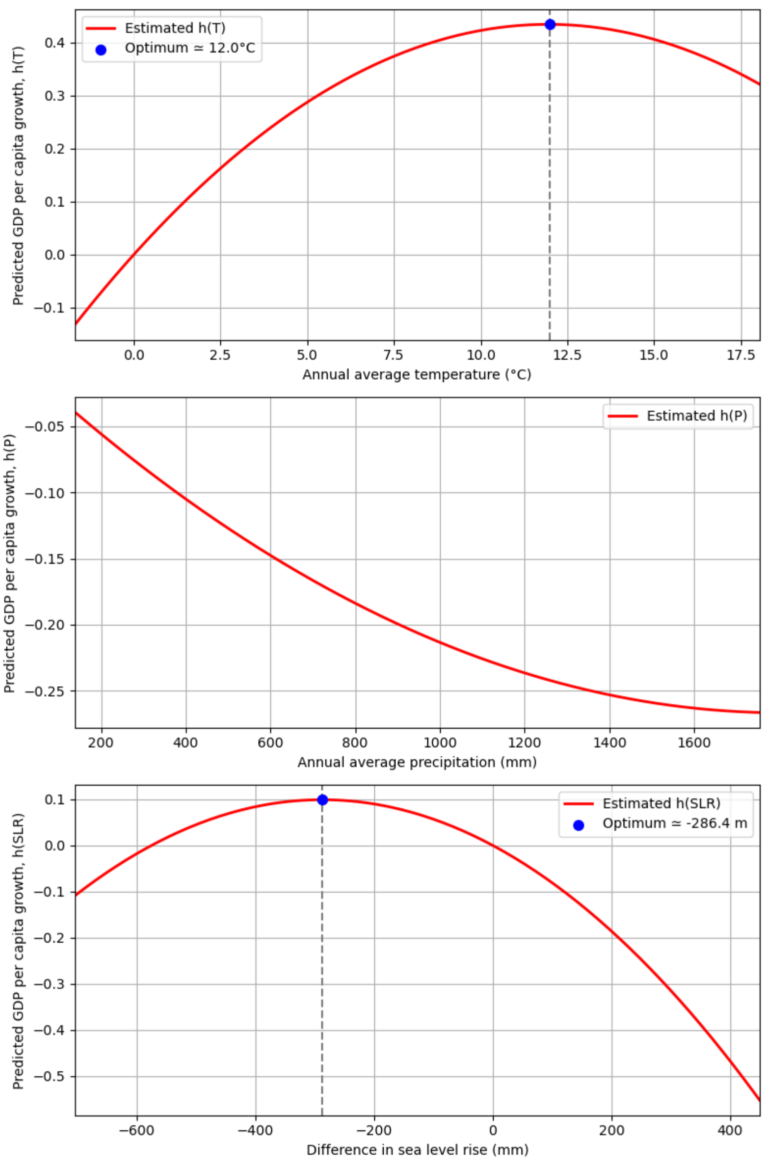

Figure 5.5: Estimated temperature–growth relationship of temperature, precipitation and SLR on GDP per capita growth.

The second subplot in figure 5.5 (above) shows the estimated effect of precipitation on GDP per capita growth. The relationship is nearly linear across the observed range, indicating that higher levels of precipitation are linked to lower growth rates in this dataset. But this plot needs to be interpreted which caution since the estimated coefficient for precipitation is -0.0003 and is not statistically significant at the 10% level ($p = 0.126$), while the coefficient for the squared term is positive 8.16×10^{-8} but not significant at all ($p = 0.366$). This means that the curve plotted in the second subplot does not exhibit a statistically significant non-linear relationship between precipitation and growth, consistent with the shape of the figure, which shows a smooth, downward-sloping curve without a clear turning point. Although the squared and linear precipitation terms are not individually significant, the joint significance test indicates that precipitation still contributes explanatory power when included. This suggests that while the effect of precipitation on its own may be modest or diffuse, it plays a role in shaping economic outcomes when considered alongside other climate variables. As shown in Chapter 6, Section 6.2, the regression using only temperature as an explanatory variable results in a much higher estimated optimum. This highlights the importance of including precipitation in the analysis, as it helps prevent overestimating the effect of temperature on economic growth.

The third subplot in Figure 5.5 shows the estimated effect of the difference in SLR since 1900 on GDP per capita growth. The relationship is non-linear and concave, with an identified optimum around −286 mm. This implies that regions which experienced a relative fall in SLR since 1900 tend to exhibit higher regional GDP growth rates. Growth rates initially rise with increasing SLR, but decline beyond the estimated optimum, indicating that higher relative SLR negatively impacts economic growth. While the overall magnitude of SLR on regional growth is smaller than the temperature effect, the direction and curvature are consistent with the literature, which finds that SLR increase economic risks (Chatzivasileiadis et al., 2023).

In addition, Figure 5.6 shows the estimated relationship between SLR and GDP per capita growth across European regions. The black line represents the predicted growth effect, which is calculated as a quadratic function of relative sea-level rise following the approach in BHM and Kotz et al. (2024). The blue shaded area shows the 90 % confidence interval around the estimate.

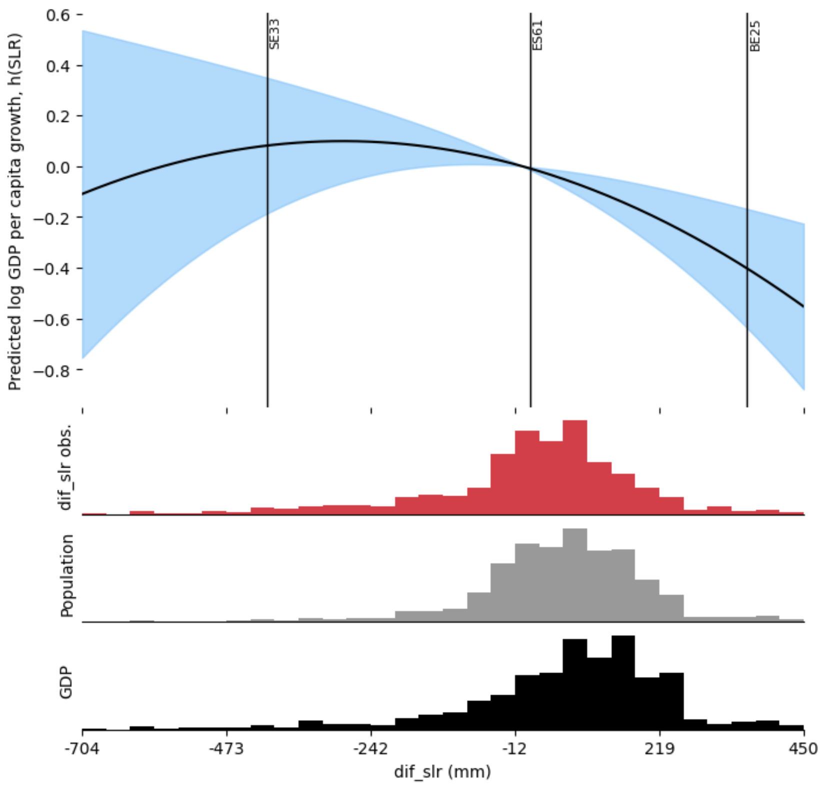

Figure 5.6: Estimated temperature–growth relationship of temperature, precipitation and SLR on GDP per capita growth.

Three regions are included in the figure to illustrate how different parts of the curve reflect real-world conditions. Övre Norrland (SE33) in northern Sweden has a relative SLR of approximately −408 mm and lies on the far left of the curve, where growth is relatively high. Andalucía (ES61) in southern Spain is located near the centre, with a value close to 12 mm, while West-Vlaanderen (BE25) in Belgium has a relative SLR of about 360 mm, positioning it on the right side of the curve, where the predicted growth effect declines. The lower panels of the figure show how observations of relative SLR are distributed across regions, weighted by SLR, population, and GDP. Most regions fall between 0 and 200 mm, suggesting that a significant share of economic activity and population is exposed to SLR.

5.2.2. Compound effects of SLR

Overall, the findings support the robustness of the non-linear temperature effect, while showing that precipitation lacks a clear threshold beyond which economic impacts become significant. In contrast, SLR has a statistically significant but smaller effect. The model reaffirms the central role of temperature in shaping long–term economic outcomes. The concave relationship remains stable across model specifications, now also confirmed when controlling for precipitation and SLR. SLR emerges as a secondary but directionally consistent factor. Although its overall magnitude is smaller, the concave pattern and statistical significance of SLR suggest that long-term exposure to SLR is associated with declining economic performance.

These findings directly answer Subquestion 2, which asked whether the relationship identified by BHM holds when considering additional climate variables. The results show that the non-linear temperature-growth relationship is not only preserved but strengthened when extended to a broader specification. The inclusion of precipitation and SLR does not fundamentally alter the curve, but it does improve the models credibility. While the inclusion of SLR leads to only a marginal increase in the model's explanatory power, as reflected in a minimal change in the R^2, its coefficients are statistically significant and directionally consistent with the expectation that higher SLR imposes economic costs. This suggests that, although SLR explains only a small portion of the overall variation in regional growth, it captures a relevant dimension of climate-related pressure. The limited contribution to the R^2 highlights that temperature remains the dominant driver, but the significance of SLR coefficients indicates that its economic impact, though more subtle, is detectable in the long-run growth data.

6
Sensitivity Analysis

In addition to Chapter 5, multiple sensitivity analyses are performed for the second subquestion to assess the stability of the estimated relationship across time and space. The jackknife method is used to sequentially exclude one region, one country and one benchmark year from the regression to test whether the results depend on any specific subset of the data. These procedures help ensure that findings are not driven by spatial outliers or by benchmark years in the historical data. Lastly, the baseline model, as conducted for subquestion 1 in the Chapter 5, Section 5.1, is re-run, but instead with temperature and precipitation as exploratory variables, only temperature is included. This last model is conducted to see the sensitivity of the estimated temperature effect when precipitation is excluded, and whether temperature alone captures some of the variation otherwise attributed to precipitation.

6.1. Subquestion 3: Robustness and sensitivity to spatial and temporal exclusions

This section outlines the results of multiple jackknife analyses to test the robustness of the temperature-growth relationship proposed by BHM, with a focus on subnational, national and temporal variation. Within countries, differences in topography, infrastructure, and adaptive capacity can lead to different responses to climatic variables and could, therefore, lead to a different impact on economic growth (Rosés et al., 2021). To examine how these differences affect the results, three complementary jackknife analyses are conducted. Each analysis reruns the model multiple times, excluding one region, one country or one year at a time, to assess how sensitive the estimated climategrowth relationship is to specific locations or time periods (Kahn et al., 2021).

First, a spatial regional–level jackknife is performed by sequentially removing each NUTS-2 region from the estimation. This procedure will highlight the influence of individual regions on the overall temperature-growth curve. Given that climate exposure varies across regions, this approach reveals whether the overall non-linear shape is disproportionately influenced by any particular region. A second spatial country–level jackknife is thereafter conducted by sequentially removing each country from the estimation to see the influence of one country on the overall temperature-growth curve. Since a single region may have only a limited effect, this country-level jackknife helps to assess whether national patterns or country–specific characteristics are driving the observed relationship. Lastly, a temporal jackknife removes one benchmark year at a time to test whether the results are robust across historical periods.

6.1.1. Jackknife visualisation and interpretation

To visualise these robustness checks, Figures 6.1, 6.2 and 6.3 plot the estimated temperature-growth curves from the spatial and temporal jackknife procedures, using the observed range of annual temperatures in the dataset.

Spatial jackknife (regional-level)

Table C.18 presents the results of the spatial jackknife analysis. The table shows the average value of each coefficient across all runs (Mean), the variation across those runs (Std. dev.), and the results of a simple statistical test to check if the average is meaningfully different from zero (t-stat and p-value). The t-statistic is calculated by dividing the mean by its standard deviation. It shows how many times the average estimate is larger than the variation across the runs. A higher t-value means that the estimate is more precisely measured. In this table, all variables show very large t-statistics and p-values close to zero, meaning that the effects are statistically significant. The temperature coefficient is positive (0.0738), and the squared term is negative (−0.00313), confirming a concave relationship between temperature and economic growth. The small standard deviations and high t-statistics show that this non-linear pattern is very stable across regions. The precipitation coefficients are much smaller in size but also highly significant. This may seem surprising, as precipitation is not significant in the main regression model as mentioned in Chapter 5, Section 5.1. However, the jackknife test reflects consistency rather than exploratory power of the variables. The small standard deviations indicate that the precipitation terms change very when a single region is excluded, which leads to large t-statistics even if their size is small. This suggests that while precipitation does not have a strong effect on economic growth in the baseline regression model, its estimated role in the model is stable and not dependent on a few influential regions.

Table 6.1: Spatial jackknife results for climate coefficients

Variable	Mean	Std. dev.	t-stat	p-value
Temperature	0.0738	0.0013	58.64	0.000
Temperature^2	-0.00313	0.000077	-40.76	0.000
Precipitation	-0.00029	0.000013	-22.28	0.000
Precipitation^2	0.0000000745	0.0000000062	11.97	0.000

Figure 6.1 presents the results of the spatial regional–level jackknife, showing the robustness of the estimated temperature-growth relationship when omitting each of the five most influential regions in turn. The characteristic concave shape of the temperature-growth curve remains consistent across all specifications, indicating that the non-linear relationship is not driven by any single region. Slight upward deviations are visible when omitting SE33 (light blue) and F1D (orange), while smaller shifts appear when excluding NO02 (red), NO05 (purple), and ITG1 (green). These differences are most apparent at the higher end of the temperature distribution, where the curves diverge modestly above 12°C.

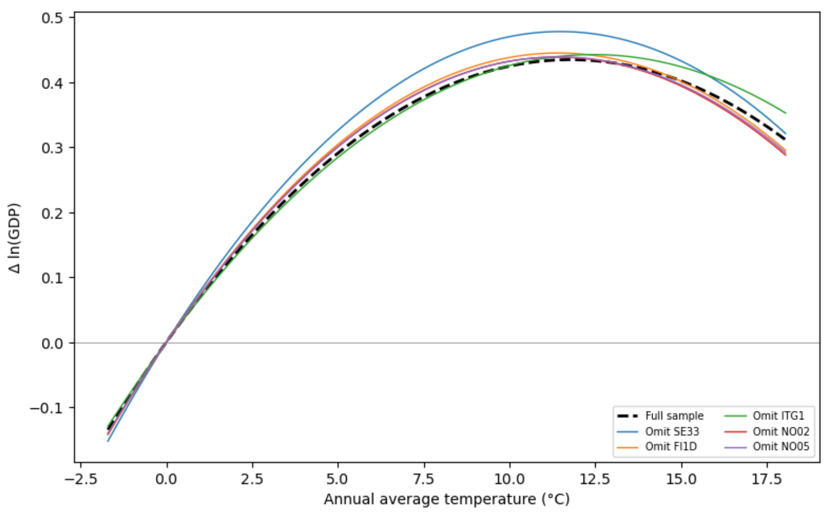

Figure 6.1: Spatial jackknife estimates of the temperature-growth curve. Each coloured line shows the result when omitting one NUTS-2 region. The dashed black line represents the full-sample estimate.

The similarity of the curves in Figure 6.1 across the observed temperature range demonstrates the robustness of the estimated non-linear temperature-growth relationship to the exclusion of individual influential regions. Table 6.2 presents the estimated optimal temperatures when excluding each of the five most influential NUTS regions from the sample. The results demonstrate that the estimated optimum clusters closely around the baseline optimum, as mentioned in Chapter 5 Section 5.1 of 11.8°C.

Table 6.2: Optimal temperature after excluding influential NUTS regions

Region	Optimal Temperature (°C)	Difference from 11.8 (°C)
NO02	11.38	-0.42
FI1D	11.42	-0.38
NO05	11.42	-0.38
SE33	11.47	-0.33
ITG1	12.44	+0.64

Excluding NO02, FI1D, NO05, and SE33 shifts the estimated optimum downward by 0.33 to 0.42°C, indicating that omitting these colder northern regions shifts the optimum towards an even cooler temperature. In contrast, excluding ITG1, a warmer southern region, increases the estimated optimum by 0.64°C. This shift reflects the influence of ITG1 on the upper end of the temperature distribution. But this shift needs to be interpreted with caution since there are fewer regions that occupy higher temperature ranges in Europe. It could mean that the optimal temperature estimate is more sensitive to the inclusion or exclusion of warmer regions due to the smaller number of observations in this part of the distribution, while the lower end is anchored by many regions with cooler climates, causing the optimum to shift downward when colder regions are excluded.

These findings align with the overall, full sample, curve observed in Figure 6.1, where curves diverge modestly above 11.8°C depending on which region is excluded. This narrow range suggests that the identified concave relationship and its estimated optimum are not completely driven by any single influential region. However, before drawing strong conclusions about robustness, it is important to consider that the sample includes 170 regions. As a result, the exclusion of a single region may have only a limited influence on the overall curve. At the same time, the fact that there are still visible deviations suggests that the estimates are not fully stable and may be sensitive even to the removal of just one region. This highlights the importance of cautious interpretation and the need to complement the regional jackknife with a country jackknife analysis.

Spatial jackknife (country-level)

Table C.20 shows the results of the country–level spatial jackknife analysis. This checks whether the estimated climate effects are heavily influenced by any single country. The estimated temperature coefficients remain, regarding the regional–level jackknife analysis, statistically significant, with a positive linear term (mean = 0.0739) and a negative squared term (mean = −0.0031), indicating a concave relationship between temperature and economic growth. However, compared to the regional–level jackknife, the standard deviations are larger and the t-statistics are notably lower. This suggests more variation in the estimated coefficients when excluding entire countries. The precipitation terms are also statistically significant, but their standard deviations are higher than in the regional analysis. These results imply that the estimated effects of climate variables are more sensitive at the county–level than at the regional–level. While the direction and significance of the effects remain stable, the country–level jackknife shows greater variability and lower precision in the estimates. This is expected, as excluding one country removes more data from the model than excluding a single region, making the results more sensitive to country–level omissions.

Table 6.3: Country jackknife summary statistics for climate coefficients

Variable	Mean	Std. Dev.	t-statistic	p-value
Temperature	0.0739	0.0089	8.33	0.0000
Temperature^2	-0.0031	0.0006	-5.26	0.0000
Precipitation	-0.0003	0.0001	-4.19	0.0000
Precipitation^2	7.23e-08	2.16e-08	3.35	0.0008

Figure 6.2 shows the jackknife estimates of the temperature-growth curve, where each line represents the fitted quadratic relationship when one of the five most influential countries is excluded from the regression. The black dashed line shows the full-sample estimate. The other lines illustrate how the

shape and peak of the curve change when leaving out Spain (ES, blue), the United Kingdom (UK, orange), Sweden (SE, green), Norway (NO, red), or Italy (IT, purple).

Excluding Spain or the United Kingdom causes the curve to shift rightward, increasing the estimated optimal temperature to around 13.6 °C and 13.9 °C, respectively. In contrast, excluding Norway significantly lowers the optimum to 9.5 °C, suggesting that this country's observations pull the curve upward in the warmer range. Excluding Sweden or Italy produces more moderate shifts. These results suggest that although the overall concave shape remains visible, both the steepness of the curve and the estimated turning point vary depending on which country is excluded. This indicates that the temperature-growth relationship is sensitive to national contributions.

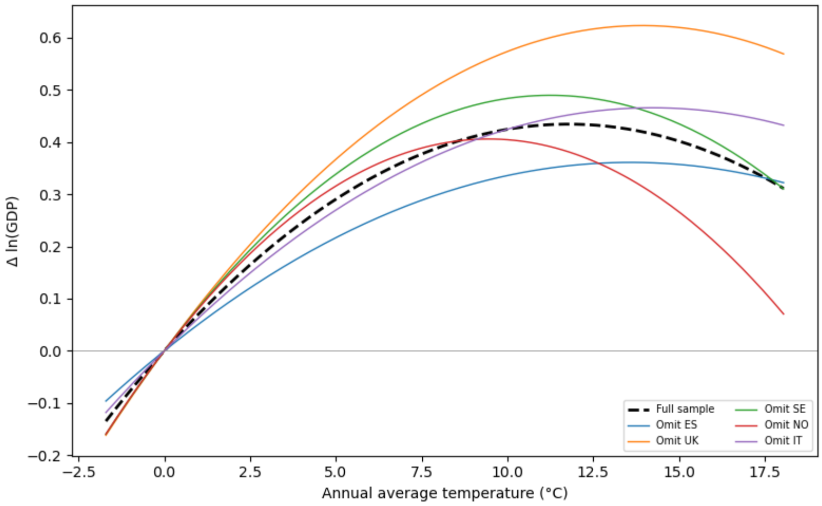

Figure 6.2: Spatial jackknife estimates of the temperature-growth curve. Each coloured line shows the result when omitting one country. The dashed black line represents the full-sample estimate.

Table 6.4 summarises the optimal temperature estimates for each specification and shows the deviation from the full-sample value:

Table 6.4: Optimal temperature after excluding influential countries

Region	Optimal Temperature (°C)	Difference from 11.8 (°C)
Omit ES	13.6	+1.8
Omit UK	13.9	+2.1
Omit SE	11.2	−0.6
Omit NO	9.5	−2.3
Omit IT	14.2	+2.4

The regional and country–level jackknife results reveal both common patterns and notable differences. In both cases, the concave temperature–growth relationship remains visible, indicating that the overall shape is not driven by a single location. However, excluding an entire country has a much stronger influence on the estimated optimum than omitting one region. Regional exclusions shift the optimum by less than ±0.6°C, while country-level exclusions lead to changes of up to +2.4°C or −2.3°C. This suggests that national data, especially from countries with broader temperature ranges, can affect the position and steepness of the curve. Beyond the curve itself, the jackknife results also provide insight into the stability of the estimated coefficients. Although the functional form remains similar, the size of the coefficients changes more when a full country is excluded than when a region is left out. While no single country dominates the results, the precision of the estimates varies with the sample. These findings suggest that the concave relationship is robust, but the estimated coefficients are not fully stable across different country samples.

While the spatial jackknife analyses focus on variation across regions and countries, it is also important to test whether the estimated climate effects are stable over time. The temporal jackknife allows for an assessment of how sensitive the results are to specific benchmark years and is discussed in the following subsection (6.1.1).

Temporal jackknife

Table C.19 reports the results of the temporal jackknife analysis, where the model is re-estimated while omitting one benchmark year at a time. The temperature coefficients remain statistically significant and continue to show a concave relationship, with a positive linear term (mean = 0.0762) and a negative squared term (mean = −0.00320). Compared to the regional–level jackknife (Table C.18), the standard deviations are noticeably larger, but only moderately higher than in the country–level jackknife (Table C.20). This results in lower t-statistics, though the estimates remain statistically significant. The precipitation term is significant, but the squared term is not, with a p-value above conventional threshold ($p = 0.312$). This suggests that the evidence for a non-linear precipitation effect is weaker over time than across space. Compared to the regional–level jackknife, which shows high stability and strong significance for all coefficients, and the country-level jackknife, which shows moderate variation, the temporal jackknife reveals greater sensitivity of the model to specific years. This implies that benchmark–year–specific events may influence the magnitude and precision of the estimated climate effects more than regional or national exclusions do.

Table 6.5: Temporal jackknife results for climate coefficients

Variable	Mean	Std. dev.	t-stat	p-value
Temperature	0.0762	0.0330	2.31	0.021
Temperature^2	-0.00320	0.00146	-2.19	0.028
Precipitation	-0.00029	0.00014	-2.10	0.036
Precipitation^2	0.0000000722	0.0000000714	1.01	0.312

Figure 6.3 presents the results of the temporal jackknife. Each line shows the estimated temperature-growth curve when one year is excluded from the sample. The dashed black line represents the full-sample estimate. This analysis tests how sensitive the model is to the influence of individual years. Figure 6.3 shows that the concave shape of the temperature-growth curve remains visible across almost all specifications, but the height and steepness of the curve vary considerably depending on the year omitted. Excluding the year 1938 (pink) has the largest effect. The curve becomes flatter and even turns negative beyond approximately 7°C, suggesting that this year strongly pulls the curve upward in the full sample. Since this study relies on benchmark data, it is not possible to explore in detail what drives the strong influence of 1938. However, this year falls just before the outbreak of the Second World War. This may have introduced irregularities in economic growth or climate sensitivity that make 1938 stand out. Still, without more detailed temporal data, the exact reason for its outsized influence remains unclear. In contrast, excluding the year 2010 (dark blue) produces the steepest and highest curve. The peak is higher and occurs earlier, with a stronger drop after the turning point. This suggests that 2010 lowers the estimated temperature effect when included in the full sample. One possible explanation is the aftermath of the 2008 financial crisis, which may have depressed growth rates in 2010 across multiple regions. Excluding the year 1925 (orange) results in a nearly flat curve, indicating that this year contributes substantially to the concavity of the full-sample estimate. A similar effect is seen when excluding 1950 (green), where the curvature also largely disappears, and the curve becomes almost linear. These cases show that some years have a stabilising effect on the turning point and overall shape of the relationship. The exclusion of 1910 (purple) and 1990 (light green) leads to moderate upward shifts, with a slightly less steep curve and higher peak. This suggests that these years moderate the estimated growth effect of temperature when included. In contrast, omitting 1960 (light blue), 1970 (brown), 1980 (grey), 2000 (cyan), and 2015 (red) results in curves with similar curvature to the full sample, but they lie consistently lower across the temperature range. This suggests that these years influence the overall level of the relationship more than its shape.

