

The London School of Economics and Political Science

Empirical Essays on the Economics of Climate Change

Manuel Linsenmeier

June 2022

A thesis submitted to the Department of Geography and Environment
of the London School of Economics and Political Science for the
degree of Doctor of Philosophy, London, June 2022

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others.

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 45,000 words.

Statement of co-authored work

Chapters 1 and 2 are single-authored. Chapters 3, 4, and 5 are co-authored.

Chapter 3 is co-authored with Ben Groom and Sefi Roth. My contribution amounts to 60% of the chapter.

Chapters 4 and 5 are co-authored with Adil Mohommad and Gregor Schwerhoff. My contribution amounts to 80% for each chapter.

Abstract

Climate change has been referred to as the world's largest externality, motivating research and policies that in recent years appear to have gained additional momentum. This thesis compiles five empirical essays on the economics of climate change. The first three chapters study the costs of climate change. The last two chapters examine policies to reduce greenhouse gas emissions. More specifically, the first two chapters identify causal effects of temperature variability to examine its possible costs under scenarios of future climate change. The third chapter studies the impacts of weather shocks in Europe paying particular attention to their heterogeneity by industry and average climate. The chapters apply novel strategies for causal identification and report evidence on new channels through which climate change affects society. The fourth chapter empirically studies the sequencing of mitigation policies by instrument type and the association between sequencing and the adoption of carbon pricing policies. The fifth chapter examines the international diffusion of carbon pricing policies and quantifies its contribution to global greenhouse gas emission reductions. Both chapters report novel evidence that speaks to current debates in research and policy about how to limit global warming.

Acknowledgements

This thesis could not have been completed without the guidance, support, and company of many people. I am deeply grateful to my supervisors Simon Dietz, Sefi Roth, and James Rising for their direction and support.

I would also like to thank my co-authors Ben Groom, Adil Mohommad, and Gregor Schwerhoff for advice, help, and encouragement along the way.

My work benefitted greatly from discussions with colleagues at the Department of Geography and Environment and the Grantham Research Institute on Climate Change and the Environment at the LSE. I am particularly grateful to Mook Bangalore, Ana de-Menezes Silva, Roger Fouquet, and Glen Gostlow. Invaluable feedback was also received at numerous seminars, workshops, and conferences where the work was presented.

Even the greatest destination cannot replace good company on a journey of three years. I am thankful to many colleagues at LSE who became friends along the way. I would like to particularly thank Sanchayan Banerjee, Louise Bernard, Ghassane Benmir, Glen Gostlow, Timo Leiter, Julien Picard, and Sandro Provenzano. I am also grateful to Antonio Avila, Giulia Romano, Anna Schulze-Tilling, Francesco Scotti, and Sarah Spycher, for their company on segments of this trip.

Finally, and above all, I am indebted to my family and especially to my partner in love and economics, Tatsiana.

Funding Acknowledgement

I gratefully acknowledge financial support from the UK Economic and Social Research Council and the Grantham Research Institute on Climate Change and the Environment.

Contents

Introduction	16
1 Temperature variability and long-run economic development	23
1.1 Introduction	24
1.2 Climate variability	28
1.2.1 How temperature variability can affect economic ac- tivity	28
1.2.2 Temperature variability: day-to-day, seasonal, and interannual	29
1.3 Data and descriptive statistics	34
1.3.1 Economic variables	34
1.3.2 Climate variables	36
1.3.3 Geographical covariates	37
1.3.4 Descriptive statistics	38
1.4 Econometric strategy	38
1.5 Results	42
1.5.1 Main results	42
1.5.2 Sensitivity analysis	45
1.5.3 Reverse causality	46
1.6 Mechanisms	47
1.6.1 Urban areas	47
1.6.2 Agriculture	48
1.7 Conclusion	48
2 Seasonal temperature variability and economic cycles	52
2.1 Introduction	53
2.2 Methods	56
2.2.1 Theoretical framework	56
2.2.2 Identification strategy	57

2.2.3	Data	61
2.3	Results	66
2.3.1	Stylised facts on seasonal economic cycles	66
2.3.2	The contribution of seasonal temperature variability	71
2.3.3	Effects by industry groups for European economies .	73
2.3.4	Results from long differences	74
2.3.5	Scenarios of future climate change	77
2.4	Conclusion	78
3	Some like it cold: The persistent cost of higher temperatures in the economic sectors of Europe	82
3.1	Introduction	83
3.2	Methods	85
3.2.1	Econometric framework	85
3.2.2	Data	87
3.3	Results	87
3.3.1	Main findings	87
3.3.2	The impact of temperature change on GVA across re- gions and sectors	90
3.3.3	Robustness	95
3.4	Conclusion	96
4	Policy sequencing towards carbon pricing: Empirical evidence from G20 economies and other major emitters	98
4.1	Introduction	99
4.2	Methods	102
4.2.1	Empirical framework	102
4.2.2	Data	103
4.2.3	Statistical methods	106
4.3	Results	108
4.3.1	The temporal sequence of climate policies	108
4.3.2	The build-up of climate policy portfolios over time .	110
4.3.3	Policy portfolios of adopters versus non-adopters of carbon pricing	111
4.3.4	Heterogeneity among adopters of carbon pricing . .	112
4.4	Discussion and Conclusions	113
5	The international diffusion of policies for climate change mit-	