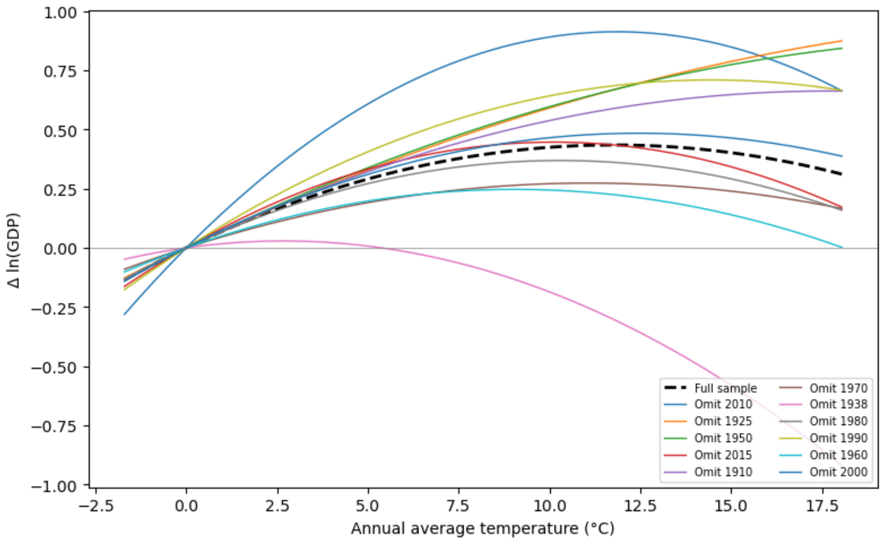

Figure 6.3: Temporal jackknife estimates of the temperature-growth curve. Each coloured line shows the result when omitting one benchmark year. The dashed black line represents the full-sample estimate.

The results of the temporal jackknife reveal that while the concave shape of the temperature-growth relationship remains broadly consistent, the estimated coefficients are not stable. Figure 6.3 shows that omitting individual years can substantially alter the height and steepness of the curve. These deviations indicate that certain years, such as 1938 and 2010, exert disproportionate influence on the estimated temperature effect. This sensitivity is reflected in the standard deviation of the temperature coefficient, which increases from approximately 0.0013 in the spatial regional–level jackknife, given in Table C.18 to over 0.0033 in the temporal one, given in Table C.19, a difference of more than 2.5 times. Similarly, the standard deviation of the squared temperature coefficient is almost twenty times larger. This means that the estimated shape of the temperature-growth curve is highly dependent on the inclusion of specific years.

The temporal jackknife shows that the concave shape of the temperature-growth curve is not preserved in all specifications. While many of the curves retain the general non-linear form, others, such as those excluding 1925 (orange) and 1950 (green), appear almost linear across the observed temperature range. This suggests that certain benchmark years play a key role in shaping the curvature of the full-sample estimate. Moreover, the large variation in curve height and steepness across years points to substantial sensitivity in the estimated coefficients. The standard deviations of the temperature and squared temperature coefficients are considerably larger than in the spatial jackknife, confirming that the model is more sensitive to temporal variation. These findings highlight that while a non-linear relationship emerges in many cases, the estimated parameters do not appear to be stable, and specific benchmark years may strongly influence the result. It does make one wonder, since one benchmark year can make such an impact, if the estimated optimum is valid. This underscores the impact of the model and the data choices, which call into question the relevance and universality of a single estimated optimal temperature (Newell et al., 2021). Importantly, these differences cannot be explained by average growth shocks in those years alone, as the model already includes year FE. These control for the mean shift in growth that may result from broad historical events such as wars or recessions. The observed variation, therefore, reflects how individual years influence the shape of the estimated relationship between temperature and economic growth, rather than the overall level of growth.

Table 6.6 presents the estimated optimal temperature after sequentially omitting each benchmark year from the sample. Deviations from this benchmark are shown in the third column.

Table 6.6: Optimal temperature after excluding influential years

Year	Optimal Temperature (°C)	Difference from 11.8 (°C)
1910	17.66	+5.86
1925	27.33	+15.53
1938	2.68	-9.12
1950	23.58	+11.78
1960	9.04	-2.76
1970	11.09	-0.71
1980	10.28	-1.52
1990	14.45	+2.65
2010	11.85	+0.05
2015	10.11	-1.69

The results show that the estimated optimum is highly sensitive to the exclusion of certain benchmark years. Removing 1925 or 1950 results in extreme shifts, increasing the optimum to 27.33 °C and 23.58 °C, respectively. These years correspond to pre- and post-war periods, indicating that the model heavily relies on their inclusion to define the concave shape of the curve. In contrast, omitting 1938 dramatically lowers the estimated optimum to 2.68 °C, further highlighting the central role of this benchmark years in shaping the temperature-growth relationship. More moderate shifts are observed when omitting years like 1960, 1970, 1980, and 2015, with changes in the estimated optimum of less than ±3°C. Notably, the exclusion of 2010 results in only a marginal increase in the optimum to 11.85°C, suggesting that this year aligns closely with the full-sample estimate. Overall, these findings confirm that while the concave relationship often reappears, the estimated turning point is not stable. The strong dependence on a few historical years points to a structural fragility in the models coefficients.

Table 6.6 summarises the estimated optimal temperature for GDP per capita growth after excluding each benchmark year from the regression. The baseline optimum, based on the full sample, is 11.8 °C. The results show that omitting certain years leads to large shifts in the estimated peak, while others have little impact. The greatest changes occur when early benchmark years are excluded. Omitting 1925 and 1950 raises the estimated optimum to 27.3 °C and 23.6 °C, respectively. These years anchor the lower part of the temperature distribution, with relatively uniform GDP levels throughout Europe. Figure 4.2 (in Chapter 4) shows how GDP per capita increased along a consistent trajectory up to 1950, after which growth patterns began to diverge more strongly across regions. Removing these early anchor years shifts the weight of the regression toward post-war decades, when GDP growth was higher and more divergent across Europe. Excluding 1938 lowers the estimated optimum to 2.7 °C. This benchmark year appears to stabilise the central part of the temperature distribution before the economic disruption of the Second World War.

For post-war years such as 1960, 1970, 1980, and 2000, the effect of exclusion on the estimated optimal temperature is relatively small. The concave shape remains visible and the estimated peak changes only slightly. Excluding more recent years like 2010 and 2015 also results in minimal changes in terms of optimal temperature, suggesting that the modern economic period has less influence on the overall temperature optimum.

Figure 6.4 shows a close–up of the estimated temperature–growth curve from the baseline model, with dashed lines indicating the average temperature of each benchmark year. The benchmark years are spread across the temperature range without a clear pattern. For example, 1980 is the coldest year in the sample, but it has little effect when excluded. In contrast, 1925 and 1950 are closer to the centre but lead to large shifts in the estimated optimum. This suggests that a years influence is not determined by its temperature alone. It likely reflects other factors in the data, such as changes in economic conditions or variation in regional growth. Figure 6.4 shows that the position of a year within the temperature distribution does not seem explain its effect, underlining the models sensitivity to specific historical years.

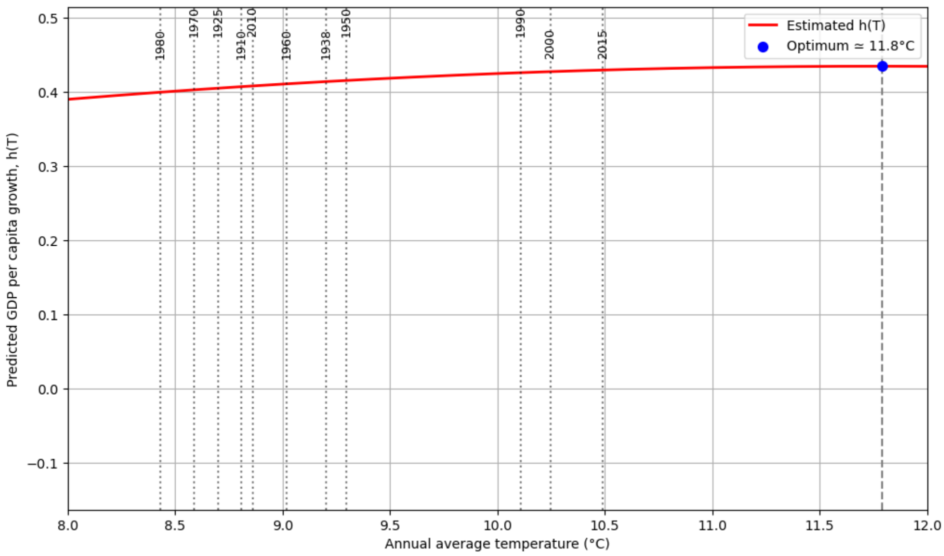

Figure 6.4: Estimated temperature-growth relationship with dashed lines showing average temperatures of benchmark years

In summary, the estimated optimal temperature is not constant across time. While the concave pattern re-emerges in most cases, the location of the turning point depends on which benchmark years are included. This underlines the importance of a long historical time span to capture a wide range of economic and climatic conditions and to avoid over-reliance on any specific period. Moreover, the strong shifts in the estimated optimum suggest that the growth–maximising temperature is not stable over time, but depends on the broader economic context of each era.

6.1.2. Insights from the jackknife analysis

The jackknife analysis confirms the robustness of the estimated non-linear relationship between temperature and GDP per capita growth. The spatial jackknife at the regional level shows that the concave shape of the temperature–growth curve is preserved across all specifications when individual NUTS-2 regions are excluded. The estimated mean coefficient for temperature is 0.0738, and for temperature squared, −0.00313, both statistically significant with very low standard deviations and p-values below 0.001. These results, presented in Table C.18, indicate that the curves shape and magnitude are not dependent on any single region. The spatial jackknife at the country level in Table C.20 confirms this robustness on a broader national scale. The coefficients remain stable and significant: 0.0739 for temperature and −0.0031 for temperature squared, with t-statistics of 8.33 and −5.26, respectively. Precipitation and its squared term also remain statistically significant across all country-level exclusions, suggesting that no single country drives the core climate–growth results either. In contrast, the temporal jackknife has a more pronounced outcome. As shown in Table C.19, the estimated coefficients are more sensitive to the exclusion of specific years. The coefficient for temperature varies more strongly, with a higher standard deviation of 0.0330, compared to 0.0013 in the spatial analysis. Its t-statistic drops to 2.31, still significant, but indicating more uncertainty. The squared temperature term, although still significant at the 5% level, has a p-value of 0.028. The squared precipitation term, however, becomes statistically insignificant, with a p-value of 0.312, suggesting that higher-order effects may be less robust over time than across space.

These temporal patterns point to the influence of economic history on the climate—growth relationship. There are significant changes in curvature and turning points when benchmark years like 1938 or 2010 are excluded. This raises the question of how much of the observable temperature–growth relationship reflects underlying climatic processes and how much is shaped or amplified by historical shocks. Given that European regions cluster on the ascending part of the temperature–growth curve, modest warming may coincide with higher growth. However, this association becomes blurred when large-scale structural disruptions such as post-war recovery, oil crises, or financial crisis happen. These shocks affect the economic baseline against which temperature effects are estimated. The concave shape remains broadly visible, but the estimated coefficients are not universal; they shift depending on which year is excluded. This shows that the full-sample estimate is not a single stable function, but the result of combining multiple, historically contingent growth paths.

To conclude, the spatial and country-level jackknife results support the robustness of the climate–growth relationship across regions. However, the temporal jackknife reveals that this relationship is sensitive to the historical time period covered. While the concave relationship proposed by BHM is supported in structure, its precise curvature and optimum vary with time. In response to Subquestion 2, this suggests that the temperature effect on economic growth is both shaped by climatic variation and moderated by broader historical context. Expanding the number of benchmark years in long-run datasets is necessary to reduce this temporal sensitivity and better isolate the climatic signal from economic noise.

6.2. Regression with temperature only

An alternative version, including only temperature as exploratory variable, of the model by BHM is carried out as an additional sensitivity analysis to determine how much of the variation in GDP growth can be accounted for by average annual temperature alone. This makes it possible to assess the stability of the non-linear temperature–growth relationship under the current panel dataset and offers a standard by which to see how the predicted relationship's shape and amplitude change when precipitation and SLR are included.

Table C.15 in reports the estimated coefficients for the temperature variables. The coefficient on linear temperature is positive and significant ($\hat{\beta}_1 = 0.0572$, $p < 0.001$), while the squared term is negative and significant at the 5% level ($\hat{\beta}_2 = -0.0021$, $p = 0.041$), confirming an concave shaped relationship between temperature and growth. This curvature is consistent with the original findings of BHM and suggests that moderate warming may support growth in colder regions, but harm hotter ones. Year FE are also included and are displayed in Table C.17 in Appendix C.

Table 6.7: Estimated coefficients for temperature effects (clustered SE)

Variable	**coef**	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
Temperature	0.0572	0.016	3.594	0.000	0.026	0.088
Temperature^2	-0.0021	0.001	-2.040	0.041	-0.004	-8.04×10^{-5}

The estimated non-linear relationship between GDP per capita growth and the annual average temperature is shown in Figure 6.5. The fitted coefficients from the regression shown in Table C.15 are used to generate the curve. The concave shape confirms that the marginal influence of temperature becomes negative after a certain point. The predicted optimum temperature for economic growth, or the turning point, is around 13.9 °C. The anticipated growth rate is maximised at this temperature. This is marginally higher than the optimal temperature estimate of BHM.

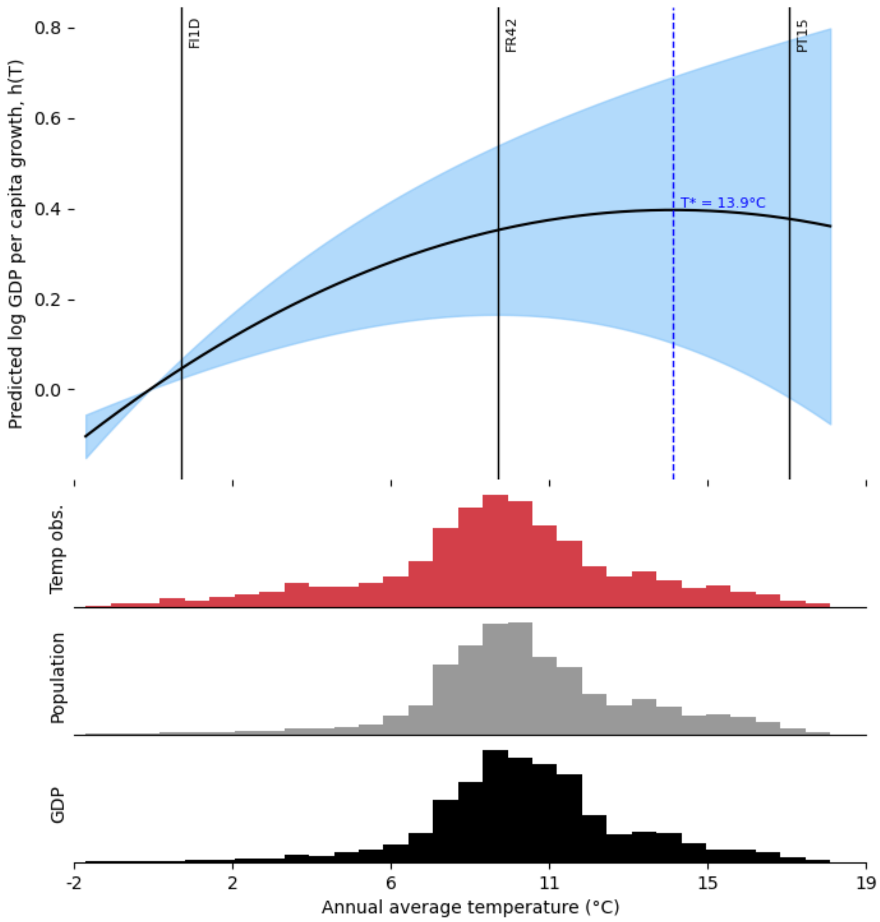

Figure 6.5: Estimated non-linear relationship between temperature and GDP per capita growth. Optimum temperature estimated at 13.9 °C.

Figure 6.5 shows the 90% confidence interval around the estimated curve and overlaying the distribution of temperature exposure across the sample. Because there are fewer data at higher temperatures, statistical uncertainty rises, as seen by the shaded region. Three histograms, number of temperature observations, population, and GDP, show how temperature values are distributed throughout the dataset beneath the curve. These verify that most of data points, population, and GDP mass are concentrated between 5 °C and 15 °C, where the fit is most dependable.
Taken together, the figures visualise the statistical findings in Tables C.15 and C.14. The concave relationship, which has a distinct peak in the middle of the observed temperature range, is not only significant but also economically meaningful. This provides a benchmark for comparing the shape and stability of the curve when introducing additional climate variables.

When only temperature is included in the regression, the predicted optimal annual average temperature for economic growth is 13.9 °C. When precipitation is included, this value drops to 11.8 °C, and when SLR is added, it slightly rises to 12 °C. The decrease from 13.9 °C to 11.8 °C suggests that the precipitation effect was partially accounted for by the original temperature coefficient. When precipitation is taken into account, the model can calculate its contribution independently, which lowers the proportion of variance that was previously attributable to temperature alone and moves the predicted optimum downward. The calculated optimal temperature is not significantly altered by the addition of SLR. This is because, unlike precipitation, which typically has a negative correlation with temperature, where hotter years tend to be drier, SLR affects economic growth in a much slower manner. Because of this, SLR adds another influence channel, but it does not as significantly confound the temperature–growth relationship as precipitation does.

7 Discussion

Using historical data from 170 European NUTS-2 areas from 1900 to 2015, this chapter discusses the empirical results from the re-analysed and extended BHM model. The first point of discussion is the robustness of the BHM model, in section 7.1. The second point of discussion is the insights from the sensitivity analyses, in section 7.2. The third point of discussion, discusses the inclusion of precipitation and SLR and their dependence on temperature, in section 7.3. The fourth point of discussion, considers the potential of an analysis which breaks GDP down per sector, in section 7.4. The fifth point of discussion, discusses the importance of adaptation in interpreting long-term estimates, in section 7.5. The final point of discussion, reflects on the broader policy implications of the findings, highlighting the importance of careful use of climate-economy estimates in policy design, in section 7.6.

7.1. Robustness of the BHM model

The study by BHM introduced a global model showing a concave temperature–growth relationship, peaking at approximately 13 °C. Applying this model to the dataset by Rosés et al. (2021) confirms that this concave relation also appears across European regions over a longer time frame. The estimated optimum in this research lies around 11.8 °C, indicating that the concave shape is preserved across both spatial and temporal expansions.

Yet this robustness appears structural rather than statistical. Small adjustments to the dataset, including the exclusion of a country or benchmark year, cause significant changes to the shape of the curve and the estimated optimum. Barker (2024) strongly criticised BHM for cherry–picking data and failing to account for sampling uncertainty. BHM estimate their model over the full period 1960–2010, with particular focus on 1980–2010 due to better data quality in those years. To test the stability of their estimated temperature–growth relationship, they divide the sample into two subperiods: 1960–1989 and 1990–2010. After doing so, they report that the results from both subperiods are nearly identical, which they interpret as evidence that the relationship is stable over time and not sensitive to the specific time frame. However, this claim is challenged by Barker (2024), who shows that shifting the cut–off year slightly (from 1990 to 1989 or 1991) can cause the results to lose statistical significance. This undermines the argument that the estimated relationship is stable and suggests that the results are more sensitive to sample choices than claimed by BHM. This thesis similarly finds that removing benchmark years like 1925 or 1938 shifts the optimum by more than 10 °C. These findings challenge the idea that a stable, generalisable optimum exists.

Also, in this research, the confidence band (in the Figure 5.2, Figure 5.4 and Figure 5.6 in Chapter 5) does not reflect the actual sensitivity of the curve: it is narrow at the colder temperatures and widens toward warmer temperatures, even though most observations lie near the middle of the temperature range. This shape is counterintuitive and suggests that the confidence band is shaped more by the programming code of the model than by the underlying sensitivity of the curve. In panel data models with FE and non–linear terms, standard errors are not easily translated into confidence bands for predicted values. FE absorb much of the variation in the climate variables, and non–linear terms cause uncertainty to vary across the curve (Blanc et al., 2017). Clustering standard errors by region helps address statistical dependence, but does not provide a straightforward way to display uncertainty across the full temperature range.

7.2. Sensitivity to historical and spatial context

The temporal structure of this study differs from BHM, who used annual growth rates. In this study, GDP per capita growth is calculated between twelve benchmark years, with uneven intervals. According to Baltagi et al. (2011), irregular spacing introduces biases and complicates identification of relationships in the data. But dividing growth by interval length could disguise historical shocks. This reflects a broader issue: how the panel data structure affects estimated climate responses. Climate evolves in a more gradual fashion over decades, while GDP reacts quickly to business cycles, wars, and financial crises (Dell et al., 2014; Kolstad et al., 2020). This difference in timescale complicates the task of identifying climate effects on economic outcomes. Standard panel models often rely on year-to-year weather variation, which may capture short–term shocks but not long-run responses (Dell et al., 2009; Mérel et al., 2021). As a result, these models risk overstating damages if they cannot account for adaptation (Mérel et al., 2021).

These structural concerns become even more visible when testing the robustness of the model through jackknife analyses. While the regional–level jackknife shows limited deviation in the estimated coefficients, the spatial country- and temporal benchmark year-level analyses tell a different story. Excluding certain countries shifts the estimated optimum by up to $+2.4\,^\circ\mathrm{C}$, and removing individual benchmark–years can shift it as far as $+15.5\,^\circ\mathrm{C}$ or as low as $-9.1\,^\circ\mathrm{C}$. This large range undermines the idea that there is a single, stable temperature optimum for economic growth.

The temporal jackknife shows that the estimated curvature of the temperature effect is highly dependent on a few benchmark years, particularly 1925, 1938, and 1950. This is despite the inclusion of year FE, which already control for mean growth shocks. These findings suggest that the model's shape and optimum temperature are not purely a result of climatic factors but are shaped by specific historical moments in European economic development. The high sensitivity to benchmark years raises questions about the reliability of the concave shape as a universal economic law. It implies that the temperature–growth relationship is contingent on data availability and historical context, rather than reflecting a fixed or universal climatic mechanism. The fact that results change drastically depending on the exclusion of one benchmark–year indicates that the temperature–growth relationship is likely contingent, not causal.

7.3. Inclusion of precipitation and SLR

Adding precipitation and SLR modifies the estimated temperature effect. The optimum falls from $13.9\,^\circ\mathrm{C}$ to $11.8\,^\circ\mathrm{C}$ when precipitation is added, and rises to $12.0\,^\circ\mathrm{C}$ when SLR is included. This supports the view that temperature absorbs the effects of other climate variables when they are omitted. Including these variables reveals their role and improves the attribution of impacts. Yet precipitation did not have a significant impact in this study, consistent with Dell et al. (2012), Tol (2021), and Khan et al. (2022), who also found weak or unstable effects. By contrast, SLR showed a small but meaningful influence. Other work supports this. Cortés Arbués et al. (2024) estimate up to 20% GDP losses in European coastal regions by 2100 under high-end SLR scenarios. However, Nováková et al. (2018) emphasise that SLR is a slow process, with long lags before impacts are measurable in economic data. This underlines the need for long time series and caution in interpretation.

Furthermore, precipitation and SLR are not independent of temperature. Rising temperatures increase evaporation and atmospheric moisture (Malhi et al., 2021), which in turn shape precipitation patterns. Temperature–driven ocean expansion and ice melt are the main causes of global SLR (Khan et al., 2022). If these climate factors cause economic damage but are themselves driven by temperature, then omitting or misattributing their effects risks biasing the estimated role of temperature. This interdependence complicates identification. FE reduce some bias, but unobserved time–varying confounders may still remain. As Dell et al. (2012) argue, without valid instruments or a model that explicitly reflects the causal chain between temperature and related variables, it is difficult to isolate temperatures independent impact. These issues highlight the need to model temperature as the central driver in a more structured way. Empirical work confirms that omitting key weather variables such as precipitation and SLR can distort temperature coefficients (Chen et al., 2019). But there are also other key climate variables, such as solar radiation, humidity, and wind speed, which influence economic outcomes and, if not included, can lead to biased results (Blanc et al., 2017).

7.4. Breakdown per sector

Sectoral vulnerability to climate change is well–documented in the more recent literature. The agricultural and coastal industrial sector are more exposed than the service sector, which is partially shielded (Tol, 2018; Kotz et al., 2022). Including sector–level data helps identify where losses occur and where adaptation is needed. Although BHM report that their concave response holds across sectors, more recent studies find sectoral variation. Zhang et al. (2018), Chen et al. (2019), and Kotz et al. (2022) highlight different climate sensitivities across economic activities. For example, labour–intensive industries may suffer productivity declines under extreme heat, while coastal infrastructure faces increased risks from SLR (Chatzivasileiadis et al., 2023).

This study could not conduct a sectoral breakdown due to data limitations. GDP data by sector were unavailable, and employment shares are not an accurate substitute. Yet Rosés et al. (2021) and Koodziejczak (2020) confirm that services dominate the European economy. Thus, results may primarily reflect the climate response of this sector. While sectoral breakdown is not required to estimate the aggregate temperature effect, it would provide insights into heterogeneity and support more tailored policy responses (Cortés Arbués et al., 2024).

7.5. Adaptation

Economic responses to temperature are shaped not only by immediate weather shocks but also by long–term behavioural and structural adjustments. Adaptation, including investments in irrigation, infrastructure, working hours, or early warning systems, can reduce vulnerability to climate impacts (Burke and Emerick, 2016; Blanc et al., 2017; Mérel et al., 2021). Ignoring these processes risks overstating damages or misinterpreting observed relationships (Kolstad et al., 2020; Tol, 2021; Kahn et al., 2021). The BHM model does not explicitly model adaptation. Instead, it assumes that adaptation is embedded in the historical temperature–growth relationship. Because the model identifies a stable concave shape across countries and over time, it implicitly includes past adaptive responses (Burke, Hsiang, et al., 2015). However, this embedded approach limits understanding of how future adaptation might change vulnerability, especially if future warming patterns or adaptive capacities differ from the past.

Mérel et al. (2021) propose a more precise approach to measuring adaptation by introducing a climate penalty term into the standard panel regression framework. This penalty term captures the squared deviation between current weather and the long–term average climate, defined as a 20–year rolling average of past weather. The idea is that economic agents make long–run decisions based past climate. When weather deviates from this expected norm, economic outcomes may suffer. Mérel et al. (2021) apply their model to the same country-level dataset used by BHM and do not find statistically significant evidence of adaptation. However, applying the same approach to regional crop yield data in the United States and France, Mérel et al. (2021) do find clear signs of adaptation. For instance, the estimated losses in US corn yields from a warming fall from 27% to 16% when adaptation is accounted for. This demonstrates that the capacity to adapt varies by context, and that national–level models may obscure such variation.

Mérel et al. (2021) show that standard panel models, such as those used by BHM, combine short– and long–term effects, which can lead to overstated estimates of long–run damages if adaptation is not properly accounted for. This ties into a broader debate in climate econometrics: whether temperature affects long–run growth rates or only the level of output in the short run. BHM assume the former. They argue that even small annual losses accumulate over time, because lower productivity reduces investment and technological progress, which slows down economic growth. In contrast, Newell et al. (2021) find little evidence for persistent growth effects. They argue that temperature shocks reduce output temporarily, but that economies recover, suggesting smaller long–term damages. The distinction matters. Growth effects imply lasting divergence, while level effects suggest short–term disruptions. BHM include lagged temperature terms by adding the average annual temperature of up to ten previous years as separate variables in their regression model. This approach tests whether temperature in earlier years continues to influence economic growth in later years. However, as more lagged years are added, the estimated effects become statistically less precise. Newell et al. (2021) interpret this increasing uncertainty as evidence against persistent growth effects. Recent studies suggest that both short–run and long–run effects exist, depending on the specific time frame, scope, and modelling approach (Chen et al., 2019; Dell et al., 2012; Kotz et al., 2024). Future research should focus on identifying when short–run shocks translate into long–term economic damages.

These findings strengthen the case for using regional rather than national data. National averages obscure within–country differences in climate exposure and adaptation (Rosés et al., 2021). Adaptation is typically implemented at subnational levels, where local conditions shape both risks and responses (Cortés Arbués et al., 2024). Regional data, such as the NUTS–2 level used in this study, are better suited to capture how adaptation varies across space. Within–country variation also improves the identification of short–run versus long–run effects (Mérel et al., 2021). This highlights a key limitation of the BHM approach: national averages can overlook critical differences in adaptation and exposure that shape climate damages.

This thesis does not apply the climate penalty term from the Mérel et al. (2021) model, because the dataset by Rosés et al. (2021) provides annual climate data but only includes economic growth observations for twelve benchmark years between 1900 and 2015. This means there are large gaps between economic observations, making it difficult to link short-term deviations in weather to economic outcomes. Adaptation to changing climate conditions may already have occurred within those long intervals, so the influence of weather shocks from 5 to 15 years earlier cannot be properly tracked. As a result, it is not possible to construct a meaningful penalty term based on rolling deviations between annual weather and long-term climate expectations. Still, understanding adaptation remains essential. As global warming accelerates, the distribution and effectiveness of adaptive capacity will strongly shape future outcomes. Yet quantifying adaptation is difficult. Lesnikowski et al. (2017) argue that without metrics for adaptation outcomes, it is hard to evaluate whether policies work. Nonetheless, research must keep addressing this gap. Without it, models risk drifting away from the mechanisms they aim to explain.

7.6. Policy implications

This discussion underscores that econometric models of climate impacts, including the extended and re–analysed BHM model, provide important insights but must be interpreted with caution. The full sample results confirm the presence of a concave temperature-growth relationship, but the estimated optimum temperature is highly sensitive to the composition of the dataset. This numerical fragility, revealed through both spatial and temporal jackknife analyses, challenges the notion of a universal economic threshold. While the shape of the curve appears mostly robust across space and time, its precise parameters are not. This distinction is essential for policymakers. Structural robustness does not imply predictive certainty. The inclusion of precipitation and SLR improves the attribution of damages and reveals the broader climate system's interdependence. However, these variables are often driven by temperature, and failing to model their causal dependencies risks underestimating systemic effects. Explicitly modelling temperature as a driver of related climate variables is essential to capture indirect pathways of impact. Policies that focus solely on direct temperature damages may overlook compound risks from other climate variables related to temperature.

The debate about whether climate change affects growth rates or only output levels carries major implications for policy design. If damages affect growth, even small effects accumulate into large future losses. If damages affect only levels, the short-run costs are real, but recovery is possible. Policymakers often rely on economic damage estimates to support emission targets or investment in adaptation. But as Pindyck (2017) warns, estimates based on empirical climateeconomy relationships may give a false sense of precision. Hsiang (2016) argue that such models can guide scenario analysis but must be paired with transparent assumptions. The results in this thesis confirm that small changes in sample structure can produce very different policy implications. Policymakers should therefore treat these models as exploratory tools rather than predictive instruments.

Also, the functional form of the model assumes a global concave relationship between temperature and growth. While intuitive and easy to estimate, this choice is restrictive. Studies such as Newell et al. (2021) and Barker (2024) question whether a quadratic curve accurately captures complex climate-economy dynamics. Different, more flexible functional forms, could reveal non-linearities, plateaus, or multiple optima that the current form conceals.

The results suggest that while the BHM model captures a robust statistical pattern, its exact parameters are highly context–sensitive. The observed concavity is preserved in the full sample, but the estimated coefficients and turning points are unstable. This makes the model structurally robust but numerically fragile. For researchers and policymakers relying on these estimates to predict economic damages from warming, this distinction is crucial. A small change in sample or specification can lead to vastly different policy implications. Therefore, the model by BHM should be interpreted as a framework for identifying patterns, not for defining precise thresholds. Its application to longer time frames and higher spatial resolution reveals both the strengths and the limits of reduced-form econometric approaches to climate impacts.

The jackknife analysis for both the 30–year climate means and the annual temperature model shows that the estimated non-linear temperature-growth relationship is robust across regions, with only minor variation when individual NUTS regions are excluded. In both specifications, the spatial standard deviations of the temperature and precipitation coefficients are very small, confirming that no single region drives the overall results. However, key differences emerge when comparing temporal sensitivity. The model based on climate means exhibits greater variability in the linear temperature term when individual years are omitted ($\sigma = 1.80 \times 10^{-1}$), while the annual temperature model displays a lower, though still notable, temporal standard deviation ($\sigma = 2.95 \times 10^{-3}$). Despite the smaller scale of coefficients in the annual model, both versions demonstrate that certain years exert stronger influence than others, especially mid-century benchmarks like 1950. Moreover, although the figures for both models use an identical x-axis range from 0°C to 35°C, the actual temperature observations in the data do not exceed approximately 18°C. As a result, the right-hand side of each graph extrapolates beyond the available data, and confidence intervals widen accordingly. This artificial extension of the domain, while useful for visual consistency, underscores the need for caution in interpreting the tails of the estimated temperature-growth curves. In sum, both specifications support a concave long-run temperature effect on economic growth, but the curvature and sensitivity to individual years vary depending on the temporal resolution of the temperature input.

8
Conclusions

This chapter concludes the summary of the findings of this thesis in Section 8.1. The implications of climate econometric models for policy design are concluded on in Section 8.2. The limitations and directions of future research are outlined in Section 8.3.

8.1. Summary of findings

This thesis examined the robustness of the non-linear relationship between temperature and economic growth proposed by BHM by extending both the temporal and spatial scope of the original model. Using a regional panel dataset covering 170 European NUTS-2 regions from 1900 to 2015, the analysis tested whether the concave shape and turning point of the temperature-growth curve persist when an additional climate variable is included and when sensitivity to different regions, countries and benchmark years is evaluated.

First, re-estimating the non-linear relationship between temperature and economic growth over the extended dataset confirms that the concave functional form identified by BHM is structurally robust. The estimated turning point lies at 11.8 °C, close to the 13.0 °C global estimate of the original model by BHM. However, the estimated coefficients are approximately six times larger, suggesting that regional–level variation within Europe may capture stronger climate sensitivities than global, national–level data. This indicates that aggregation can dampen the observed temperature effect and supports the use of regionally disaggregated data.

Second, the inclusion of SLR does not fundamentally change the concave shape but shifts the estimated turning point. The temperature–only regression estimates an optimum at 13.9 °C, which drops to 11.8 °C when precipitation is included, and slightly rises to 12.0 °C with SLR. This shift confirms that part of the temperature effect may be absorbing the influence of other, not included, variables. While precipitation remains statistically insignificant on its own, it is jointly significant when included with other climate variables. SLR has a small but statistically significant concave effect, consistent with findings that it damages growth through physical capital losses in coastal areas (Cortés Arbués et al., 2024). These results highlight the importance of accounting for multiple and interdependent climate channels.

Third, sensitivity analyses reveal strong structural differences between spatial and temporal variation. The regional–level jackknife shows that the estimated temperature-growth curve is highly robust to the exclusion of individual regions. Estimated turning points remain within a narrow band, approximately ± 0.6 °C around the baseline estimate. When entire countries are excluded, however, the results become less stable. The country–level jackknife reveals that omitting certain countries shifts the estimated optimum by up to $+2.4$ °C. Still, the concave shape persists across all country–level subsamples, supporting the structural robustness of the relationship. However, the temporal jackknife results show much greater instability. Removing specific benchmark years, particularly 1925, 1938, or 1950, produces significant shifts in the estimated optimum temperature, up to $+15.5$ °C higher or -9.1 °C lower than the estimated optimum of 11.8 °C. This reveals that the estimated relationship is not a universal function, but rather a collection of temporally distinct curves shaped by both climate and macroeconomic conditions. The full–sample estimate is therefore not a stable optimum, but a historical average of multiple context–dependent relationships.

8.2. Econometric models for policy design

This thesis confirms that the temperature-growth relationship is sensitive to regional and temporal conditions, suggesting that policy should not rely on a single global damage estimate. Instead, regionally tailored estimates are needed to assess vulnerability, prioritise adaptation, and design more accurate cost–benefit analyses of climate policy. Econometric models offer a key tool for this purpose by providing empirical estimates of climate effects that can be directly integrated into policy frameworks (Baltagi et al., 2011). They support policy design by revealing how temperature and related variables affect growth trajectories across space and time. When used to estimate marginal effects, evaluate persistence, and simulate policy scenarios, these models can improve Integrated Assessment Models, whose damage functions have historically lacked empirical foundations. Econometric models might provide evidence that helps clarify whether small changes in climate conditions may lead to large cumulative economic damages. This distinction is critical for setting mitigation targets and justifying early investment in adaptation.

8.3. Limitations and future research

This thesis has shown that the estimated relationships of the original study by BHM are robust to modifications to model specification and data. Reanalysis and extension tests have provided insights into the stability and potential generalisability of the findings. Through a re–analysis and extension of the model by BHM, applied to climatic and economic data from 170 NUTS–2 regions in Europe between 1900 and 2015 provided by the Climatic Research Unit, University of East Anglia (n.d.) and Rosés et al. (2021) datasets, this thesis has investigated whether the climate–related economic effects identified in the original work hold across different regional contexts and over an extended time frame. However, several limitations affect the interpretation of the results and provide directions for future work.

A first limitation concerns the use of benchmark year growth. While this method allows long-term economic data to be matched with historical climate conditions, it makes the model sensitive to the choice of years. These shifts indicate that the results reflect a mix of climatic and historical effects. Benchmark year models are useful for detecting long-run patterns, but they do not isolate short–run shocks. Future work should compare results across models using benchmark year growth and annual data to test how stable the estimated effects are over different time structures.

A second limitation relates to variable inclusion. The estimated temperature effect changes when precipitation and SLR are added to the model. The linear term increases from 0.0572 in the temperature–only model (Table C.15) to 0.0737 when precipitation is included (Table C.5), but slightly decreases again to 0.0727 with the addition of SLR (Table C.10). This pattern suggests that temperature partly reflects the influence of these omitted variables, and that including them helps isolate their effect on economic growth. However, precipitation and SLR are not independent of temperature. Their inclusion improves attribution, but also introduces identification problems. Precipitation is difficult to model due to high variability and weak effects in models that use annual averages. SLR moves slowly and affects specific locations. The interdependence of climate variables requires more advanced modelling strategies. Future studies should treat temperature as a driver of other variables and consider system–based approaches that reflect causal relations.

A third limitation is the lack of sector–level detail. Because GDP data by sector are not available at the NUTS–2 level over this full time span, the results mostly reflect average effects across the service–dominated European economy (Rosés et al., 2021). This hides differences between sectors with higher exposure, such as agriculture and coastal infrastructure. Studies using sectoral output, crop yields, or firm-level productivity could offer more precise and informative insight into how different activities respond to climate risks. This would improve the relevance of climate damage estimates for adaptation planning. Adaptation is not modelled directly in this thesis, which is in line with the model by BHM. Although long time series can reflect past adjustments, the results do not distinguish between damages avoided through adaptation and those that occurred despite it. Future work should try to include measures of adaptation directly, for example by using climate penalty terms or comparing regions with different capacities to adapt. This would allow researchers to test whether adaptation has occurred and where it is most effective.

This study used the dataset by Rosés et al. (2021), which provides detailed regional variation across 170 European NUTS–2 regions. However, the analysis focused on the aggregate relationship between temperature and economic growth, as captured by the concave curve. As a result, the regional heterogeneity present in the data was not fully explored. Future studies could follow the approach of Cortés Arbués et al. (2024) by zooming in on specific regions to examine where climate change has caused the strongest disruptions. This approach could help policymakers understand regional damage hotspots and plan investments to promote climate adaptation, even exploring scenarios of drastic relocation of capital and labour.

Finally, the shape of the estimated curve follows a fixed quadratic form. While this is common in the literature, it may hide more complex dynamics. Future research should test alternative functional forms to allow for multiple optima, thresholds, or plateau effects. This could reveal whether the concave shape truly holds, or whether the response varies in different parts of the temperature distribution. Testing functional forms is important when using these models to inform climate damage functions or policy thresholds.

In summary, this study shows that climate affects regional growth in a non–linear way, but that the strength and shape of this effect depend on model and data structure, included variables, and historical context. Future research should aim to separate short– and long–run effects, capture adaptation explicitly, and improve spatial and temporal resolution. These improvements are necessary to develop robust, transparent, and context–sensitive climate-economy estimates that can guide effective decision–making (Cortés Arbués et al., 2024; Baltagi et al., 2011).

References

Acevedo, Sebastian et al. (2020). "The effects of weather shocks on economic activity: what are the channels of impact?" In: *Journal of Macroeconomics* 65, p. 103207.

Andersen, Thomas Barnebeck, Carl-Johan Dalgaard, and Pablo Selaya (2016). "Climate and the emergence of global income differences". In: *The Review of Economic Studies* 83.4, pp. 1334–1363.

Aral, Mustafa M and Jiabao Guan (2016). "Global sea surface temperature and sea level rise estimation with optimal historical time lag data". In: *Water* 8.11, p. 519.

Baltagi, Badi H and Badi H Baltagi (2011). *What is Econometrics?* Springer.

Barker, David (2024). "Global Non-Linear Effect of Temperature on Economic Production: Comment on Burke, Hsiang, and Miguel." In: *Econ Journal Watch* 21.1.

Bester, C Alan, Timothy G Conley, and Christian B Hansen (2011). "Inference with dependent data using cluster covariance estimators". In: *Journal of Econometrics* 165.2, pp. 137–151.

Blanc, Elodie and Wolfram Schlenker (2017). "The use of panel models in assessments of climate impacts on agriculture". In: *Review of Environmental Economics and Policy*.

Burke, Marshall and Kyle Emerick (2016). "Adaptation to climate change: Evidence from US agriculture". In: *American Economic Journal: Economic Policy* 8.3, pp. 106–140.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel (2015). "Global non-linear effect of temperature on economic production". In: *Nature* 527.7577, pp. 235–239.

Carleton, Tamma A and Solomon M Hsiang (2016). "Social and economic impacts of climate". In: *Science* 353.6304, aad9837.

Chang, Jun-Jie, Zhifu Mi, and Yi-Ming Wei (2023). "Temperature and GDP: A review of climate econometrics analysis". In: *Structural Change and Economic Dynamics* 66, pp. 383–392.

Chatzivasileiadis, Theodoros et al. (2023). "Actualised and future changes in regional economic growth through sea level rise". In: *arXiv preprint arXiv:2401.00535*.

Chen, Xiaoguang and Lu Yang (2019). "Temperature and industrial output: Firm-level evidence from China". In: *Journal of Environmental Economics and Management* 95, pp. 257–274.

Clemens, Michael A (2017). "The meaning of failed replications: A review and proposal". In: *Journal of Economic Surveys* 31.1, pp. 326–342.

Climatic Research Unit, University of East Anglia (n.d.). *High-resolution gridded datasets*. URL: `https://crudata.uea.ac.uk/cru/data/hrg/`.

Cortés Arbués, Ignasi et al. (2024). "Distribution of economic damages due to climate-driven sea-level rise across European regions and sectors". In: *Scientific Reports* 14.1, pp. 1–15.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2008). *Climate change and economic growth: Evidence from the last half century*. Tech. rep. National Bureau of Economic Research.

— (2009). "Temperature and income: reconciling new cross-sectional and panel estimates". In: *American Economic Review* 99.2, pp. 198–204.

— (2012). "Temperature shocks and economic growth: Evidence from the last half century". In: *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

— (2014). "What do we learn from the weather? The new climate-economy literature". In: *Journal of Economic literature* 52.3, pp. 740–798.

Diffenbaugh, Noah S and Marshall Burke (2019). "Global warming has increased global economic inequality". In: *Proceedings of the National Academy of Sciences* 116.20, pp. 9808–9813.

Duchenne-Moutien, Ramona A and Hudaa Neetoo (2021). "Climate change and emerging food safety issues: a review". In: *Journal of food protection* 84.11, pp. 1884–1897.

Hinkel, Jochen et al. (2014). "Coastal flood damage and adaptation costs under 21st century sea-level rise". In: *Proceedings of the National Academy of Sciences* 111.9, pp. 3292–3297.

Holgate, Simon J et al. (2013). "New data systems and products at the permanent service for mean sea level". In: *Journal of Coastal Research* 29.3, pp. 493–504.

Hsiang, Solomon (2016). "Climate econometrics". In: *Annual Review of Resource Economics* 8.1, pp. 43–75.

Incropera, Frank P (2015). *Climate change: a wicked problem: complexity and uncertainty at the intersection of science, economics, politics, and human behavior*. Cambridge University Press.

Kahn, Matthew E et al. (2021). "Long-term macroeconomic effects of climate change: A cross-country analysis". In: *Energy Economics* 104, p. 105624.

Kalkuhl, Matthias and Leonie Wenz (2020). "The impact of climate conditions on economic production. Evidence from a global panel of regions". In: *Journal of Environmental Economics and Management* 103, p. 102360.

Khan, Majid and Abdul Rashid (2022). "(A) symmetry effects of climate changes on economic growth: a panel data analysis". In: *International Review of Economics* 69.4, pp. 571–607.

Kirezci, Ebru et al. (2020). "Projections of global-scale extreme sea levels and resulting episodic coastal flooding over the 21st Century". In: *Scientific reports* 10.1, p. 11629.

Koodziejczak, Wodzimierz (2020). "Employment and gross value added in agriculture versus other sectors of the European Union economy". In: *Sustainability* 12.14, p. 5518.

Kolstad, Charles D and Frances C Moore (2020). "Estimating the economic impacts of climate change using weather observations". In: *Review of Environmental Economics and Policy*.

Kotz, Maximilian, Anders Levermann, and Leonie Wenz (2022). "The effect of rainfall changes on economic production". In: *Nature* 601.7892, pp. 223–227.

— (2024). "The economic commitment of climate change". In: *Nature* 628.8008, pp. 551–557.

Lesnikowski, Alexandra et al. (2017). "What does the Paris Agreement mean for adaptation?" In: *Climate Policy* 17.7, pp. 825–831.

Magnan, Alexandre K et al. (2021). "Estimating the global risk of anthropogenic climate change". In: *Nature Climate Change* 11.10, pp. 879–885.

Malhi, Gurdeep Singh, Manpreet Kaur, and Prashant Kaushik (2021). "Impact of climate change on agriculture and its mitigation strategies: A review". In: *Sustainability* 13.3, p. 1318.

McIntosh, Craig T and Wolfram Schlenker (2006). "Identifying non-linearities in fixed effects models". In: *UC-San Diego Working Paper*.

Mérel, Pierre and Matthew Gammans (2021). "Climate econometrics: Can the panel approach account for long-run adaptation?" In: *American Journal of Agricultural Economics* 103.4, pp. 1207–1238.

Newell, Richard G, Brian C Prest, and Steven E Sexton (2021). "The GDP-temperature relationship: implications for climate change damages". In: *Journal of Environmental Economics and Management* 108, p. 102445.

Nicholls, Robert J et al. (2021). "A global analysis of subsidence, relative sea-level change and coastal flood exposure". In: *Nature Climate Change* 11.4, pp. 338–342.

Nováková, Monika and Richard SJ Tol (2018). "Effects of sea level rise on economy of the United States". In: *Journal of Environmental Economics and Policy* 7.1, pp. 85–115.

OECD (2019). *Responding to Rising Seas: OECD Country Approaches to Tackling Coastal Risks*. Paris: OECD Publishing. DOI: 10.1787/9789264312487-en. URL: https://doi.org/10.1787/9789264312487-en.

Pauling, Andreas et al. (2006). "Five hundred years of gridded high-resolution precipitation reconstructions over Europe and the connection to large-scale circulation". In: *Climate dynamics* 26, pp. 387–405.

Pindyck, Robert S (2017). "The use and misuse of models for climate policy". In: *Review of Environmental Economics and Policy*.

Rahimi, Reyhaneh et al. (2020). "Compound inundation impacts of coastal climate change: Sea-level rise, groundwater rise, and coastal precipitation". In: *Water* 12.10, p. 2776.

Rosés, Joan R and Nikolaus Wolf (2021). "Regional growth and inequality in the long-run: Europe, 1900–2015". In: *Oxford Review of Economic Policy* 37.1, pp. 17–48.

Somanathan, Eswaran et al. (2021). "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing". In: *Journal of Political Economy* 129.6, pp. 1797–1827.

Stern, Nicholas (2008). "The economics of climate change". In: *American Economic Review* 98.2, pp. 1–37.

Tol, Richard SJ (2018). "The economic impacts of climate change". In: *Review of environmental economics and policy*.

— (2021). "The economic impact of weather and climate". In.

Twardosz, Robert, Adam Walanus, and Izabela Guzik (2021). "Warming in Europe: recent trends in annual and seasonal temperatures". In: *Pure and Applied Geophysics* 178.10, pp. 4021–4032.

Uhlig, Harald (2012). “Economics and reality”. In: *Journal of Macroeconomics* 34.1, pp. 29–41.
United Nations (n.d.). *Sustainable Development Goals*. URL: https://sdgs.un.org/goals.
Xia, Yang et al. (2018). “Assessment of the economic impacts of heat waves: a case study of Nanjing, China”. In: *Journal of Cleaner Production* 171, pp. 811–819.
Zaveri, Esha, Jason Russ, and Richard Damania (2020). “Rainfall anomalies are a significant driver of cropland expansion”. In: *Proceedings of the National Academy of Sciences* 117.19, pp. 10225–10233.
Zhang, Peng et al. (2018). “Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants”. In: *Journal of Environmental Economics and Management* 88, pp. 1–17.

A

Appendix A - Data

A.1. Economic and climate data description and visualisation

xxxx

A.1.1. Tables before processing

The main output variable is regional GDP measured in 2011 international dollars, adjusted for purchasing power parity. This means that one dollar of output buys roughly the same amount of goods and services in every region. Additional variables include annual population (in thousands), land area (in square kilometres) and the employment in agriculture, industry and services (in percentage).