igation	116
5.1 Introduction	117
5.2 Methods	120
5.2.1 Empirical analysis of policy diffusion	120
5.2.2 Modelling the effect of policy diffusion on GHG emissions	124
5.2.3 Data	125
5.3 Results	126
5.3.1 Descriptive evidence	126
5.3.2 Model estimates	128
5.3.3 Emission reductions	134
5.4 Discussion and Conclusions	139
Concluding remarks	146
References	164
A Appendix of Chapter 1	165
A.1 Map of annual mean temperature	165
A.2 Omitted variable biases and multicollinearity	166
A.3 Estimated coefficients across bins of annual mean temperature	169
A.4 Alternative measures for seasonal variability	170
A.5 Discontinuities of treatment	171
A.6 Climate data for 1955-1984	175
A.7 Nightlights for 2015-2019	176
A.8 Sensitivity analysis	177
A.9 Control variables	180
A.10 Urban versus rural areas	182
A.11 Agricultural land use	183
B Appendix of Chapter 2	184
B.1 Detrending of time-series	184
B.2 Summer peak vs winter peak countries	185
B.3 Descriptive evidence on statistical associations	186
B.4 Estimation results to explore alternative channels	187
B.5 Specification chart	188
B.6 Results by industry group	189
B.7 Estimation results for different samples	190
B.8 Robustness check with seasons based on quarters	191

B.9	Past climate trends	192
B.10	Estimation results by season	193
B.11	Climate projections	194
B.12	Projected changes to seasonal economic cycles for RCP4.5 and RCP8.5	196
B.13	Map of projected changes to seasonal economic cycles . . .	197
C	Appendix of Chapter 3	198
C.1	Data	198
C.2	Additional results	201
C.2.1	Spatial spill-overs	201
C.2.2	Weighted industry results	202
C.2.3	Seasonal effects	202
C.2.4	Adaptation	202
C.2.5	Degree-day model	203
C.3	Additional robustness checks	204
C.3.1	Sample of countries and years	204
C.3.2	Alternative standard errors	204
C.3.3	Alternative weather data	206
D	Appendix of Chapter 4	207
D.1	Causality and the empirical framework	207
D.2	Descriptive statistics	208
D.3	Results from regression analysis	209
D.4	Policy portfolios by the end of 2020	211
D.5	Sample of countries	212
D.6	Instrument types and sectors	214
D.7	Instrument types and instrument categories	215
E	Appendix of Chapter 5	219
E.1	Quantification of the global benefits of diffusion	219
E.1.1	Back-of-the-envelope calculations	219
E.1.2	Monte-Carlo simulations	221
E.2	Additional results	222

List of Figures

1.1	Calculation of my three measures of temperature variability: day-to-day, seasonal, and interannual variability.	31
1.2	Geographical distribution of temperature and its variability: day-to-day, seasonal, and interannual variability	33
1.3	Geographical distribution of VIIRS nightlights	35
2.1	Time-series of quarterly GDP for the USA before and after detrending	62
2.2	Map of the size of the seasonal economic cycle as percentage of annual GDP	64
2.3	Frequency of patterns of quarterly economic production	67
2.4	Stylised facts identified in previous studies and their frequency in the sample of countries	68
2.5	Frequency of patterns of economic production in summer and winter	70
2.6	Geographical distribution of winter-peak and summer-peak countries	71
2.7	Projections of Δ GDP for alternative scenarios of future climate change by country	78
3.1	Main estimate (quadratic model) and the underlying variation in the data	89
3.2	Marginal effects of an increase in annual mean temperature obtained with alternative model specifications	90
3.3	Geographical distribution of estimated marginal effects of annual mean temperature across Europe	92
3.4	Estimated marginal effects by industry at three different levels of temperature	93
4.1	Causal diagram illustrating alternative pathways of causality	103