Table A.1: Descriptive Statistics (part 1)

	count	mean	std	min	5%	25%	50%	75%	95%	
year	29766.0	1960.9917	34.9434	1900.0000	1907.0000	1931.0000	1961.0000	1991.0000	2015.0000	2
temperature	29766.0	9.2904	3.2626	-3.3171	3.3975	7.6285	9.2217	11.0333	15.0651	
temp_max	29766.0	13.4258	3.5880	1.0818	6.7562	11.5111	13.3667	15.6530	19.3964	
precipitation	29766.0	641.9627	204.6717	93.7083	386.0955	503.7450	603.3263	736.4878	1055.6521	1

Table A.2: Descriptive Statistics (part 2)

	count	mean	std	min	5%	25%	50%	75%	
gdp_1990	2046.0	21892.2762	36382.7394	44.3411	789.9337	3623.9108	9760.2065	25190.6133	81316.
gdp_2011	2046.0	35574.2528	60963.7541	42.5413	1110.0423	4864.0178	14396.3184	41627.9092	137768.
population	2046.0	1831.4582	1788.3697	20.0608	273.1268	705.1220	1268.1855	2339.4982	5140.
emp_agrishare	2036.0	0.2257	0.2168	0.0003	0.0121	0.0439	0.1348	0.3830	0.
emp_indshare	2036.0	0.3165	0.1177	0.0508	0.1441	0.2276	0.3023	0.3929	0.
emp_servshare	2036.0	0.4579	0.2074	0.0499	0.1619	0.2758	0.4362	0.6431	0.
area	2052.0	19614.5258	25739.3625	419.2000	1784.0000	5291.1001	10831.8398	23863.6504	70273.

Number of unique NUTS regions: 246

A.1.2. Histograms before processing

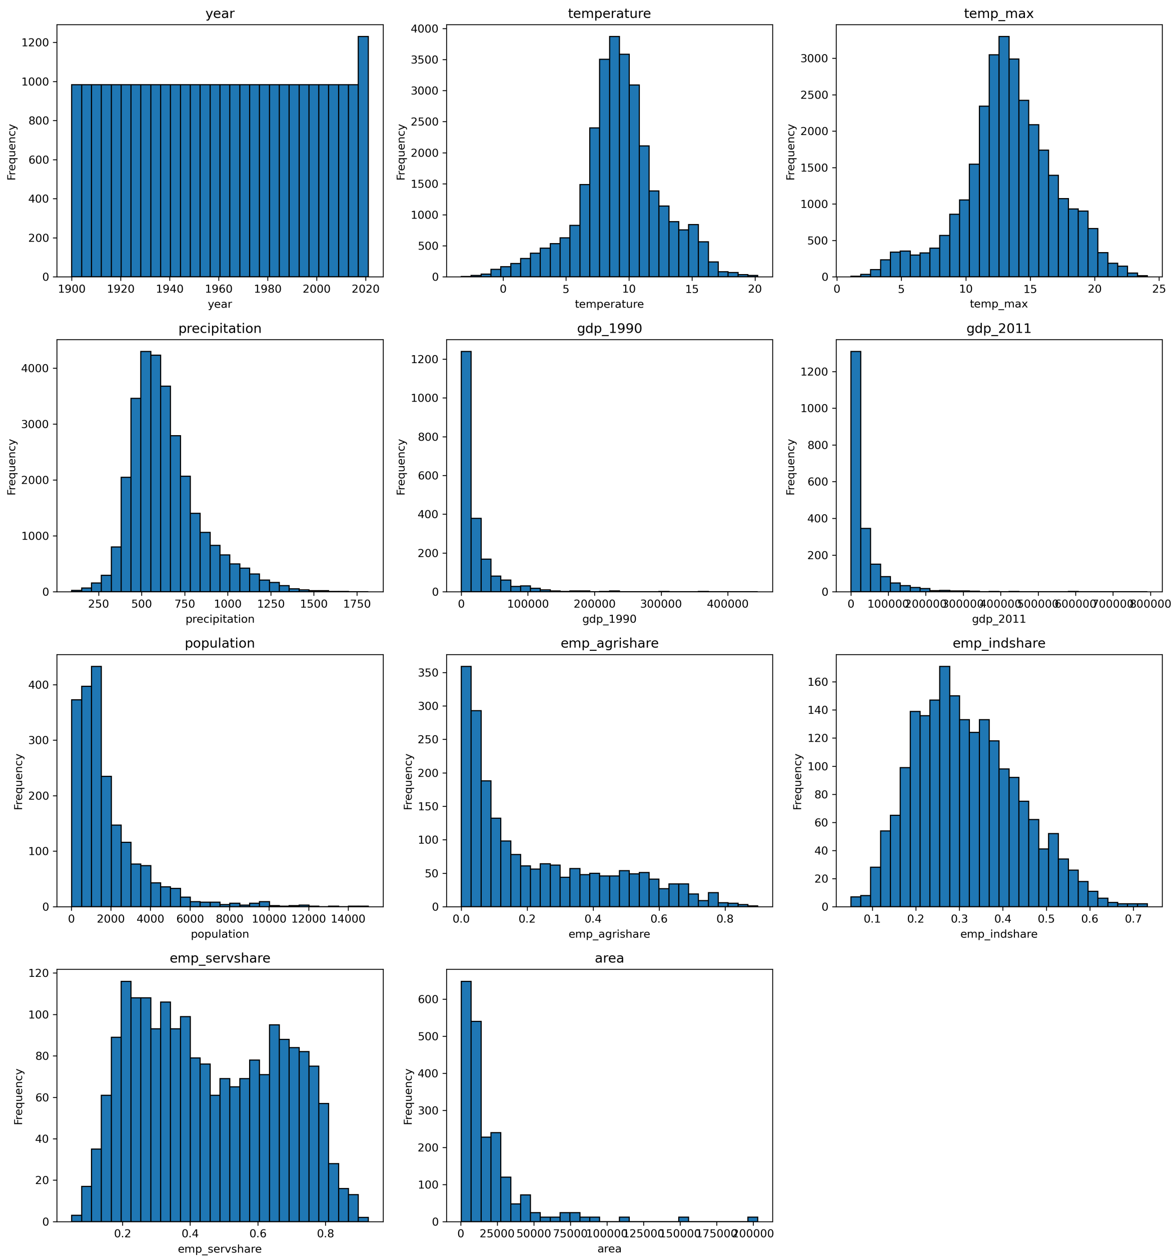

Figure A.1: Histograms roses and wolf data no prep

A.1.3. Regional GDP per capita and growth rate (for 170 regions)

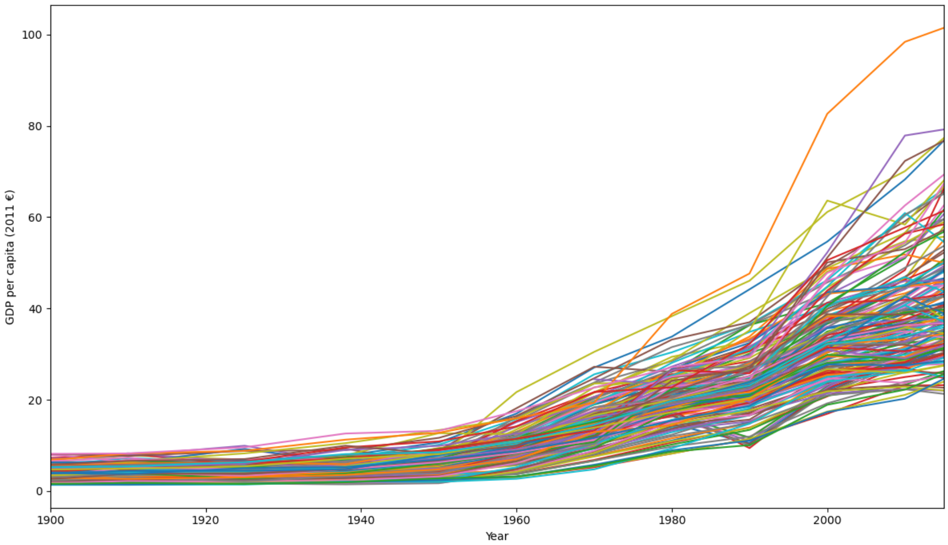

Figure A.2: Evolution of GDP per capita over time by (NUTS-2) region

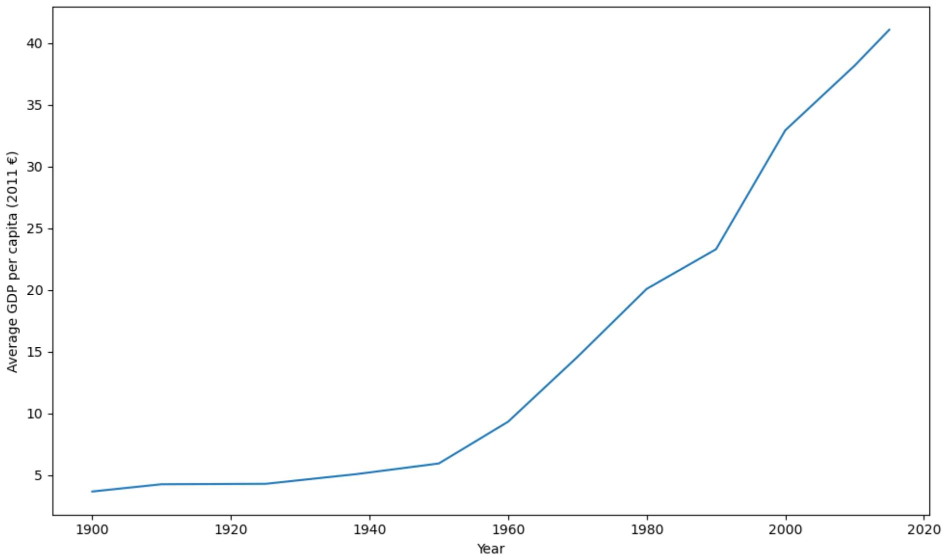

Figure A.3: Average GDP per capita over time across 170 regions

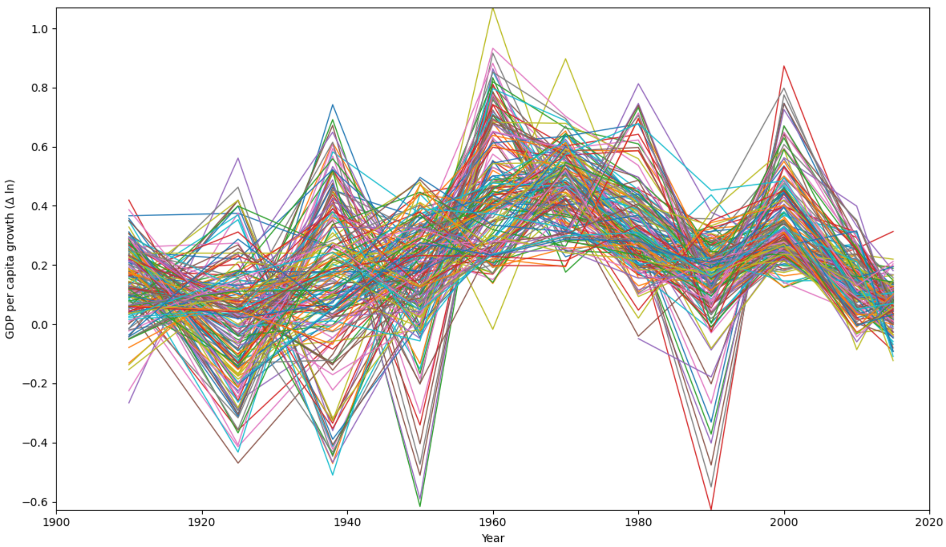

Figure A.4: Evolution of GDP per capita growth over time by (NUTS-2) region

Figure A.5: Average GDP per capita growth over time across 170 regions

A.1.4. Rolling means temperature and precipitation

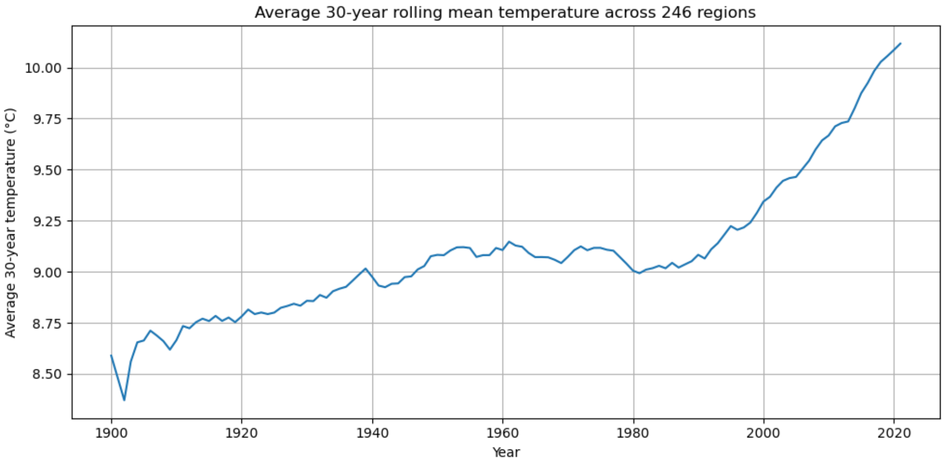

Figure A.6: Average all regions (246) rolling mean for temperature

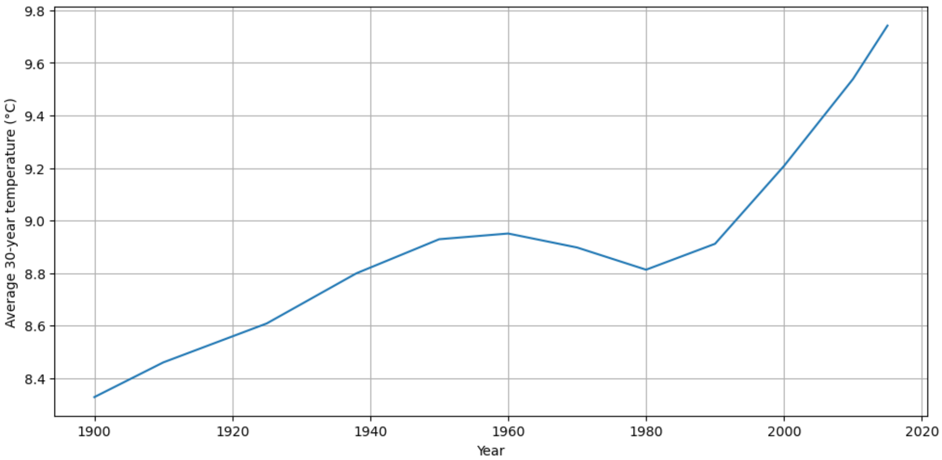

Figure A.7: Average 30-year rolling mean temperature across 170 regions

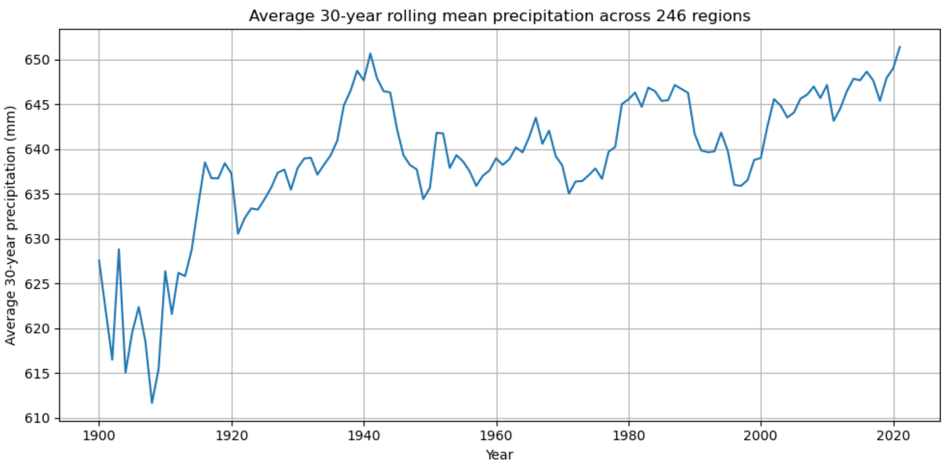

Figure A.8: Average all regions (246) rolling mean for precip

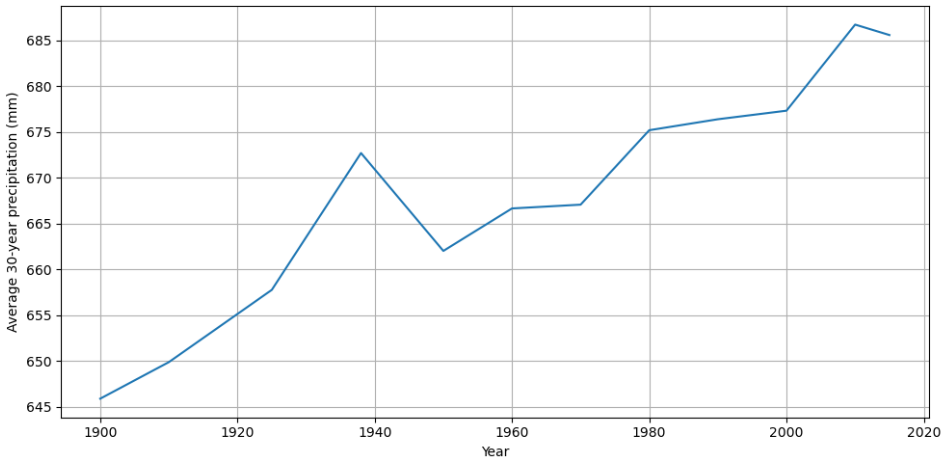

Figure A.9: Average 30-year rolling mean precipitation across 170 regions

A.1.5. SLR data preparation

SLR data before preperation

Table A.3: Descriptive Statistics of SLR Variables

	count	mean	std	min	5%	25%	50%	75%	95%	ma
year	2249.0	1966.7692	38.6771	1900.0000	1900.0000	1938.0000	1970.0000	2000.0000	2020.0	2020
AverageRSLR	757.0	7002.6667	101.2444	6624.3999	6835.3354	6960.2998	7004.7998	7049.3335	7146.0	7460
MaxRSLR	757.0	7069.5020	111.1896	6655.0000	6894.8000	7012.0000	7070.0000	7128.0000	7220.6	7612

Regions with 1 observation in AverageRSLR: 82
Regions with 1 observation in MaxRSLR: 82

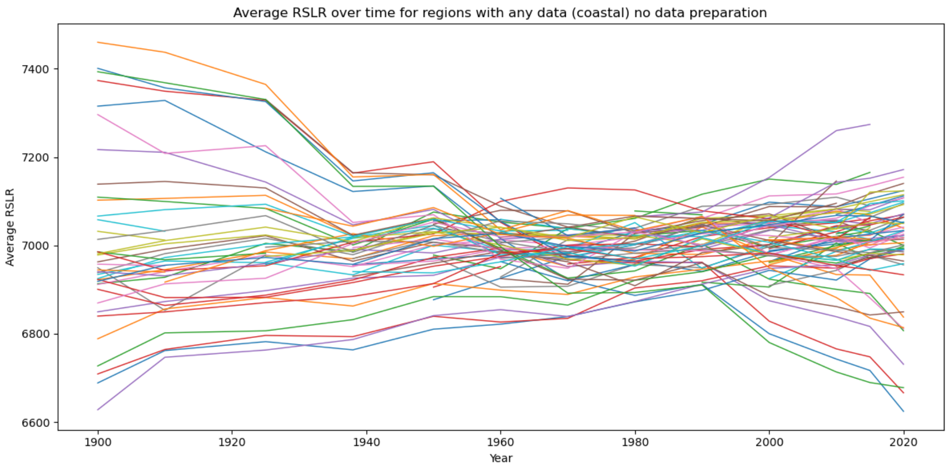

Figure A.10: All regions with SLR data no prep

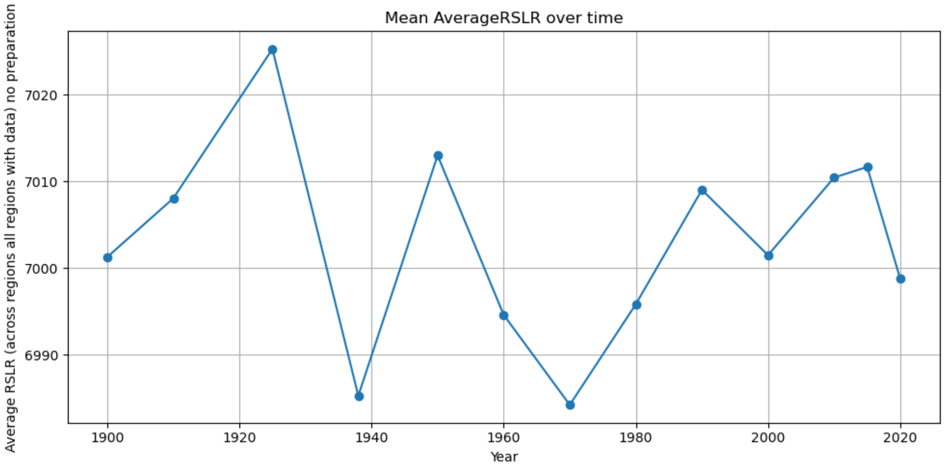

Figure A.11: Average over all regions with SLR data no prep

A.1.6. Interpolation SLR data

Table A.4: Missing values per variable

Variable	Number of missing values
nuts	0
geometry	0
year	0
AverageRSLR	150
MaxRSLR	150

spatial and temporal interpolation

Table A.5: Remaining missing values after spatial and temporal interpolation

Variable	Missing values
nuts	0
geometry	0
year	0
AverageRSLR	5
MaxRSLR	5

Table A.6: NUTS regions with remaining missing values

NUTS	Missing count
ES12	1
ES13	1
ES21	1
FR23	1
FR61	1

Table A.7: Years with missing values per NUTS region

NUTS	Year
ES13	1900
ES12	1900
FR23	1900
FR61	1900
ES21	1900

The procedure identifies regions where only one year of sea level data is missing. If this missing year is 1900 and data for 1910 is available, the value for 1900 is filled using the value from 1910. The script then reports which regions have been updated and displays the new values for 1900. As a result, no missing values remain in the dataset. The final classification still includes 82 coastal regions, defined as those with any non-zero sea level values, and 88 non-coastal regions, where all sea level values remain zero. This gives 170 regions in total.

Visualisation SLR

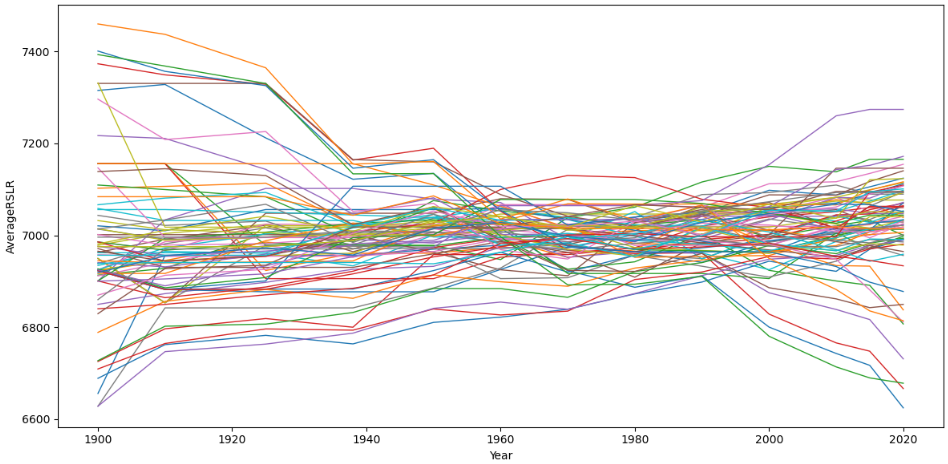

Figure A.12: Evolution 170 regions with SLR data after interpolation

Figure A.13: Average across 170 regions with SLR data after interpolation

Difference in SLR

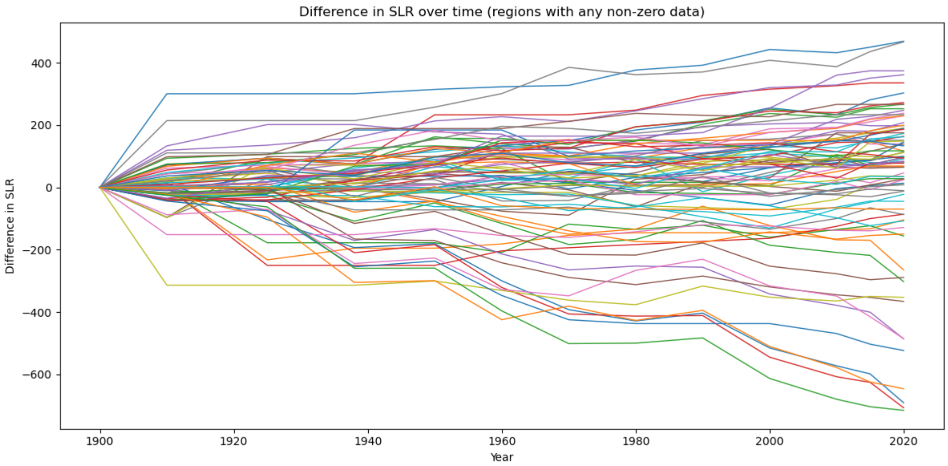

Figure A.14: Average over all regions with SLR data (coastal) after interpolation

Since 1900 for all regions is zero it does not make sense and is removed. Note, the colours of the lines are random in each plot.

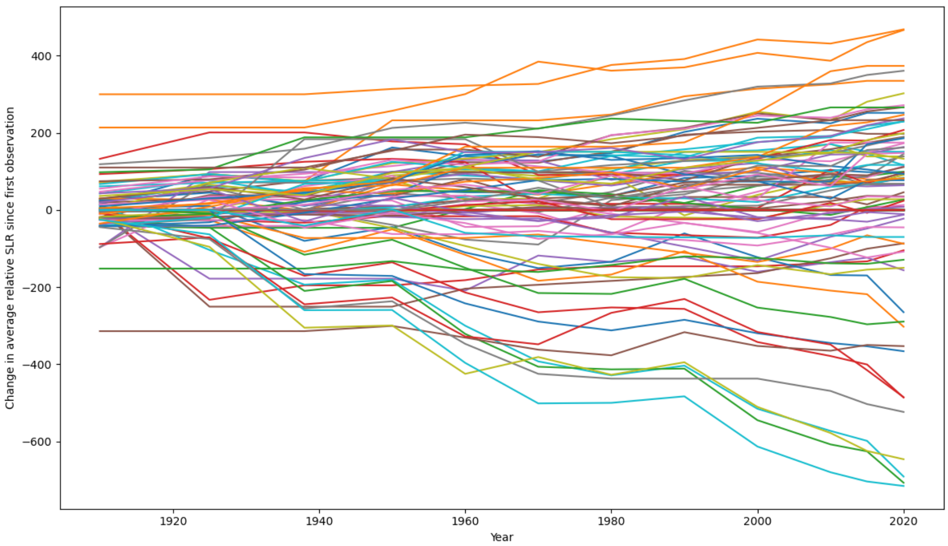

Figure A.15: Average over all regions with SLR (coastal) data after interpolation

Average SLR of coastal regions over time

Figure A.16: Average over coastal regions with SLR data

A.1.7. Scatter plots

Figure A.17 presents a series of scatterplots showing the relationship between the squared change in relative sea level rise (dif_slr_sq) and GDP per capita growth across all benchmark years from 1910 to 2015. Each subplot corresponds to a specific year and visualises the cross-sectional distribution of NUTS-2 regions. Across all years, the relationship between dif_slr_sq and economic growth appears weak and largely unstructured. Most observations are tightly clustered near the origin on the x-axis, reflecting that squared SLR changes are small for the majority of regions. A small number of regions, particularly in more recent years, show much larger values of dif_slr_sq (exceeding 5×10^5), indicating relative SLR changes likely caused by strong subsidence or uplift.

Despite these larger values, there is no clear visual indication of a consistent relationship between squared sea-level rise and GDP growth. This suggests that, on its own, the intensity of relative sea-level change does not systematically relate to regional economic performance. However, the increasing dispersion of values over time hints at growing heterogeneity in regional sea-level experiences.

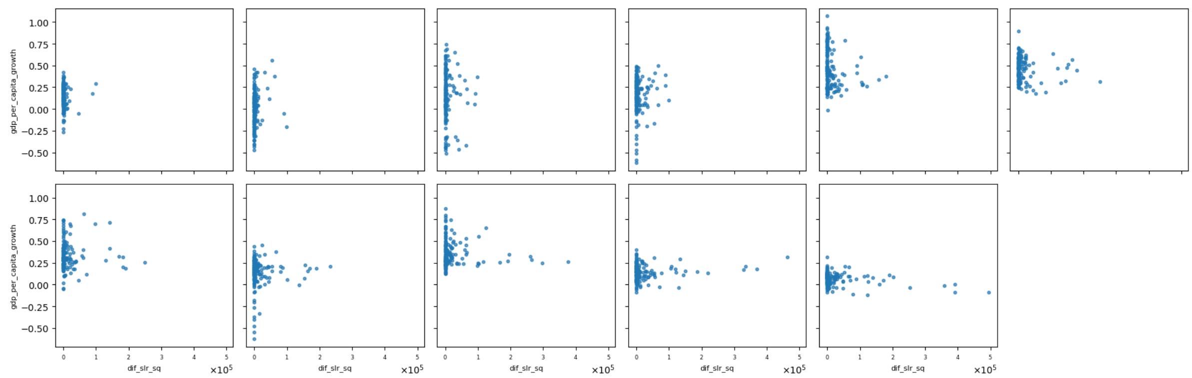

Figure A.17: Scatter plots for benchmark years GDP growth and the difference in SLR squared

Scatterplots for all exploratory variables used in the analysis (both in linear and squared form), can be found in Appendix A.

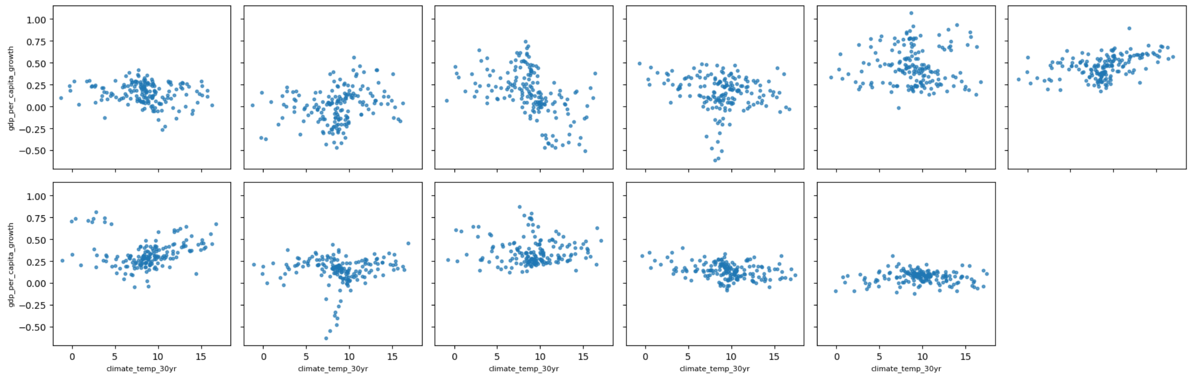

Figure A.18: scatter on benchmark years climate temp

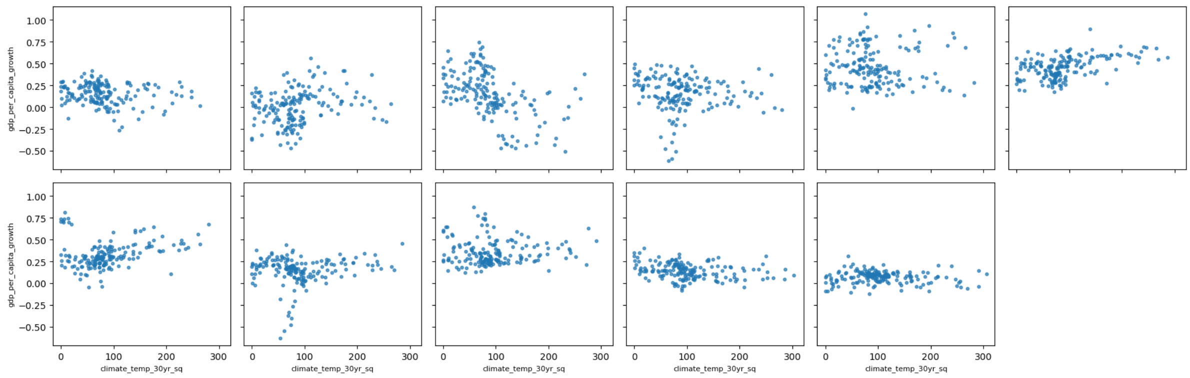

Figure A.19: scatter on benchmark years climate temp squared

A.2. pairwise correlation including squares

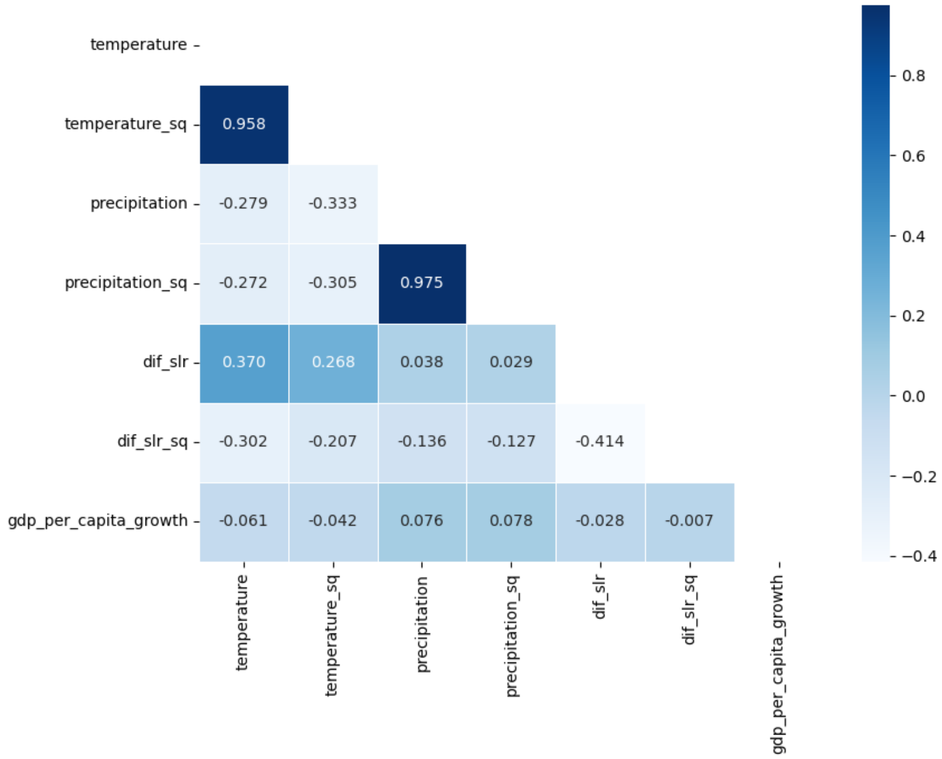

Figure A.20: Lower triangle of the pairwise correlation matrix between GDP per capita growth and climate variables

B
Appendix B - Model

B.1. Context - BHM model

This section describes the theoretical and empirical foundation of the BHM model. The model starts from a Cobb-Douglas production function where temperature influences the productivity of labour and capital (Burke, Hsiang, et al., 2015). Temperature enters the model through productivity terms that respond to climate conditions. When this framework is aggregated from individual production units to countries or regions, it leads to a smooth and concave relationship between temperature and economic output. The aggregation assumes that the shape of the temperature exposure distribution remains stable across time. As a result, changes in the annual average temperature reflect changes in productivity. This leads to the expectation that economic output increases with temperature up to an optimum level and then declines.

B.2. Research motivation - modelling approach

This research applies the BHM model to a different empirical setting to test whether the concave temperature-growth relationship is robust. The model is estimated for 170 NUTS-2 regions in Europe between 1900 and 2015. This allows for a long-run perspective and analysis at a finer spatial scale than in the original study (Rosés et al., 2021). To assess whether including more climate risks affects the main results, SLR is added as an additional variable. This makes it possible to consider potential risks from coastal flooding (Kirezci et al., 2020; Kotz et al., 2022).

Initially, a second model by Mérel et al. (2021) was considered. This model includes a penalty term for deviations from long-run climate norms to examine the role of adaptation. The idea was to implement both models and compare their findings. However, the Mérel et al. (2021) model was ultimately not included in the empirical analysis. The reason for this is that the adaptation penalty term proved difficult to interpret over the full historical period, and the primary focus of this thesis remained on long-run structural climate risks. Still, the approach by Mérel et al. (2021) is briefly discussed to reflect its conceptual relevance to the broader research design.

B.3. Methods, data and limitations

In order to do a successful robustness analysis, this section describes the methods, data sources, and limitations of two climate-econometric models: BHM and Mérel et al. (2021). As Clemens (2017) clarifies, a robustness test differs from replication in that it does not seek to reproduce identical results, but rather to assess the sensitivity of a model's conclusions under altered conditions. In line with this distinction, the BHM model is re-estimated using revised code and applied to an alternative dataset comprising 170 NUTS-2 regions in Europe over the period 1900–2015, thereby combining elements of reanalysis and extension testing. This thesis aims to determine whether the relationships identified in the original study, specifically the global concave effect of temperature on economic output (Burke, Hsiang, et al., 2015), are comparable when applied to a different regional and temporal context. The Mérel et al. (2021) model is discussed as a relevant alternative but is not implemented.

The BHM model, titled 'Global Non-Linear Effect of Temperature on Economic Production', employs a global panel dataset covering 166 countries over a time frame of 50 years, using econometric techniques to identify the relationship between temperature fluctuations and economic performance. The model accounts for long-run adaptation and non-linearity in economic responses to climate change.

The model by Mérel et al. (2021), titled 'Climate Econometrics: Can the Panel Approach Account for Long-Run Adaptation?', proposes an alternative way to estimate adaptation by including a climate penalty term. It relies on a fixed-effects panel regression with long-run climate norms and their deviation as explanatory variables. The model uses data from 12 regions in the United States and 21 in France, subdivided into 88 departments at the NUTS-3 level, over a 60-year period.

B.4. Data

This study utilises two datasets to examine the relationship between climate change and economic performance at the regional level: the Roses and Wolf V6 regional GDP dataset (Rosés et al., 2021) and high-resolution climatic data from the Climatic Research Unit (CRU) at the University of East Anglia (Climatic Research Unit, University of East Anglia, n.d.). These two datasets are merged at the NUTS-2 level.

B.4.1. Economic data - Roses Wolf V6 dataset

The Roses Wolf V6 dataset provides long-term regional GDP estimates across 16 European countries, covering 170 regions from 1900 to 2015. This dataset follows the NUTS-2 classification (as of 2010) and offers a detailed reconstruction of economic activity using historical national accounts, employment structures, and sectoral productivity estimates (Rosés et al., 2021).

Variables in the Roses Wolf V6 dataset

The variables in the dataset are as follows: country, indicating the sovereign state to which the region belongs based on current international borders; NUTS codes, the statistical classification code assigned to each region under the Nomenclature of Territorial Units for Statistics (NUTS) system used for regional economic analysis in Europe; region, specifying the name of each region; regional GDP (1990 PPP), which is regional GDP adjusted for purchasing power parity in 1990 international dollars; regional GDP (2011 PPP), which is regional GDP adjusted for PPP in 2011 international dollars; area (km^2), indicating the total land area of each region in square kilometres; population (1000s), which reports the population of each region measured in thousands; employment share in agriculture, the percentage of the regional workforce employed in the agricultural sector; employment share in industry, the percentage of the workforce employed in the industrial sector; and employment share in services, the percentage of the workforce employed in the services sector.

B.4.2. Climatic data - CRU high-resolution gridded dataset

The CRU high-resolution gridded dataset uses monthly temperature and precipitation data from the CRU TS (time-series) dataset provided by the Climatic Research Unit at the University of East Anglia. The CRU dataset offers gridded climate data at a 0.5ř 0.5ř spatial resolution, covering global climate observations from 1901 to the present.

The variables used in this research are: the variable `pre`, which represents mean precipitation, and `tmp`, which represents mean temperature.

B.5. Principal mathematical formulas from BHM

1. Macro–level output as a function of average temperature

The original study by BHM presents a macroeconomic production function (Equation 1 in the main paper and Equation 7 in the supplementary materials) that links a countrys total output to its annual average temperature. This expression is derived by summing production over all industries, locations, and times, and rewriting it in terms of average temperature. For country L in year τ, output is defined as:

$$Y_{L\tau}(\overline{T}_{L\tau}) = \sum_i \int_{-\infty}^{+\infty} f_i(T)\, g_i(T - \overline{T}_{L\tau})\, \mathrm{d}T, \tag{B.1}$$

Where $Y_{L\tau}$ is total output, $\overline{T}_{L\tau}$ is the countrys average temperature, $f_i(T)$ is the temperature–dependent productivity of a micro–level unit in industry i, and $g_i(T - \overline{T}_{L\tau})$ is the distribution of temperature exposure relative to the mean. This integral aggregates all micro–level responses to derive total output. The relationship implies that national output is a smooth concave function of annual temperature, reflecting underlying non–linear temperature sensitivities.

BHM assume that output contributions from industries and locations are additive, and that the shape of the distribution $g_i(\cdot)$ is stationary over time, shifting only with the mean. Capital and labour are assumed not to relocate in response to short–term temperature changes, making $\overline{T}_{L\tau}$ a sufficient statistic for that years exposure.

The same logic applies to regional analysis. For a subnational region R, the expression becomes:

$$Y_{R\tau}(\overline{T}_{R\tau}) = \sum_i \int_{-\infty}^{+\infty} f_i(T)\, g_i^{(R)}(T - \overline{T}_{R\tau})\, \mathrm{d}T \tag{B.2}$$

Where $\overline{T}_{R\tau}$ is the regions annual average temperature, and $g_i^{(R)}$ is the region–specific exposure distribution. The meaning is similar: regional output depends on how often different temperatures occur and how each temperature affects productivity. Regions tend to have narrower distributions than countries, so their aggregate response may reflect $f_i(T)$ more directly.

Since micro–level data are not available, only the relationship between regional output $Y_{R\tau}$ and average temperature $\overline{T}_{R\tau}$ is estimated. The function $f_i(T)$ remains theoretical. Regions are treated as economically independent, even though in reality there may be spillovers through trade or migration. This simplification allows each region to be analysed as a small open economy.

Micro-level productivity

The original study by BHM assumes a highly non–linear response of output to instantaneous temperature. A simplified form of this response, shown in their supplementary materials as equation 8, is a piecewise linear function with a turning point $\tilde{T}$ where productivity is maximised:

$$f_i(T) = \begin{cases} c_1 + b_1 T, & \text{if } T < \tilde{T}, \\ c_2 + b_2 T, & \text{if } T \geq \tilde{T}, \end{cases} \tag{B.3}$$

Continuity at $\tilde{T}$ is enforced by $c_1 + b_1\tilde{T} = c_2 + b_2\tilde{T}$, with slopes $b_1 > 0$ and $b_2 < 0$. In addition, $|b_2| > |b_1|$, meaning the decline after the threshold is steeper than the increase below it. This formulation captures the idea that output improves with temperature in colder conditions but decreases more sharply once it gets too hot. It implies a well–defined optimum beyond which productivity drops quickly.

This micro–level function $f_i(T)$ forms the basis of the macro model. When aggregated over the distribution of daily temperatures, as in equation B.1, it produces a smooth concave annual relationship between average temperature and output. This is a key component in deriving macro–level climate impacts from micro–level physiological or economic responses.

Empirical support for this non–linear shape is provided in Figure B.1 and Figure B.2. These show fitted pairwise linear regressions that approximate the shape of $f_i(T)$. Figure B.1 plots annual temperature against GDP per capita. It identifies a turning point at approximately 11.35 °C, where output peaks before declining more sharply.

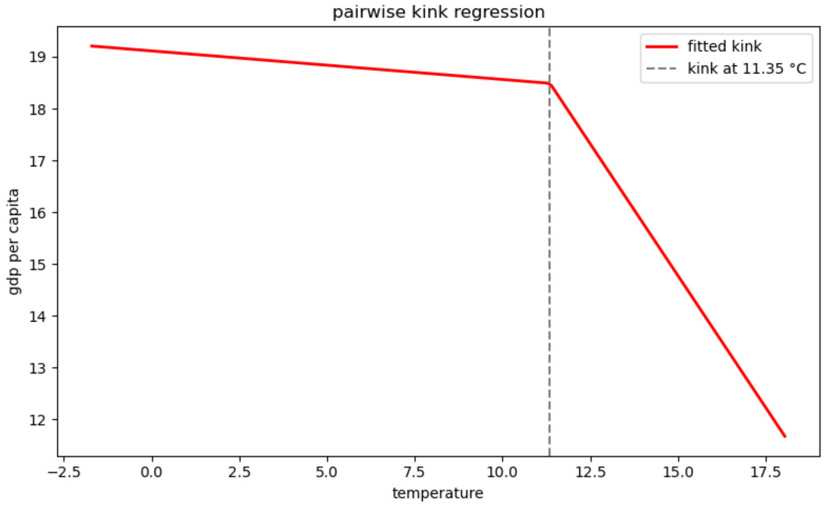

Figure B.1: Pairwise linear fit between annual temperature and GDP per capita. The kink is estimated at 11.35 °C.

Figure B.2 repeats this analysis using a rolling average climate temperature. This represents longer–term climatic exposure. The turning point in this case occurs slightly earlier, at 10.44 °C. Both figures confirm that output first rises with temperature, reaches a maximum, and then declines.

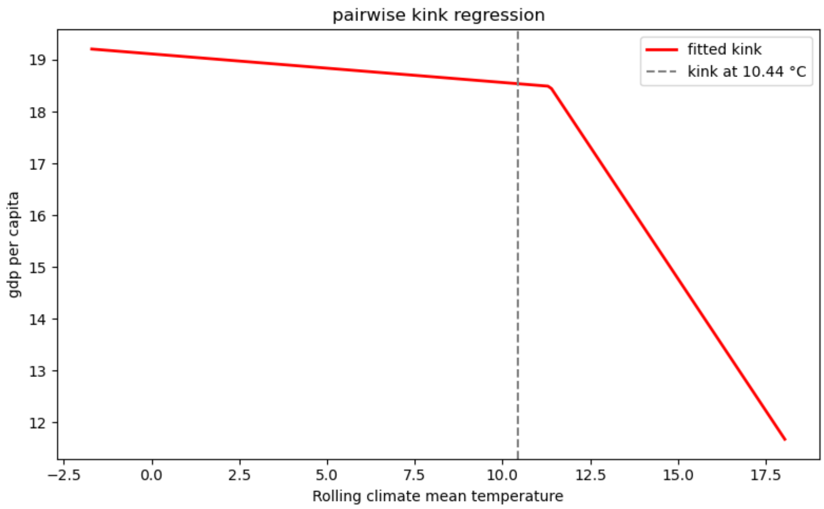

Figure B.2: Pairwise linear fit between climate mean temperature and GDP per capita. The kink is estimated at 10.44 °C.

While the piecewise linear form in equation B.3 is a simplification, it remains analytically useful and consistent with the observed data. BHM assume one dominant threshold $\tilde{T}$ and apply the same relationship across all micro units within each industry i. Capital and labour do not reallocate in response to short–term weather changes, and each unit directly experiences the effects of $f_i(T)$. This isolates the instantaneous effect of temperature on productivity.

The same micro–level shape is assumed to hold within each region. Although actual thresholds or slopes might differ across regions, the analysis follows BHM in keeping a common structure and estimating aggregate effects. In principle, each region R could have its own function $f_i^{(R)}(T)$, reflecting local adaptation. However, the model assumes a general form to assess whether average productivity outcomes vary with local conditions. The key assumption remains that each unit responds non–linearly to temperature, and that this response does not adjust over short time frames.

Non-linear aggregate response and optimal temperature

After integrating the micro–level function over the full temperature distribution, national output $Y_{L\tau}$ becomes a smooth concave function of the annual average temperature $\overline{T}_{L\tau}$. BHM derive an expression for the marginal effect of a small change in average temperature on output per unit of productive mass. Using the piecewise linear micro function $f_i(T)$, the marginal effect is written as:

$$\frac{\partial}{\partial \overline{T}_{L\tau}}\left(\frac{Y_i}{M_i}\right) = b_1 m_{i1}(\overline{T}_{L\tau}) + b_2 m_{i2}(\overline{T}_{L\tau}), \tag{B.4}$$

where M_i is the total mass of productive units in industry i, and m_{i1} and m_{i2} are the shares of time that units in industry i spend below and above the threshold temperature $\tilde{T}$, respectively. By definition, $m_{i2} = 1 - m_{i1}$.

The coefficients b_1 and b_2 reflect how productivity responds to temperature in the cooler and hotter parts of the year. If most temperatures fall below $\tilde{T}$, so m_{i1} is large, the marginal effect of warming is likely positive. If the climate is already warm, and m_{i2} dominates, the marginal effect is likely negative. In a moderate climate where m_{i1} and m_{i2} are balanced, the net effect of a small temperature change is close to zero.

This expression helps explain why the overall temperature–output relationship is non–linear. It shows that the marginal effect depends on the countrys climate profile through m_{i1} and m_{i2}. If b_1 and b_2 are fixed across all locations, then variation in marginal effects across countries or regions must come from differences in their average temperature. This supports the idea that hotter areas are more vulnerable to warming than cooler areas, not because of income differences, but because of the temperature distribution they experience.

The interpretation relies on two assumptions. First, the shape of the exposure distribution $g_i(T - \overline{T})$ does not change except for a horizontal shift. This means changes in average temperature shift the distribution left or right, without altering its variance or skewness. Second, the slopes b_1 and b_2 are constant within each segment. These assumptions simplify the model and make the integration analytically tractable.

The optimum temperature $\bar{T}^*$ is the level of $\overline{T}$ where the marginal effect becomes zero. Setting the derivative in equation B.4 equal to zero gives the condition:

$$b_1 m_1(\bar{T}^*) + b_2 m_2(\bar{T}^*) = 0. \tag{B.5}$$

This point defines the turning point of the concave function. At this temperature, the positive and negative effects of temperature exposure balance each other out.

A similar relationship holds at the regional level. Letting R represent a region, the marginal effect of regional temperature on regional output is:

$$\frac{\partial}{\partial \overline{T}_{R\tau}}\left(\frac{Y_R}{M_R}\right) = b_1 m_1^{(R)}(\overline{T}_{R\tau}) + b_2 m_2^{(R)}(\overline{T}_{R\tau}), \tag{B.6}$$

where M_R is total productive mass in region R, and $m_1^{(R)}$ and $m_2^{(R)}$ are region–specific exposure shares. This structure allows the model to be applied to smaller geographic units, using the same underlying assumptions and logic as at the national level.

temporary and lasting effects on economic growth

BHM investigate whether temperature shocks affect only the level of output or also influence its long–term growth path. To explore this, they use a simplified Solow–style growth model in their supplementary materials. When changes in the capital stock are modest from year to year, output can be approximated as linear in the stock of productive units M, with the slope depending on temperature through the function $\psi(\overline{T})$.

The total output at time t is given by:

$$Y_t = \psi(\overline{T}_t)\gamma M_t \tag{B.7}$$

where γ is a constant and $\psi(\overline{T}_t)$ captures how productivity changes with average temperature.

Capital accumulation is governed by:

$$\frac{\mathrm{d}M_t}{\mathrm{d}t} = sY_t - \delta M_t, \tag{B.8}$$

where s is the savings rate and δ is the depreciation rate. A fraction s of current output is reinvested, while a fraction δ of capital depreciates. If temperature is abnormally high in year t, then $\psi(\overline{T}_t)$ is low, leading to reduced output and therefore less investment. This reduces the capital available in the following year.

Combining these two expressions, the output in period t can be written in terms of the previous period's capital stock M_{t-1} and past temperature:

$$Y_t = \psi(\overline{T}_t)\gamma \left(M_{t-1} + s\psi(\overline{T}_{t-1})\gamma M_{t-1} - \delta M_{t-1}\right) \tag{B.9}$$

This equation shows that temperature in period $t-1$ affects output in period t through its influence on capital accumulation. A heat shock reduces productivity and savings, which lowers future capital and thus future output.

These dynamics imply that temperature affects both the level and the growth of output. In the short run, $\psi(\overline{T}_t)$ directly affects Y_t. In the longer run, that effect carries over through capital investment. If lost output is not fully compensated by higher savings, then even a temporary shock leads to a lasting gap in capital and output. The economy may not return to its previous path immediately.

In their main paper, BHM find that temperature affects GDP per capita growth, not just output levels. This raises the question: is this a persistent effect, or just a reflection of short–term fluctuations? The Solow–style model supports the idea that part of the effect is lasting. If shocks were fully temporary, then using GDP growth (the first difference of log output) would show no effect. But the findings indicate otherwise. This suggests that temperature shocks are partly persistent.

Whether the same mechanism holds for regions is a more complex question. In the dataset from Rosés et al. (2021), regions are not closed economies. Capital and labour can move between regions, and central governments can provide support. For example, a poor harvest in one region might be offset by national transfers or outside investment. This weakens the persistence of local shocks.

Still, equation B.9 can be applied to regional output Y_R and capital M_R if one assumes regions behave like small open economies. A heat shock that lowers output in region R would reduce investment and capital in the next year. However, outside capital might also flow in if profit opportunities exist. The model does not account for this, so applying it to regions involves an extra assumption: that regional capital accumulation responds mainly to local output, with limited outside influence.

Empirical panel regression model

To estimate the relationship between temperature and economic performance, BHM apply a panel regression using country–level data from 1960 to 2010. The outcome variable is the annual growth rate of GDP per capita, approximated by the first difference in the natural logarithm of GDP per capita. The main regression model is written as:

$$\Delta \ln Y_{i,t} = h(T_{i,t}) + \lambda_1 P_{i,t} + \lambda_2 P_{i,t}^2 + \mu_i + \nu_t + \theta_{i1} t + \theta_{i2} t^2 + \varepsilon_{i,t}, \tag{B.10}$$

where i refers to countries and t to years. The function $h(T_{i,t})$ captures the effect of annual temperature. In the baseline model, it is specified as a quadratic:

$$h(T_{i,t}) = \beta_1 T_{i,t} + \beta_2 T_{i,t}^2, \tag{B.11}$$

which allows for a concave response of economic growth to temperature. The term $P_{i,t}$ denotes annual precipitation, which is included as a control variable. Like temperature, it enters the model as a quadratic term, with coefficients λ_1 and λ_2.

The term μ_i represents country fixed effects. These control for all time–invariant characteristics of each country, such as average income, political institutions, or geographic conditions. The year fixed effects ν_t account for global shocks or trends that affect all countries in a given year, such as financial crises or commodity price shifts.

Country–specific linear and quadratic time trends are included as $\theta_{i1}t + \theta_{i2}t^2$. These capture long–run shifts in national growth patterns, for example due to technological change, structural reforms, or convergence processes. The term $\varepsilon_{i,t}$ is the idiosyncratic error. Standard errors are clustered by country to correct for serial correlation in residuals over time.

This empirical model provides a flexible structure to estimate non–linear effects of temperature, while accounting for country heterogeneity, time trends, and global shocks. It allows the marginal effect of temperature to vary with baseline climate conditions and controls for potential confounding factors through fixed effects and time trends.

Interpretation of the empirical specification

This model estimates how deviations in annual temperature influence economic growth within each country. The fixed effects and time trends ensure that the temperature effect $h(T_{i,t})$ is identified from short–term fluctuations around each countrys long–run growth path. In other words, the model asks whether a country grows faster or slower than usual in years that are warmer or colder than its historical average.

The function $h(T)$ is quadratic, which means that the estimated relationship between temperature and growth is concave. The coefficient β_1 reflects the initial slope of the response at moderate temperatures, while a significantly negative β_2 indicates diminishing returns and potential losses at high temperatures. The curve implied by this form identifies a specific temperature at which growth is maximised, and beyond which it declines.

This regression is the central empirical tool used by BHM to quantify the non–linear effect of temperature on output growth. It translates the theoretical framework into an econometric model that can be estimated using historical data. The results, especially the estimated $h(T)$, provide evidence for the concave shape and allow for projections under future warming scenarios. The use of country fixed effects μ_i and year fixed effects ν_t ensures that the estimated temperature effect is based on within–country variation over time. This avoids confounding due to cross–country differences, such as the tendency for cooler countries to be richer. Similarly, country–specific time trends $\theta_{i1}t + \theta_{i2}t^2$ control for gradual structural changes, such as institutional reform or long–term development. The estimation relies on the assumption that, conditional on the included controls, temperature deviations are exogenous with respect to economic growth. This means that there are no omitted variables that simultaneously influence both temperature and output. The inclusion of year fixed effects controls for global shocks, while country–specific trends address slow–moving unobserved factors. This setup provides a strong identification strategy, though it assumes that countries do not fully adapt to unusual temperatures in the short run. Short–term responses such as emergency policy, migration, or increased energy use are treated as part of the economic outcome. BHM argue that there is little evidence for full adaptation in the historical sample, so treating temperature variation as unanticipated is reasonable. They also test alternative specifications using flexible functional forms, such as splines or higher–order polynomials, and find similar patterns. This supports the use of a quadratic form for $h(T)$. Finally, it is assumed that the effect of temperature in year t appears primarily in year ts growth. BHM later test this by including lagged temperature values and find some evidence for persistent impacts. This suggests that temperature shocks may not only affect current output but also carry over into future periods.

To apply the panel regression model to a regional dataset, the same structure as equation B.10 is used, but with NUTS–2 regions as the observational units. Let r denote region and t denote year. The model now includes region fixed effects α_r to control for time–invariant differences across regions, such as baseline income levels, geography, or climate. The terms ϕ_{r1} and ϕ_{r2} are region–specific linear and quadratic time trends, capturing gradual changes in regional growth trajectories. Year fixed effects γ_t are also included to account for shocks common to all regions in a given year.

The interpretation remains similar: the model relates deviations in a regions economic growth to deviations in annual temperature. By including region fixed effects α_r, the regression isolates within–region variation and removes long–run differences between regions, such as those between northern and southern parts of Europe. The region–specific time trends account for slow–moving developments, including industrial change, migration, or convergence. The key identifying assumption is that, once fixed effects, trends, and precipitation controls are included, remaining temperature variation within each region is uncorrelated with other omitted determinants of growth. This is more plausible when the dataset includes many regions and many years. However, spatial correlation in temperature shocks can still bias standard errors. A heatwave or cold spell might affect multiple neighbouring regions at once. To account for this, inference should be based on standard errors clustered at a higher level, such as countries or climate zones.

Estimating the model at the regional level offers a useful test of robustness. If the concave relationship between temperature and growth found at the country level is a general feature of economic systems, then similar patterns should emerge across regions. If instead the results differ, this may point to factors such as national policy, labour mobility, or inter–regional trade modifying the economic impact of climate. The model structure allows for such a comparison, helping to understand whether the temperature–growth link holds across different spatial scales.

B.6. Principal mathematical formulas from Mérel et al. (2021)

1. Model with long–run adaptation

In their framework, Mérel et al. (2021) model the relationship between regional economic performance and annual climate by allowing for non–linear effects and deviations from long–run climatic norms. For region i in year t, the model is defined as:

$$y_{it} = \alpha_i + \beta_1 x_{it} + \beta_2 x_{it}^2 + \beta_3 (x_{it} - \mu_i)^2 + \varepsilon_{it} \tag{B.12}$$

The dependent variable y_{it} is the natural logarithm of GDP per capita. The variable x_{it} refers to the annual average of a weather indicator, such as temperature or precipitation. The term μ_i is the long–run average climate for region i, and it captures the conditions to which the region is assumed to have adapted. The fixed effect α_i accounts for time–invariant regional characteristics. The parameters β_1 and β_2 allow for a concave effect of weather on output, and β_3 captures whether output is reduced in years when weather deviates from the expected local climate.

The deviation term $(x_{it} - \mu_i)^2$ implies that weather shocks can reduce output when they move away from historical norms. The parameter β_3 is expected to be less than or equal to zero. If $\beta_3 = 0$, there is no penalty for deviations and agents are fully adapted. If $\beta_3 < 0$, output falls when conditions depart from long–term averages, indicating incomplete adaptation. The error term ε_{it} is assumed to be strictly exogenous, such that $E[\varepsilon_{it} \mid \alpha_i, \mu_i, x_{i1}, \ldots, x_{iT}] = 0$.

This model distinguishes between short–term and long–term responses. When $x_{it} = \mu_i$, the deviation term drops out and the expression simplifies to the long–run relationship:

$$y_{it}^{\mathrm{LR}} = \alpha_i + \beta_1 \mu_i + \beta_2 \mu_i^2$$

which defines the output level under full adaptation to the local climate. When actual weather differs from μ_i, output may decline due to the penalty from the deviation term. In this way, the model captures how agents adjust to their long–term climate but remain exposed to temporary anomalies.

To estimate this model using the dataset by Rosés et al. (2021), the dependent variable is constructed as the log of regional GDP per capita. The weather variable x_{it} can be temperature or precipitation. The climate mean μ_i is calculated as a rolling 20–year average of past values. For early years with limited data (19001919), all available prior observations are used. The squared deviation $(x_{it} - \mu_i)^2$ is included to estimate β_3. A significantly negative estimate for β_3 indicates that output is lower when weather deviates from the local climate.

Region fixed effects α_i are included to control for unobserved time–invariant heterogeneity. Identification of parameters comes from interannual variation in weather. While equation B.12 is shown for a single climate variable, it can be extended to multiple variables. For instance, both temperature and precipitation can be included with separate quadratic and deviation terms to estimate joint adaptation responses.
2. Fixed–effects panel model without adaptation

Many empirical studies on the economic effects of climate use a panel regression without explicitly modelling adaptation. This so–called naïve approach relates annual weather to economic outcomes using region fixed effects, while assuming that the effect of weather is the same for all regions, regardless of their usual climate. In terms of equation B.12, the naïve model sets $\beta_3 = 0$, which removes the adaptation term $(x_{it} - \mu_i)^2$. The simplified model is stated in Mérel et al. (2021) as equation 6:

$$y_{it} = \alpha_i + b_1 x_{it} + b_2 x_{it}^2 + e_{it}. \tag{B.13}$$

The error term e_{it} absorbs unexplained variation, including the effects of unmodelled adaptation. This specification estimates the effect of weather fluctuations on economic output, using only within–region variation over time. The fixed effects α_i control for time–invariant differences between regions, such as geography, infrastructure, or long–term income levels.

This approach is commonly used because it avoids bias from comparing regions with systematically different climates. For example, richer regions are often located in cooler climates. A simple cross–sectional regression of GDP on temperature might then mistakenly attribute differences in income to temperature rather than to other factors. By using fixed effects, the model removes these time–invariant confounders and estimates the effect of weather shocks within each region.

The inclusion of a quadratic term x_{it}^2 allows for non–linear responses to temperature. The key identifying assumption is that annual weather deviations are uncorrelated with other determinants of growth, after controlling for region and year fixed effects. This implies that x_{it} is as good as randomly assigned once these controls are included. Conceptually, the model compares a regions output in a warmer or cooler year to its average outcome, treating the region as its own control.

However, a central limitation of this model is that it does not account for long–run adaptation. The coefficients b_1 and b_2 are assumed to be constant across regions, even though some regions may be better adapted to certain climate conditions than others. As Mérel et al. (2021) note, with region fixed effects included, there is no remaining variation in long–run climate μ_i that can identify adaptation separately. This means that b_1 and b_2 reflect short–term responses to weather under the current adaptation state, but they do not capture the structural effect of being in a hotter or cooler climate.

The assumption for consistent estimation is that omitted factors like μ_i affect output only through the fixed effect α_i, and that there are no omitted time–varying variables that correlate with both x_{it} and y_{it}. If in fact the marginal effect of weather depends on μ_i, as in equation B.12, then the naïve model is mis–specified. In this case, b_1 and b_2 estimate a combination of short–run and long–run effects, rather than a clean short–run response.

While the naïve panel model does not provide unbiased estimates of the parameters in equation B.12, it still recovers some average effect of weather. The interpretation of b_1 and b_2 becomes more complicated when their values reflect both adaptation and vulnerability. Mérel et al. (2021) show that when the temperature distribution differs across regions, the non–linear shape of the fitted curve is informed partly by cross–sectional differences, even in a fixed–effects specification.

Intuitively, regions with hotter climates provide more data points at the upper end of the temperature range, while cooler regions provide more at the lower end. As a result, the estimated curvature in x_{it} reflects not just how a single region responds to different weather, but also how different regions contribute to the shape of the overall response. The estimated coefficients b_1 and b_2 are therefore biased versions of β_1 and β_2 when $\beta_3 \neq 0$, and the direction of the bias depends on the distribution of temperature variation in the sample.

Implementing equation B.13 with the dataset from Rosés et al. (2021) involves estimating a fixed effects panel regression with log GDP per capita as the dependent variable. For each region i and year t, y_{it} is defined as $\ln(\text{GDP}_{it}/\text{Population}_{it})$. Explanatory variables include annual average temperature x_{it} and its square x_{it}^2. Precipitation and its square are added as additional controls. Fixed effects α_i account for region–specific time–invariant factors, while year fixed effects control for shocks common to all regions.

The resulting coefficients b_1 and b_2 describe the average temperature–output relationship observed in the panel. For example, a negative b_1 combined with a positive b_2 implies a concave shape, where output decreases with rising temperature up to a certain point and then the decline slows. These coefficients must be interpreted with care: they reflect the effect of short–term weather fluctuations, not necessarily the long–term impact of climate change. If adaptation plays a role, as suggested by theory, then further analysis is needed to isolate its influence. The next section introduces a decomposition that clarifies how the panel estimates relate to underlying adaptation dynamics.
3. Decomposition of panel estimates into short– and long–run components

Mérel et al. (2021) derive a key result that links the fixed–effects panel model to both short–run and long–run climate responses. Under certain conditions, the coefficients estimated from equation B.13 can be expressed as a convex combination of the true short–run and long–run parameters from equation B.12. This decomposition formalises how much of the estimated effect reflects immediate reactions to weather versus structural adaptation to climate.

In the case of a quadratic temperature effect, the fixed–effects estimator converges to a weighted average of location–specific short–run responses and the general long–run relationship, as shown in equation 11 of Mérel et al. (2021):

$$\hat{\beta} = (1 - \bar{\theta})\beta^{LR} + \bar{\theta} \sum_i \lambda_i \beta_i^{SR} \tag{B.14}$$

The parameter $\bar{\theta}$ represents the share of total variation that comes from within–region annual weather variation, while $(1 - \bar{\theta})$ reflects variation across long–term climate. The decomposition implies that if local weather variability dominates, the fixed–effects panel estimate will lean towards the short–run response. Conversely, when cross–sectional climate differences are more prominent, the estimate will be closer to the long–run adaptation outcome.

This decomposition offers both theoretical insight and a practical diagnostic tool. It shows that fixed–effects estimates from many panel studies do not capture a pure short–run response but rather a mix. The framework clarifies when and how panel estimates may be interpreted as long–run effects. Specifically, when local weather is symmetrically distributed around each regions mean climate μ_i, the panel estimate $\hat{\beta}$ becomes a convex combination of the true long–run parameter β^{LR} and a weighted average of region–specific short–run parameters β_i^{SR}.

The short–run coefficients β_i^{SR} are derived from equation B.12 and depend on the regions average climate:

$$\beta_i^{SR} = (\beta_1 - 2\beta_3\mu_i,\ \beta_2 + \beta_3)$$

The average short–run effect $\beta^{SR}(\bar{\mu})$ in equation B.14 is evaluated at the mean climate $\bar{\mu} = \sum_i \lambda_i \mu_i$, with weights λ_i proportional to each regions contribution to total variance. At $x = \bar{\mu}$, the estimated panel curve is tangent to the long–run curve, which means the marginal effect matches the long–run response at this point. Elsewhere, it leans towards the short–run outcome.

When this model is applied to the European dataset by Rosés et al. (2021), it helps to evaluate whether the estimated temperature–growth relationship is driven more by long–run adaptation or short–run weather shocks. Europe includes a wide range of climates, from Mediterranean to Northern continental zones. This geographic diversity increases the share of cross–sectional variation, which may lower $\bar{\theta}$ and thus align panel estimates more closely with long–run effects. In contrast, precipitation tends to vary more across years and less across space. As a result, fixed–effects estimates for precipitation are likely to place greater weight on short–run fluctuations, yielding larger values of $\bar{\theta}$. Without explicitly modelling adaptation, such estimates may understate how regions respond to long–term shifts in precipitation patterns. The value of $\bar{\theta}$ depends on the empirical ratio between within–region weather variability and between–region climate differences. A high ratio implies that short–run responses dominate the estimate, while a low ratio suggests that the estimate reflects long–run adaptation. This decomposition framework can be used to guide interpretation and support more targeted modelling choices.

Where needed, models can be extended to include climate normals, climate–weather interactions, or explicit adaptation terms. These allow a clearer separation of long–term and short–term effects. This matters for policy relevance: understanding whether the estimated effect of a 1řC increase in annual temperature refers to a weather shock or a structural change determines the usefulness of the estimate for long–run climate impact projections.

B.7. Exploring regional heterogeneity in climate sensitivity and adaptation across Europe

This section explores the extent to which the models by BHM and Mérel et al. (2021) capture patterns of regional heterogeneity in climate sensitivity and long-run adaptation. Both models are applied to a harmonised regional panel dataset, allowing for a direct comparison of how each framework identifies differential climate responses across European NUTS–2 regions over the long term.

The model developed by BHM assumes that economic output responds to annual temperature in a concave manner, where productivity increases up to an optimal level and declines thereafter. The specification includes fixed effects to account for time-invariant differences across regions, such as geography or baseline income, and uses year-to-year variation in temperature to estimate the economic response. A quadratic term for temperature captures the non-linear relationship, consistent with the idea of an optimal climatic range.

Although the model is designed to estimate short-run effects of weather variation, the use of non-linear terms means that it partially reflects long-run differences in regional climate. As discussed in McIntosh et al. (2006) and emphasised by BHM, panel models that include higher-order temperature terms draw on both within-region variation and cross-sectional climate differences. This allows regions with different average temperatures to inform different parts of the estimated curve, such that the regression reflects not only short-run sensitivity to weather but also long-run differences in climate exposure.

Adaptation is not explicitly estimated in this model, but its effect is indirectly captured through the variation in climate across regions. If regions have adapted to their historical climate, this will be reflected in their contribution to the shape of the estimated temperature–growth curve. However, the model does not separate adaptation as a distinct mechanism and does not estimate how much of the observed effect is due to structural adjustment versus transitory weather responses.

The regression follows the specification in equation B.10, where both temperature and precipitation are entered in quadratic form. Region fixed effects absorb time-invariant characteristics such as institutional quality or geographic location. Year fixed effects control for common shocks, such as financial crises or global technological changes. Region-specific trends may be added to account for long-term development paths, such as those shaped by industrialisation or European integration.

In this model, heterogeneity in regional climate sensitivity can emerge through two main mechanisms. First, regions at different points on the estimated temperature curve will experience different marginal effects. Second, fixed effects and time trends allow each region to have its own baseline and trajectory, even though the temperature–growth relationship is estimated uniformly across the sample. As a result, while the model does not include explicit adaptation terms, it captures spatial variation in climate sensitivity and allows for an assessment of how warming may affect regions differently.

The following section explores the model by Mérel et al. (2021), which explicitly separates short-run responses from long-run climatic adaptation and offers a complementary approach to understanding regional heterogeneity in the economic impacts of climate.

First, regions differ in their long-run average temperature and therefore occupy different positions on the estimated global concave response function $h(T)$. For example, a southern European region with a relatively warm climate is more likely to fall on the right side of the temperature curve, where additional warming is associated with reduced economic performance. In contrast, a colder northern region is more likely to lie on the left side of the curve, where marginal increases in temperature may be beneficial or less detrimental. As a result, the marginal effect of temperature varies by region, even though the parameters β_1 and β_2 are estimated uniformly across all regions. The local sensitivity is given by the derivative

$$\frac{\partial h}{\partial T} = \beta_1 + 2\beta_2 T$$

which depends directly on the regional temperature level T.

Second, structural heterogeneity can be investigated by interacting temperature terms with region-specific characteristics. Following the approach described in BHM, interactions can be included between average temperature $\overline{T}_i$, average income $\overline{Y}_i$, or dummy variables representing different region groups, such as richer and poorer regions. These interaction terms allow formal testing of whether the temperature-growth curve differs across types of regions. If such terms are statistically significant, this suggests systematic variation in temperature sensitivity. If they are not, the evidence supports the use of a common functional form.

In the original work by BHM, income-based interaction terms were not statistically significant. This indicated that income levels did not meaningfully alter the shape of the temperature-growth relationship in their global dataset. A similar empirical approach is applied here to examine whether such structural heterogeneity is present across European NUTS–2 regions. This provides a basis for evaluating whether the estimated climate response function holds uniformly across regions or differs by local conditions.

An additional strategy, introduced in the supplementary materials of BHM, investigates the influence of structural characteristics by modifying the baseline specification. This alternative model, presented as equation 17 in the supplementary materials, includes interaction terms between annual temperature and two structural variables: the regions long-run average temperature and its average income. The aim is to assess whether the observed global concave response is primarily driven by climatic differences or income disparities. The temperature response function is defined as:

$$h(T_{it}) = \beta_1 T_{it} + \beta_2 (T_{it} \cdot \overline{T}_i) + \beta_3 (T_{it} \cdot \overline{Y}_i), \tag{B.15}$$

where T_{it} denotes the annual average temperature in region i and year t, $\overline{T}_i$ is the regions long-run mean temperature, and $\overline{Y}_i$ is its average income level over the sample period. These interaction terms allow the marginal effect of temperature to vary according to structural features of the region, providing a more flexible functional form.

The empirical results reported by BHM indicate that the coefficient β_2 is negative and statistically significant, while β_3 is not. This suggests that differences in long-run climate conditions explain more of the variation in economic sensitivity to temperature than income levels do. The finding implies that regions with warmer climates experience more negative marginal effects of temperature than cooler regions, regardless of income. Consequently, the global concave relationship estimated in the baseline model is largely shaped by climatic variation, and the assumption of a common functional form across regions is supported, provided that regions are already positioned differently along the temperature curve due to historical climate adaptation.

B.8. Regional heterogeneity in climate sensitivity and adaptation in the model by Mérel et al. (2021)

The model developed by Mérel et al. (2021) builds on the fixed–effects panel framework by explicitly including a term for adaptation. Their central insight is that standard panel regressions, such as the one used by BHM, may conflate short–run and long–run responses to temperature depending on the structure of the data. In particular, if economies adjust gradually to changes in their average climate, the long–run response to sustained warming may differ from the immediate effect of a weather anomaly.

The climate penalty term introduced by Mérel et al. (2021) captures the extent to which annual weather deviates from a regions long–term climatic norm. This term allows the model to distinguish between transitory weather shocks and structural changes in the baseline climate. When the average temperature $\overline{T}_i$ of a region rises gradually, the deviation between T_{it} and $\overline{T}_i$ shrinks over time, implying that the magnitude of the climate penalty should decline. The penalty term $(T_{it} - \overline{T}_i)^2$ therefore acts as an indicator of adaptation. A similar term can be constructed for precipitation, $(P_{it} - \overline{P}_i)^2$, to reflect deviations from long–run rainfall patterns.

The key innovation of this approach lies in its ability to separate the economic impact of expected climate conditions from the effect of unexpected weather variation. Over time, as agents adapt to a shifting climate, the difference between weather and climate is expected to decrease. If the estimated coefficient β_3 on the deviation term is negative and significant, it implies that output declines when weather deviates from the expected climate, suggesting that adaptation remains incomplete.

The model by Mérel et al. (2021) is designed specifically to measure regional heterogeneity in both sensitivity and adaptation. It captures two dimensions of regional climate response. First, the long–run marginal effect of temperature, which applies when $T_{it} = \overline{T}_i$, is given by

$$\frac{\partial y_{it}^{\mathrm{LR}}}{\partial T_{it}} = \beta_1 + 2\beta_2 \overline{T}_i,$$

which reflects the slope of the fully adapted climateoutput relationship at the regional norm. Second, when $T_{it} \neq \overline{T}_i$, the short–run effect includes an additional term:

$$\frac{\partial y_{it}}{\partial T_{it}} = \beta_1 + 2\beta_2 T_{it} + 2\beta_3 (T_{it} - \overline{T}_i),$$

which accounts for the immediate response to weather anomalies. If a year is unusually warm, the marginal loss is greater than what the long–run curve would predict. Conversely, cooler-than-usual years may attenuate the negative effect.

This specification provides a richer framework to analyse heterogeneity in climate impacts. Regions differ not only in their baseline response to temperature, determined by $\overline{T}_i$, but also in how strongly they react to deviations from this norm. A significantly negative β_3 indicates that output is more sensitive in years when the weather departs from typical conditions. If long–run adaptation has occurred, then much of the short–term fluctuation should be absorbed by this term, resulting in a smoother long–run response. The sign and magnitude of β_3 across regions offer insight into adaptation patterns. A smaller absolute value of β_3 may suggest successful adaptation, particularly in historically warmer regions. In contrast, larger values may indicate vulnerability to weather anomalies, especially in regions with historically cooler climates. The model by Mérel et al. (2021) thus offers a more nuanced lens through which to examine spatial variation in climate sensitivity and the role of adaptation over time.

C

Appendix C - Results

This appendix presents additional tables and diagnostic statistics supporting the main regression results of this thesis. It begins with an overview of the core findings from the global model developed by BHM, which identified a concave relationship between temperature and economic growth using national–level data from 1960 to 2010. The next sections re-analyse and extend this model in three steps. Subquestion 1 examines whether the concave pattern remains stable when applied to a long-term panel of European regions from 1900 to 2015. Subquestion 2 adds SLR to the model to assess whether the inclusion of additional climate risks alters the estimated relationship. Subquestion 3 evaluates the robustness of the temperature-growth curve by estimating a simplified model with only temperature and applying spatial, temporal, and country–level jackknife procedures. Together, these extensions test the consistency of the BHM findings across space, time, and model complexity.

C.1. Main table from the work by BHM

Table C.1: Regression estimates: main specification and robustness (1–5)

	(1) Base	(2) >20yrs	(3) No Oil	(4) No US/China	(5) ContYr FE
Temp.	0.0127^{***}	0.0135^{***}	0.0128^{***}	0.0128^{***}	0.0142^{***}
	(0.0038)	(0.0038)	(0.0036)	(0.0038)	(0.0037)
Temp. sq.	-0.0005^{***}	-0.0005^{***}	-0.0005^{***}	-0.0005^{***}	-0.0005^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Precip.	0.0145	0.0148	0.0130	0.0148	0.0124
	(0.0100)	(0.0100)	(0.0101)	(0.0101)	(0.0106)
Precip. sq.	-0.0047^{*}	-0.0049^{*}	-0.0040	-0.0048^{*}	-0.0041
	(0.0026)	(0.0026)	(0.0025)	(0.0026)	(0.0027)
Constant	1.4575^{**}	0.0740	1.4522^{**}	1.4707^{**}	-0.0362
	(0.6444)	(0.0633)	(0.6228)	(0.6507)	(0.0411)
Observations	6584	6477	6090	6484	6584
R squared	0.286	0.278	0.275	0.284	0.367
Optimum	13.06	13.39	12.64	13.09	14.92

Unless otherwise indicated, all models include country fixed effects, year fixed effects, and quadratic country time trends, with errors clustered at the country level. Temperature is measured in °C and precipitation in metres. Columns: (1) main specification, (2) as in column 1 but excluding countries with fewer than 20 years of growth data, (3) as in column 1 but dropping large oil exporting countries, (4) as in column 1 but dropping United States and China, (5) as in column 1 but adding continent-by-year fixed effects, (6) as in column 1 but adding continent-by-year fixed effects and dropping country time trends, (7) as in column 1 but dropping year fixed effects, (8) as in column 1 but only linear time trend, (9–10) as in column 1 but adding 1 or 3 lags of per capita growth (that is, lagging the dependent variable), (11) as in column 1 but using growth data from Penn World Tables. Asterisks indicate statistical significance at the 1% (***), 5% (**), and 10% (*) levels.

Table C.2: Regression estimates: robustness specifications (6–11)

	(6) ContYr + noTrend	(7) No YrFE	(8) LinearTime	(9) LDV 1lag	(10) LDV 3lags	(11) PWT
Temp.	0.0133***	0.0103***	0.0128***	0.0087**	0.0062**	0.0072*
	(0.0034)	(0.0039)	(0.0043)	(0.0039)	(0.0038)	(0.0039)
Temp. sq.	−0.0004***	−0.0004***	−0.0005***	−0.0004***	−0.0003***	−0.0004***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Precip.	0.0084	0.0159	0.0137	0.0165*	0.0201**	0.0195*
	(0.0098)	(0.0107)	(0.0100)	(0.0095)	(0.0098)	(0.0109)
Precip. sq.	−0.0021	−0.0045*	−0.0035	−0.0047*	−0.0052**	−0.0038
	(0.0023)	(0.0026)	(0.0023)	(0.0024)	(0.0024)	(0.0028)
Constant	−0.0819**	−0.8024**	−0.7693***	−11.3227***	−28.3451***	0.0643
	(0.0370)	(0.3366)	(0.0517)	(0.7957)	(2.0763)	(0.0467)
Observations	6584	6584	6584	6418	6086	6627
R squared	0.267	0.240	0.219	0.286	0.289	0.220
Optimum	17.40	12.74	13.40	11.92	9.98	9.88

Note: Same model form as Table C.1. Asterisks: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.1.1. of regression coefficients: BHM global model vs. European panel replication (this thesis)

Table C.3: Comparison of regression coefficients: BHM global model vs. European panel replication

Variable	**BHM (2015) Global (Col. 1)**	**Your European Panel Replication**	**Comment**
Temperature	0.0127*** (0.0038)	0.0737*** (0.018)	Stronger positive effect in Europe
Temperature^2	−0.0005*** (0.0001)	−0.0031** (0.001)	More curvature ⇒ lower optimum temperature
Precipitation	0.0145 (0.0100)	−0.0003 (0.0002)	Sign flips, small and insignificant
Precipitation^2	−0.0047* (0.0026)	+7.24e–08 (ns)	Sign flips, both effects negligible in Europe
Intercept	1.4575** (0.6444)	−0.0038 (0.003)	Intercept not interpretable due to fixed effects
Fixed effects	Country, year	Region (NUTS), year	Consistent, adapted to spatial resolution
Clustered errors	Country level	Region level	Same principle, adjusted for unit of analysis
Sample	Global, 1960–2010	Europe only, extended panel	Different climate zone and longer time horizon

Note: Standard errors in parentheses. Asterisks indicate statistical significance at the 1% (***), 5% (**), and 10% (*) levels. ns = not significant.

C.2. Subquestion 1: Annual temperature and precipitation

This section presents the estimation results of the BHM model extended with annual temperature and precipitation variables over the full period from 1900 to 2015. The model includes region-specific fixed effects (μ_i), year fixed effects (ν_t), and region-specific time trends of the form $\theta_{i1}t + \theta_{i2}t^2$.

Table **??** summarises the overall regression fit. The model explains 51.7% of the variation in log GDP per capita growth ($R^2 = 0.517$), while the adjusted $R^2 = 0.329$ accounts for the large number of fixed effects and interaction terms. The F-statistic of 5.624×10^{12} ($p < 0.001$) indicates strong joint significance of the included variables.

Table C.4: OLS regression summary (clustered standard errors)

Statistic	**Value**
Dependent variable	`ln_growth`
Model	OLS
Method	Least Squares
Observations	1,875
Degrees of freedom (residual)	1,350
Degrees of freedom (model)	524
R^2	0.517
Adjusted R^2	0.329
F-statistic	5.624×10^{12}
Prob (F)	0.000
Log-likelihood	834.02
AIC	-618.0
BIC	2,289.0
Covariance type	Cluster

Table C.5 provides the estimated coefficients. Temperature is significantly associated with economic growth in a non-linear, concave manner, as shown by the positive linear term and negative squared term. Precipitation coefficients are less precisely estimated.

Table C.5: Estimated coefficients from OLS regression

Variable	**Coef.**	**Std. Err.**	**t**	**P>|t|**	**[0.025**	**0.975]**
Temperature	0.0737	0.018	4.066	0.000	0.038	0.109
Temperature^2	-0.0031	0.001	-2.805	0.005	-0.005	-0.001
Precipitation	-0.0003	0.000	-1.483	0.138	-0.001	0.000092
Precipitation^2	7.24×10^{-8}	9.08×10^{-8}	0.797	0.425	-1.06×10^{-7}	2.50×10^{-7}

A joint significance test for the temperature and precipitation terms is reported in Table C.6, confirming their collective relevance with $p = 0.0018$.

Table C.6: Joint significance test for temperature and precipitation terms

Test Statistic	**Value**	**Description**
F-statistic	4.49	Test statistic for joint significance
p-value	0.0018	Probability of observing F under H_0
Numerator df	4	Number of restrictions tested
Denominator df	169	Residual degrees of freedom

Year fixed effects are presented in Table C.17. The year 1960 exhibits the largest positive deviation from the baseline, while 2015 shows the most negative deviation.

Table C.7: Estimated year fixed effects from OLS regression

Year	Coef.	Std. Err.	t	P>|t|	[0.025	0.975]
1925	-0.1648	0.026	-6.391	0.000	-0.215	-0.114
1938	-0.0451	0.028	-1.584	0.113	-0.101	0.011
1950	-0.0317	0.022	-1.430	0.153	-0.075	0.012
1960	0.2823	0.019	14.809	0.000	0.245	0.320
1970	0.2621	0.020	12.974	0.000	0.223	0.302
1980	0.1342	0.027	4.945	0.000	0.081	0.187
1990	-0.1067	0.018	-5.877	0.000	-0.142	-0.071
2000	0.1045	0.014	7.623	0.000	0.078	0.131
2010	-0.0923	0.016	-5.944	0.000	-0.123	-0.062
2015	-0.2213	0.017	-13.353	0.000	-0.254	-0.189

The model's residuals are further evaluated in Table C.8. The Durbin–Watson statistic of 2.947 indicates no problematic autocorrelation. While the skewness is near zero, the Jarque–Bera test ($p < 0.001$) and a kurtosis of 5.513 confirm heavy-tailed residuals. The condition number of 4.47×10^{17} reflects multicollinearity, which is expected given the large number of fixed effects. These issues are mitigated by the use of clustered standard errors.

Table C.8: Diagnostic statistics for temperature and precipitation model

Statistic	Value
Durbin–Watson	2.947
Jarque–Bera p-value	< 0.001
Skewness	0.082
Kurtosis	5.513
Condition number	4.47×10^{17}

C.3. Subquestion 2: Adding sea level rise to temperature and precipitation

This section builds on the model in Section 5.1 by including SLR alongside temperature and precipitation. The aim is to examine whether compound climate effects meaningfully alter the climate–growth relationship and provide additional explanatory value beyond temperature and precipitation alone.

Table C.9 reports the regression summary. The R^2 remains high at 0.518, with an adjusted R^2 of 0.330values very similar to the specification without SLR ($R^2 = 0.517$, adjusted $R^2 = 0.329$). The joint F-statistic of 4.415×10^{12} ($p < 0.001$) confirms overall model significance.

Table C.9: OLS regression summary with temperature, precipitation, and sea level rise

Statistic	**Value**
Dependent variable	`ln_growth`
Model	OLS
Method	Least Squares
Observations	1,859
Degrees of freedom (model)	520
Degrees of freedom (residual)	1,338
R^2	0.518
Adjusted R^2	0.330
F-statistic	4.415×10^{12}
Prob (F-statistic)	0.000
Log-likelihood	826.67
Akaike Information Criterion (AIC)	-611.3
Bayesian Information Criterion (BIC)	2,269.0
Covariance type	Clustered (by region)

The estimated coefficients in Table C.10 again confirm a concave temperature-growth relationship. Precipitation and SLR effects are weaker, though SLR and its squared term are jointly significant at the 10% level.

Table C.10: Estimated coefficients for temperature, precipitation, and sea level rise (clustered SE)

Variable	**coef**	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
Temperature	0.0727	0.021	3.547	0.000	0.033	0.113
Temperature^2	-0.0030	0.001	-2.455	0.014	-0.005	-0.001
Precipitation	-0.0003	0.000	-1.532	0.126	-0.001	8.24×10^{-5}
Precipitation^2	8.16×10^{-8}	9.03×10^{-8}	0.904	0.366	-9.53×10^{-8}	2.58×10^{-7}
Sea level rise	-0.0007	0.000	-2.326	0.020	-0.001	-0.000
Sea level rise^2	-1.199×10^{-6}	6.96×10^{-7}	-1.723	0.085	-2.56×10^{-6}	1.65×10^{-7}

Table C.11 shows the results of a joint F-test for the six climate variables. The test strongly rejects the null hypothesis that temperature, precipitation, and SLR (and their squared terms) have no joint effect on growth ($p = 8.07 \times 10^{-5}$).

Table C.11: Joint significance test for temperature, precipitation, and SLR terms

Test component	**Value**
F-statistic	5.08
p-value	8.07×10^{-5}
Degrees of freedom	6 (numerator), 168 (denominator)

Note: The null hypothesis is that all six climate terms (temperature, temperature^2, precipitation, precipitation^2, SLR, and SLR^2) jointly have no effect on GDP per capita growth. The low p-value suggests the null hypothesis can be rejected.

Table C.12 displays regression diagnostics. Residuals are non-normal (Jarque–Bera $p < 0.001$), with moderate skewness and high kurtosis. The Durbin–Watson statistic (2.954) indicates no first-order autocorrelation. The condition number remains high due to multicollinearity among fixed effects and polynomial terms.

Table C.12: Regression diagnostic statistics (temperature, precipitation, and sea level rise model)

Statistic	Value	Statistic	Value
Omnibus	112.491	DurbinWatson	2.954
Prob(Omnibus)	0.000	JarqueBera (JB)	457.118
Skewness	−0.091	Prob(JB)	5.47×10^{-100}
Kurtosis	5.422	Condition Number	9.08×10^{17}

The year fixed effects shown in Table C.13 are stable across specifications and align with patterns seen in earlier models. For example, 1960 and 1970 again show strong positive deviations consistent with post-war economic expansion. The stability of these values supports the robustness of the model in isolating climate effects from broader global shocks.

Table C.13: Estimated year fixed effects (temperature, precipitation, and sea level rise model)

Year	coef	std err	z	P>|z|	[0.025	0.975]
1925	−0.1616	0.026	−6.239	0.000	−0.212	−0.111
1938	−0.0437	0.029	−1.518	0.129	−0.100	0.013
1950	−0.0241	0.023	−1.060	0.289	−0.069	0.020
1960	0.2887	0.020	14.640	0.000	0.250	0.327
1970	0.2654	0.020	13.036	0.000	0.225	0.305
1980	0.1363	0.028	4.933	0.000	0.082	0.190
1990	−0.1060	0.018	−5.814	0.000	−0.142	−0.070
2000	0.1069	0.014	7.913	0.000	0.080	0.133
2010	−0.0870	0.016	−5.506	0.000	−0.118	−0.056
2015	−0.2145	0.017	−12.744	0.000	−0.248	−0.182

C.4. Subquestion 3: Robustness and temperature-only specification

This section investigates the robustness of the concave temperature-growth relationship estimated by the BHM model by first estimating a reduced model that includes only annual temperature and its square. It then presents results from spatial, temporal, and country-level jackknife exercises, which test the sensitivity of the findings to the exclusion of individual regions or years.

Table C.14 presents the regression output of the reduced specification. The model explains approximately 51.3% of the variation in log GDP per capita growth ($R^2 = 0.513$), with an adjusted R^2 of 0.325. The F-statistic (2.619×10^{12}, $p < 0.001$) confirms overall significance. The residual standard errors are clustered at the region level to account for autocorrelation and heteroskedasticity.

Table C.14: OLS regression summary with clustered standard errors

Statistic	**Value**
Dependent variable	`ln_growth`
Model	OLS
Method	Least Squares
Observations	1,870
Degrees of freedom (model)	519
Degrees of freedom (residual)	1,350
R^2	0.513
Adjusted R^2	0.325
F-statistic	2.619×10^{12}
Prob (F-statistic)	0.000
Log-likelihood	824.27
AIC	-608.5
BIC	2,269.0
Covariance type	Clustered (by region)

Table C.15 reports the coefficient estimates. The temperature effect is positive and significant, while the squared term is negative and significant, indicating a concave relationship in line with the BHM model.

Table C.15: Estimated coefficients for temperature effects (clustered SE)

Variable	**coef**	**std err**	**z**	**P>|z|**	**[0.025**	**0.975]**
Temperature	0.0572	0.016	3.594	0.000	0.026	0.088
Temperature^2	-0.0021	0.001	-2.040	0.041	-0.004	-8.04×10^{-5}

Diagnostic tests are summarised in Table C.16. The Durbin–Watson statistic (2.953) suggests no first-order autocorrelation. Jarque–Bera and kurtosis statistics indicate the presence of heavy tails, which is common in large panels. The high condition number (6.08×10^{18}) reflects model complexity and multicollinearity, mitigated through the use of cluster-robust standard errors.

Table C.16: Regression diagnostic statistics

Statistic	**Value**	**Statistic**	**Value**
Omnibus	116.362	Durbin-Watson	2.953
Prob(Omnibus)	0.000	Jarque-Bera (JB)	492.681
Skewness	-0.073	Prob(JB)	1.04×10^{-107}
Kurtosis	5.510	Condition Number	6.08×10^{18}

Table C.17 provides the estimated year fixed effects for this specification. The year 1960 and 1970 show strong positive deviations, while 2015 again shows a large negative deviation, consistent with previous models.

Table C.17: Estimated year fixed effects (reference year omitted)

Year	coef	std err	z	P>|z|	[0.025	0.975]
1925	−0.1427	0.024	−5.874	0.000	−0.190	−0.095
1938	−0.0079	0.025	−0.321	0.748	−0.056	0.040
1950	−0.0069	0.018	−0.386	0.699	−0.042	0.028
1960	0.2898	0.020	14.609	0.000	0.251	0.329
1970	0.2877	0.018	16.214	0.000	0.253	0.323
1980	0.1543	0.026	5.985	0.000	0.104	0.205
1990	−0.0877	0.016	−5.519	0.000	−0.119	−0.057
2000	0.0999	0.013	7.614	0.000	0.074	0.126
2010	−0.1000	0.014	−6.943	0.000	−0.128	−0.072
2015	−0.2196	0.016	−13.440	0.000	−0.252	−0.188

C.4.1. Jackknife robustness checks

The next three tables show the results of spatial, temporal, and country-level jackknife robustness checks. These tests investigate the sensitivity of the estimated coefficients by iteratively excluding key spatial units or time periods from the regression.

Table C.18: Spatial jackknife results for climate coefficients

Variable	Mean	Std. dev.	t-stat	p-value
Temperature	0.0738	0.0013	58.64	0.000
Temperature^2	-0.00313	0.000077	-40.76	0.000
Precipitation	-0.00029	0.000013	-22.28	0.000
Precipitation^2	0.0000000745	0.0000000062	11.97	0.000

Table C.19: Temporal jackknife results for climate coefficients

Variable	Mean	Std. dev.	t-stat	p-value
Temperature	0.0762	0.0330	2.31	0.021
Temperature^2	-0.00320	0.00146	-2.19	0.028
Precipitation	-0.00029	0.00014	-2.10	0.036
Precipitation^2	0.0000000722	0.0000000714	1.01	0.312

Table C.20: Country jackknife summary statistics for climate coefficients

Variable	Mean	Std. Dev.	t-statistic	p-value
Temperature	0.0739	0.0089	8.33	0.0000
Temperature^2	-0.0031	0.0006	-5.26	0.0000
Precipitation	-0.0003	0.0001	-4.19	0.0000
Precipitation^2	7.23e-08	2.16e-08	3.35	0.0008

D

Appendix D - Re-analysis and Extension of the BHM Model with Climate Variables

D.1. Constructing Rolling Means

This section documents exploratory work aimed at capturing long-run climate conditions through 30-year rolling means of temperature and precipitation. For each region i and year $t \geq 1930$, the rolling averages are defined as

$$\overline{T}_{i,t}^{30} = \frac{1}{30} \sum_{s=t-29}^{t} T_{i,s}^{\text{raw}}, \qquad \overline{P}_{i,t}^{30} = \frac{1}{30} \sum_{s=t-29}^{t} P_{i,s}^{\text{raw}}.$$

For earlier years, from 1900 to 1929, the averages are calculated using all available years up to t:

$$\overline{T}_{i,t}^{30} = \frac{1}{t-1900+1} \sum_{s=1900}^{t} T_{i,s}^{\text{raw}}, \qquad \overline{P}_{i,t}^{30} = \frac{1}{t-1900+1} \sum_{s=1900}^{t} P_{i,s}^{\text{raw}}.$$

This method avoids interpolation across wide time gaps and maintains the historical signal. To match the structure of the economic dataset by Rosés et al. (2021), the rolling means are calculated first on the full annual series and only then linked to the benchmark years. This ensures that each benchmark year reflects a full 30-year climatic context rather than a shorter or uneven period. Unlike BHM, who apply population weights at the grid-cell level, this analysis uses pre-aggregated regional averages. Since population-weighted re-aggregation is not possible after this step, further weighting would not change the values.

These rolling means are not included in the main analysis since they do not conceptually align with the rest of the research design. The model used in this thesis focuses on annual climate differences to estimate growth effects. Long-run means are not directly comparable to the dynamic structure of the benchmark growth data. Still, they are presented here for exploratory purposes.

D.2. Exploratory Data Analysis

To understand the properties of the climate variables and assess their suitability, an exploratory data analysis is carried out. Summary statistics are calculated for per capita GDP growth $\Delta \ln \text{GDPpc}_{i,t}$, rolling mean temperature $\overline{T}_{i,t}^{30}$, rolling mean precipitation $\overline{P}_{i,t}^{30}$, and sea level rise $\Delta\text{SLR}_{i,t}$, including squared terms. These results are included in Appendix A.

Pairwise correlations are computed between GDP growth and each climate variable, and scatterplots are used to explore the shape of the relationships. High correlations between each variable and its squared term are found: $r = 0.96$ for temperature and $r = 0.98$ for precipitation. These are expected due to their mathematical construction but raise concerns about multicollinearity in regressions. Correlations between GDP growth and each climate variable (including squared terms) range from -0.04 to $+0.03$, indicating that bivariate linear relationships are weak. This suggests that any effect of climate is likely concave or conditional on other variables, and therefore best analysed using multivariate regression models. The correlation between temperature and sea level rise is moderate, with $r = 0.38$, possibly reflecting shared global warming trends.

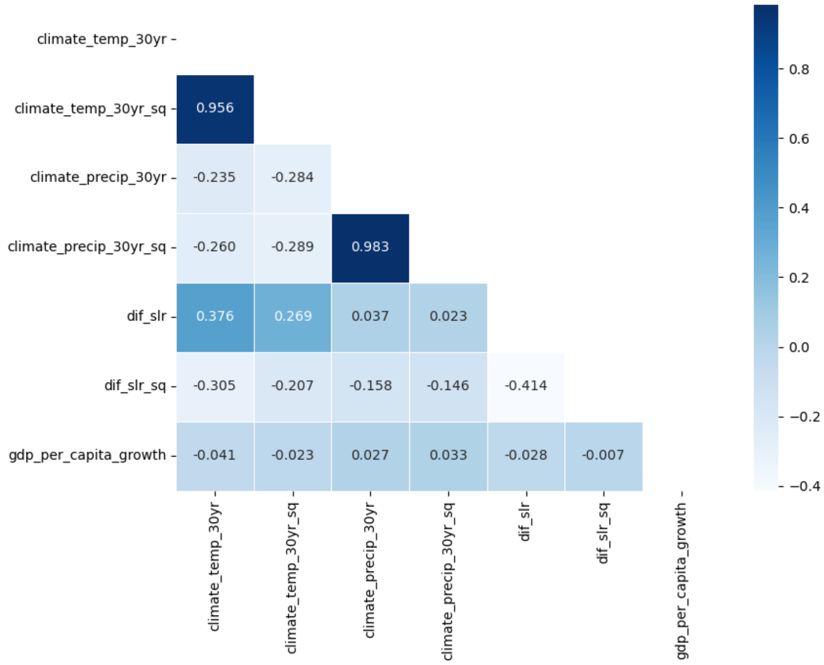

Figure D.1: Lower triangle of the pairwise correlation matrix between GDP per capita growth and climate variables

D.3. Regression annualised Log-Growth from benchmark GDP

To calculate economic growth between benchmark years, the natural logarithm of GDP per capita is used. For each region, the annualised log-growth rate is computed by subtracting the lagged log-GDP per capita from the current value and dividing by the year difference. This approach allows for irregular benchmark intervals and ensures that growth is expressed in annual terms:

$$\text{ln\_growth}_{it} = \frac{\ln(\text{GDP}_{it}) - \ln(\text{GDP}_{i,t-1})}{\text{year}_t - \text{year}_{t-1}}$$

D.3.1. Empirical strategy

The following panel regression model is estimated on data from 1900 to 2015:

$$\Delta \ln(\text{GDP}_{it}) = \beta_1 T_{it} + \beta_2 T_{it}^2 + \gamma_1 P_{it} + \gamma_2 P_{it}^2 + \text{region and year fixed effects} + \varepsilon_{it}$$

Here, T_{it} and P_{it} represent thirty–year rolling means of temperature and precipitation, respectively, for region i in year t.

D.3.2. Main Results

The regression is based on 1,875 observations and includes 524 parameters, leaving 1,350 degrees of freedom. The model explains about 57% of the variation in growth ($R^2 = 0.569$), with an adjusted $R^2 = 0.401$. The climate variables are clustered at the region level to account for intra-region correlation.

Table D.1: OLS regression GDP growth annualised summary statistics

Statistic	Value
Dependent variable	ln_growth
Model	OLS
No. Observations	1875
Degrees of freedom (model)	524
Degrees of freedom (residual)	1350
R^2	0.569
Adjusted R^2	0.401
F-statistic	1.195×10^{13}
Prob (F-statistic)	0.00
Log-Likelihood	5354.7
Covariance Type	Clustered

D.3.3. Coefficient estimates

Table D.2: Estimated coefficients for climate variables in regression GDP growth annualised

Variable	Coef.	Std. Err.	z	p	95% CI []
`climate_temp_30yr`	0.0306	0.010	3.074	0.002	0.011 – 0.050
`climate_temp_30yr_sq`	−0.0005	0.000	−1.320	0.187	−0.001 – 0.000
`climate_precip_30yr`	0.0001	0.000	0.980	0.327	−0.000 – 0.000
`climate_precip_30yr_sq`	-1.11×10^{-7}	8.06×10^{-8}	−1.377	0.168	-2.69×10^{-7} – 4.7×10^{-8}

The temperature coefficient is positive and significant, while its squared term is negative but not significant. This supports a concave relationship between temperature and growth, although the curve's shape is not precisely estimated. Precipitation terms are small and not statistically significant.

D.3.4. Year FE

Table D.3: Estimated year fixed effects GDP growth annualised (baseline omitted)

Year	Coef.	Std. Err.	t	p	95% CI Low	95% CI High
1925	-0.0158	0.002	-8.022	0.000	-0.020	-0.012
1938	-0.0065	0.002	-3.304	0.001	-0.010	-0.003
1950	-0.0070	0.002	-4.080	0.000	-0.010	-0.004
1960	0.0260	0.002	14.062	0.000	0.022	0.030
1970	0.0275	0.002	15.655	0.000	0.024	0.031
1980	0.0179	0.003	5.165	0.000	0.011	0.025
1990	-0.0030	0.003	-1.040	0.298	-0.009	0.003
2000	0.0120	0.002	7.777	0.000	0.009	0.015
2010	-0.0161	0.002	-10.298	0.000	-0.019	-0.013
2015	-0.0207	0.003	-6.676	0.000	-0.027	-0.015

Positive fixed effects in 1960 and 1970 likely reflect post-war economic growth, while negative values in 2010 and 2015 capture the effects of the financial crisis. Most years show statistically significant differences from the baseline, except 1990.

D.3.5. Temperature Optimum

The fitted growth curve is given by:

$$h(T) = \beta_1 T + \beta_2 T^2$$

with optimum temperature:

$$T^* = -\frac{\beta_1}{2\beta_2} = -\frac{0.0306}{2 \times (-0.000518)} = 29.6\,^\circ\mathrm{C}$$

This peak implies that regions experience maximum growth near 29.6 °C under long-term climate conditions. Beyond this, temperature has a negative effect.

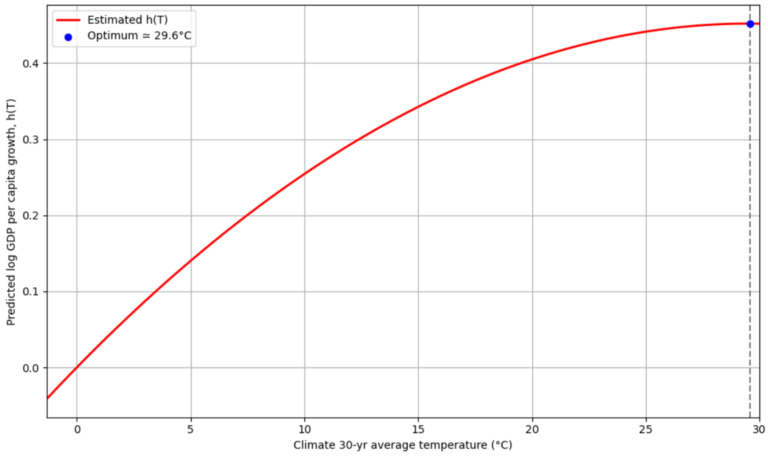

Figure D.2: Estimated long-run temperature optimum of 29.6 °C (blue dot)

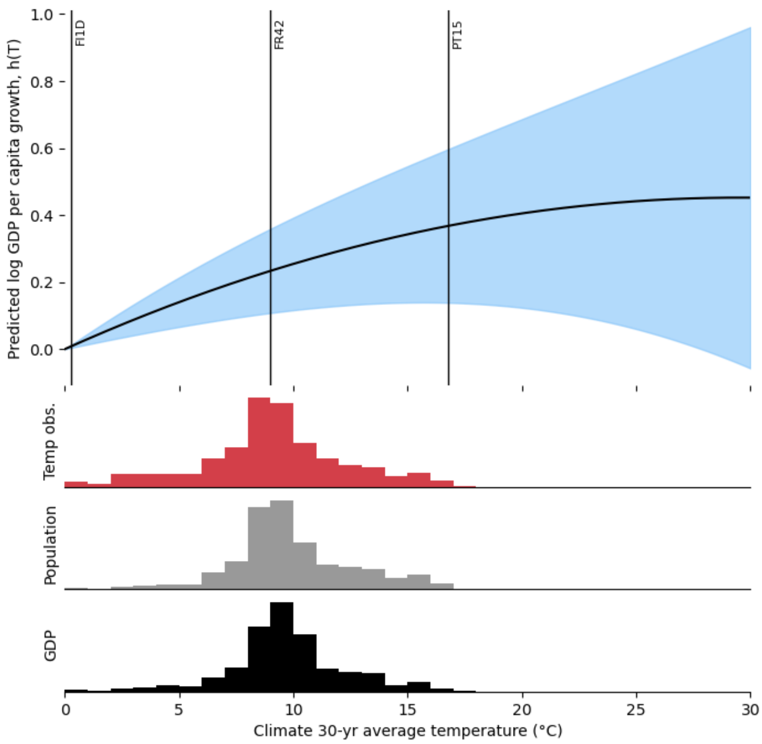

Figure D.3: Concave growth curve with 90% confidence band and distributions of temperature, population and GDP per capita

The confidence band is narrowest where data are dense (5–17 °C), and widens at extreme temperatures due to fewer observations. These patterns suggest that moderate warming can support growth in cooler regions, but excessive heat remains economically harmful.

D.3.6. Non-Annualised Growth Model Results

The results below are based on GDP per capita growth computed between benchmark years, without annualisation. The panel consists of 1,875 regional observations. Ordinary least squares regression is applied with fixed effects for year and region, and standard errors are clustered at the regional level.

Table D.4: OLS regression summary statistics (non-annualised growth)

Statistic	Value
Dependent variable	`ln_growth`
Model	OLS
Method	Least Squares
No. of observations	1,875
Degrees of freedom (model)	524
Degrees of freedom (residual)	1,350
R^2	0.519
Adjusted R^2	0.332
F-statistic	5.751×10^{12}
Prob(F-statistic)	0.00
Log-Likelihood	837.69
Covariance type	Clustered

The estimated coefficients for the climate variables are summarised below. The positive and significant linear temperature coefficient and the negative significant squared term support a concave relationship. Precipitation terms are statistically insignificant.

Table D.5: Estimated coefficients for climate variables (non-annualised model)

Variable	Coef.	Std. Err.	t	p	95% CI Low	95% CI High
climate_temp_30yr	0.4136	0.104	3.984	0.000	0.210	0.617
climate_temp_30yr_sq	-0.0080	0.004	-2.079	0.038	-0.015	-0.000
climate_precip_30yr	0.0010	0.002	0.642	0.521	-0.002	0.004
climate_precip_30yr_sq	-9.68×10^{-7}	8.93×10^{-7}	-1.084	0.278	-2.72×10^{-6}	7.82×10^{-7}

Year fixed effects are included to control for macroeconomic shocks and long-run trends. The coefficients for 1960 and 1970 are large and positive, indicating robust growth in those years. The years 2010 and 2015 show significant negative effects, consistent with post-crisis stagnation.

Table D.6: Year fixed effects in non-annualised GDP growth model (baseline omitted)

Year	Coef.	Std. Err.	t	p	95% CI Low	95% CI High
1925	-0.1421	0.026	-5.492	0.000	-0.193	-0.091
1938	-0.0093	0.024	-0.389	0.697	-0.056	0.038
1950	-0.0213	0.020	-1.069	0.285	-0.060	0.018
1960	0.2883	0.019	15.049	0.000	0.251	0.326
1970	0.3123	0.019	16.028	0.000	0.274	0.350
1980	0.2235	0.038	5.935	0.000	0.150	0.297
1990	0.0094	0.031	0.299	0.765	-0.052	0.071
2000	0.1399	0.016	8.709	0.000	0.108	0.171
2010	-0.1610	0.017	-9.552	0.000	-0.194	-0.128
2015	-0.2909	0.034	-8.629	0.000	-0.357	-0.225

The annualised and non-annualised regressions both show a concave temperature–growth relationship, but they differ in both statistical precision and model fit. In both models, the linear temperature coefficient is positive and significant, while the squared term is negative, supporting the concave structure proposed by BHM. However, the curvature is more precisely estimated in the non-annualised model, where both terms are statistically significant ($\beta_1 = 0.4136$, $\beta_2 = -0.0080$, $p = 0.038$), whereas the squared term in the annualised model is not ($\beta_2 = -0.0005$, $p = 0.187$). Despite these differences, the turning point estimates are similar, and both models indicate that growth is maximised around $T^* \approx 29.6\,°\mathrm{C}$ under long-run temperature exposure. Model fit statistics reveal another important distinction: the annualised regression achieves a substantially higher log-likelihood (LL = 5354.7) than the non-annualised one (LL = 837.69), and explains a larger share of the variation in growth ($R^2 = 0.569$ versus $R^2 = 0.519$). This suggests that annualising growth may reduce noise caused by irregular time intervals and improve statistical performance. Nevertheless, both specifications produce consistent results: temperature remains the dominant climate variable, precipitation has no significant effect, and year fixed effects reflect known macroeconomic shifts. The choice between annualised and non-annualised growth affects the precision and efficiency of the estimates but not the overall conclusions about the climate–economy relationship.

D.4. Climate temporal and spatial jackknife analysis (non-annualised)

To assess the robustness of the estimated long-run climate-growth relationship, a baseline OLS panel regression model is first estimated. The model relates GDP per capita growth to quadratic terms in 30-year average temperature and precipitation, capturing potential non-linear effects. It includes region fixed effects to control for time-invariant unobserved heterogeneity across NUTS regions, year fixed effects to account for global shocks, and region-specific linear and quadratic time trends to absorb heterogeneous long-term trajectories. The regression is specified as follows:

$$\begin{aligned}\ln(\text{growth}_{it}) = {} & \beta_1 \cdot \text{climate\_temp\_30yr}_{it} + \beta_2 \cdot \text{climate\_temp\_30yr\_sq}_{it} \\ & + \beta_3 \cdot \text{climate\_precip\_30yr}_{it} + \beta_4 \cdot \text{climate\_precip\_30yr\_sq}_{it} \\ & + \alpha_i + \gamma_t + \delta_i \cdot t + \theta_i \cdot t^2 + \varepsilon_{it}\end{aligned} \tag{D.1}$$

Here, α_i denotes region fixed effects, γ_t year fixed effects, and δ_i and θ_i allow for region-specific linear and quadratic trends over time. Standard errors are clustered at the regional level to account for serial correlation within NUTS units. Based on this baseline model, a jackknife procedure is conducted to test the sensitivity of the results. In the spatial jackknife, the model is re-estimated repeatedly, each time excluding one NUTS region. In the temporal jackknife, the same procedure is applied by sequentially omitting one year at a time. The resulting variation in the estimated temperature-growth curves is used to evaluate the robustness and stability of the non-linear climate response across space and time.

Table D.7: Jackknife summary of climate coefficient estimates

Variable	Spatial Mean	Spatial Std	Temporal Mean	Temporal Std
`climate_temp_30yr`	4.131×10^{-1}	7.06×10^{-3}	4.357×10^{-1}	1.798×10^{-1}
`climate_temp_30yr_sq`	-7.949×10^{-3}	2.63×10^{-4}	-7.704×10^{-3}	4.669×10^{-3}
`climate_precip_30yr`	1.051×10^{-3}	1.09×10^{-4}	1.125×10^{-3}	1.516×10^{-3}
`climate_precip_30yr_sq`	-9.803×10^{-7}	6.18×10^{-8}	-1.000×10^{-6}	8.46×10^{-7}

Table D.7 presents the jackknife estimates of the spatial and temporal stability of the climate coefficients. The spatial mean of the 30-year average temperature coefficient is approximately 0.413, with a low spatial standard deviation of 0.007, indicating consistent estimates across different spatial subsets. Its squared term shows a negative mean of −0.00795, also with a small standard deviation, suggesting a robust concave relationship between temperature and economic growth. The temporal means, derived by sequentially omitting different time periods, show a similar pattern: the temperature coefficient has a slightly higher mean of 0.436 and a larger standard deviation of 0.180, reflecting more sensitivity to changes over time than across space. The squared term remains negative with a mean of −0.0077 and a standard deviation of 0.0047, again supporting the presence of a stable non-linear temperature-growth relationship. For precipitation, the spatial mean is 0.00105, with a small standard deviation, while its squared term is slightly negative. Temporally, the precipitation coefficient shows a mean of 0.00113 and slightly higher variability. The squared term for precipitation is near zero in both spatial and temporal dimensions, with low standard deviations, suggesting a minimal but stable non-linear effect. Overall, the results confirm that the estimated climate coefficients are generally robust across both spatial and temporal jackknife replications, with consistent signs and magnitudes that support the existence of a non-linear climateeconomy relationship.

Figure D.4 illustrates the robustness of the estimated non-linear temperature growth relationship under the spatial jackknife procedure. The plot shows how the predicted change in log GDP ($\Delta \ln(\text{GDP})$) varies with the 30-year average temperature, comparing the full-sample estimate (black dashed line) to estimates obtained by systematically omitting individual regions.

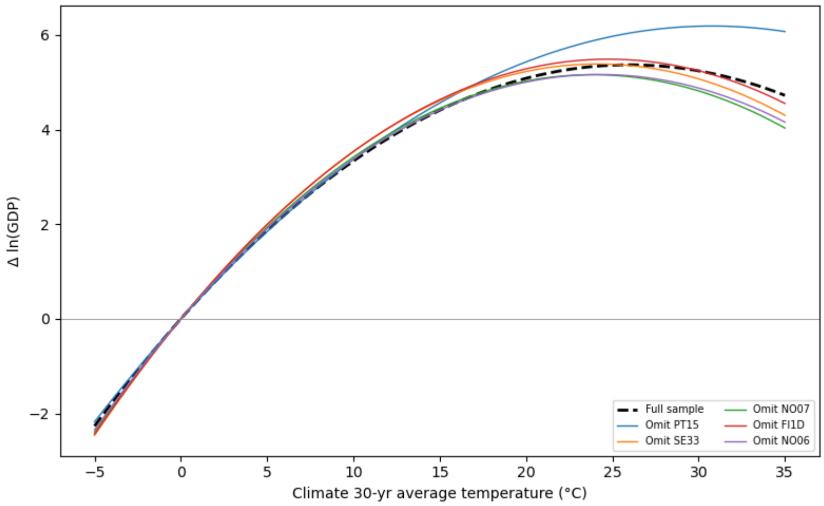

Figure D.4: Spatial jackknife curves of the estimated climate-growth relationship. Each line represents the predicted change in log GDP as a function of 30-year average temperature, estimated by omitting one region at a time. The dashed black line is the estimate using the full sample.

The curves all follow a concave pattern, confirming the existence of a global non-linear relationship where economic productivity increases with temperature up to an optimum point, after which it declines. The general similarity across the jackknife curves indicates that the estimated response is not driven by any single region. However, some variation is visible in the upper range of the temperature distribution, particularly for the curve omitting region NO07, which produces a slightly higher predicted growth at extreme temperatures. Figure D.5 presents the results of the temporal jackknife analysis, where each year except 1910 is excluded once from the sample and the model is re-estimated. The curves display the predicted relationship between long-run temperature and economic growth, showing how the estimated climate-growth function shifts depending on which individual year is left out.

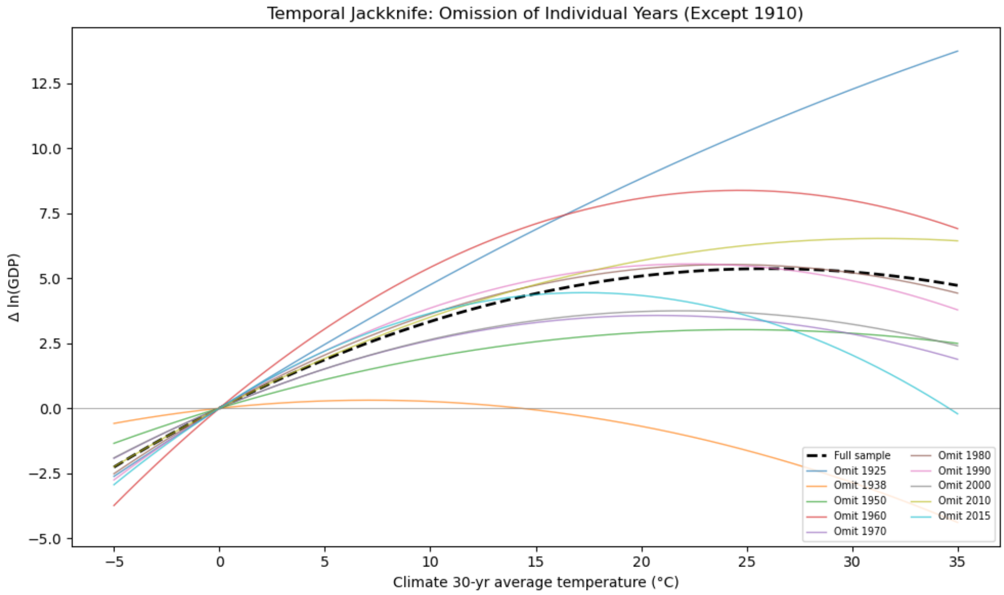

Figure D.5: Temporal jackknife estimates of the temperature-growth relationship. Each line shows the predicted change in log GDP as a function of 30-year average temperature, estimated by omitting one year at a time. All years except 1910 are omitted once. The dashed black line shows the full-sample estimate.

The dashed black line represents the full-sample estimate, which features a concave shape. Most omitted-year curves follow a similar pattern, although with notable differences in slope and curvature. For instance, omitting 1938 or 1925 results in strongly upward-sloping curves with little evidence of curvature, while omitting 1960 or 2015 leads to a more pronounced peak and faster decline at higher temperatures. The most dramatic deviation is seen when 1950 is excluded, producing a curve that flattens and turns negative across most of the temperature range.

D.4.1. annual temperature instead of climate means

Table D.8: Jackknife summary of climate coefficient estimates

Variable	Spatial Mean	Spatial Std	Temporal Mean	Temporal Std
Temperature	5.990760×10^{-3}	1.220536×10^{-4}	6.242053×10^{-3}	2.951481×10^{-3}
Temperature sq	-2.389015×10^{-4}	7.581468×10^{-6}	-2.481640×10^{-4}	1.466765×10^{-4}
Precipitation	-2.490379×10^{-5}	1.217707×10^{-6}	-2.472306×10^{-5}	1.403412×10^{-5}
Precipitation sq	5.338095×10^{-9}	5.849296×10^{-10}	5.176842×10^{-9}	7.295679×10^{-9}

Table D.8 reports jackknife summary statistics for the four climate coefficients. Spatial mean and spatial standard deviation give the average and dispersion of each coefficient when omitting one region at a time. Temporal mean and temporal standard deviation give the average and dispersion when omitting one year at a time. Values are shown in scientific notation.

The spatial means for both temperature and temperature squared closely match the fullsample estimates, and their spatial standard deviations are very small. This indicates that no single region drives the overall temperature effects. In contrast, temporal standard deviationsespecially for the linear temperature termare an order of magnitude larger than their spatial counterparts. This suggests that the estimated temperature-growth relationship is more sensitive to the exclusion of particular years than to particular regions. Precipitation coefficients show low spatial and temporal variability, implying robust linear and quadratic precipitation effects across both space and time. Overall, the results demonstrate strong regional robustness alongside notable year-to-year variation in temperature sensitivity.

Figure D.6 shows the full sample temperature-growth curve as a black dashed line and five jack-knife curves in colour. Each coloured line represents the estimated relationship when one region is omitted

(ITG1, SE33, FI1D, NO02, NO07). The horizontal axis is the annual average temperature in °C and the vertical axis is the predicted change in log GDP per capita.

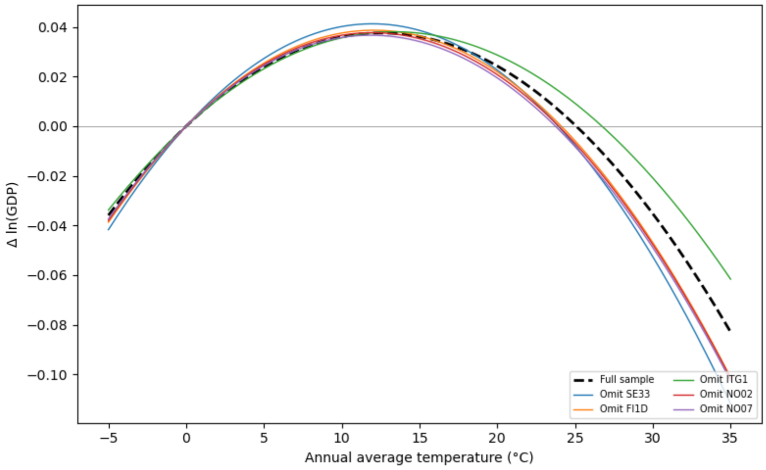

Figure D.6: Spatial Jack-Knife: drop influential regions

Omission of region ITG1 (green) shifts the peak to a higher temperature and flattens the curve at warm values, indicating that this region pulls the global optimum downward. The curves for SE33 (blue), FI1D (red) and NO07 (purple) lie close to the full sample estimate, with only minor deviations in curvature and peak height, suggesting those regions have limited influence. Excluding NO02 (orange) produces a slightly lower peak and steeper decline at high temperatures, showing some sensitivity to that regions data. Overall robustness of most curves around the dashed line confirms that the non–linear temperature-growth relationship is not driven by any single region, although ITG1 and NO02 exert noticeable influence on the estimated optimum and curvature.

Figure D.7 shows the full sample temperature-growth curve as a black dashed line and five jack-knife curves in colour. Each coloured line represents the estimated relationship when one influential year is omitted (2010 in blue, 1925 in red, 1950 in orange, 2015 in green, 1970 in purple). The horizontal axis is the annual average temperature in °C, and the vertical axis is the predicted change in log GDP per capita.

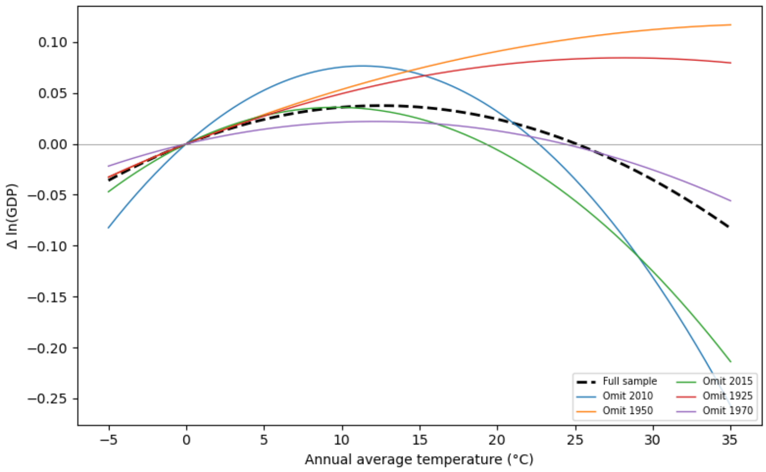

Figure D.7: Temporal Jack-Knife: drop influential years

Omitting year 2010 (blue) produces a pronounced peak near 10°C and a steep decline at higher temperatures, indicating that early twenty-first-century data dampen curvature. Exclusion of 1925 (red) shifts the peak upward to around 25°C and flattens the decline, showing that interwar data pull the optimum downward. Dropping 1950 (orange) yields a nearly monotonic increase in growth with temperature, suggesting that mid-century observations are crucial for detecting any downturn. Omission of 2015 (green) lowers the apex slightly and steepens the descent beyond the optimum, implying that recent warming years soften the downturn. Excluding 1970 (purple) produces only minor adjustments, indicating limited influence. Overall, sensitivity to specific years underscores the importance of temporal coverage in estimating the non–linear climate-growth relationship.