4.2	Map of the sample of countries included in the analysis . . .	104
4.3	Adoption of policies with different instrument types and sectors over time in different countries	107
4.4	Policy sequences of countries with a carbon price by industry and by country	109
4.5	Development of countries' climate policy portfolios over time	110
4.6	Scatterplots of statistical associations between the year of the adoption of carbon pricing, the size of climate policy portfolios in that year, and the average carbon price in the first year of implementation	112
5.1	Map with the time of adoption of the first carbon pricing policy by country	126
5.2	Descriptive evidence on possible leader-follower relationships among adopters of carbon pricing	127
5.3	Direct and indirect emission reductions from a back-of-the-envelope calculation	135
5.4	Direct and indirect emission reductions from Monte Carlo simulations	137
5.5	Frequency distribution of emission reductions from Monte Carlo simulations	137
5.6	Geographical coverage of carbon pricing policies from Monte Carlo simulations for 2020-2050 with and without diffusion	139
A.1	Global map of annual mean temperature	165
A.2	Estimated marginal effects of annual mean temperature and day-to-day, seasonal, and interannual temperature variability at different levels of annual mean temperature	169
A.3	Histograms of first-differences in temperature variability . .	171
B.11	Time-series of seasonal production for the USA and Brazil .	184
B.31	Scatter plot of seasonal differences of temperature and log GDP	186
B.51	Specification chart	188
B.91	Change in seasonal temperatures between 1981-2000 and 2001-2020 for winter and summer months	192
B.111	Future projections of ΔT for RCP4.5 for 2041-2070 and 2071-2100	194

B.112	Future projections of ΔT for RCP8.5 for 2041-2070 and 2071-2100	195
B.121	Projections of Δ GDP for two alternative scenarios of future climate change	196
B.131	Map of projections of Δ GDP for the RCP8.5 scenario	197
C.11	Map of the spatial distribution of annual mean temperature	198
C.12	Map of the spatial distribution of Gross Value Added in Europe	199
C.21	Estimated marginal effects of model with spatial lag	201
C.22	Estimated marginal effects by industry at three different levels of temperature	202
C.23	Estimated marginal effect of seasonal temperature at two levels of annual mean temperature	202
C.24	Curvature of predicted effects estimated with polynomial degree-day models	203
C.25	Estimated effects from degree-day model	203
C.31	Estimated marginal effect of increasing annual mean temperature by one unit for subnational GVA using two alternative weather data sets	206
D.11	Causal diagram with possibly confounding country characteristics	207
D.41	Carbon pricing policies and the number of instrument types in countries' policy portfolios	211
D.61	Frequency of instrument types in different sectors in the sample of countries	214
E.21	Cumulative baseline hazard of the Cox proportional hazard model	222
E.22	Scatter plot of economy-wide emission-weighted average carbon prices over time	222
E.23	Map with the time of adoption of the first carbon tax policy by country	223
E.24	Map with the time of adoption of the first ETS policy by country	223
E.25	Map of the sample of 179 countries used in this study	223
E.26	Histogram of estimated baseline hazard adjusted for covariates in 2020 for the sample of 179 countries	224

List of Tables

1.1	Descriptive statistics	37
1.2	Results of a linear model estimated with SFD	44
1.3	Results of regressions addressing concerns of reverse causality	47
2.1	Descriptive statistics	65
2.2	Results of regressions using a global sample of GDP of 81 countries	72
2.3	Results of regressions using a sample of GVA by industry groups of 35 European economies	74
2.4	Results of regressions with long differences	76
5.1	Descriptive statistics	126
5.2	Results of estimation of Cox proportional hazard models with different metrics used for the construction of the spatial lag	130
5.3	Results of estimation of Cox proportional hazard models with different policies, with stratification and an additional control variable, and with a placebo spatial lag	133
A.1	Results of a simple model estimated with levels and SFD	166
A.2	Variance inflation factors	167
A.3	Pearson correlation coefficients	168
A.4	Results of a model estimated with SFD with two alternative measures for seasonal variability of temperature	170
A.5	Results of a model estimated with SFD with subsamples	172
A.6	Results of a model estimated with SFD with and without the inclusion of spatial lags	173
A.7	Results of a model estimated with SFD where spatial differences are calculated by skipping Δi observations	174

A.8	Results of a model estimated with SFD for different climate periods	175
A.9	Results of a model estimated with SFD using different versions and time periods of the nightlights data	176
A.10	Results of the sensitivity analysis	179
A.11	Results of models estimated with different control variables .	180
A.12	Results of regressions with subsampling observations based on population density	182
A.13	Results of regressions to examine the role of agriculture . . .	183
B.21	Results of balancing tests for summer-peak and winter-peak countries	185
B.41	Results of regressions to explore possible channels through which temperature variability might affect seasonal economic cycles	187
B.61	Results of regressions using a sample of GVA by industry groups of 35 European economies	189
B.71	Results of regressions using the global samples of GDP of countries that are used for the cross-sectional and the long differences estimation.	190
B.81	Results of regressions using a global sample of GDP of countries using differences between quarter with maximum and minimum temperature for identification	191
B.101	Results of regressions with long differences by season	193
C.11	Definition of industry groups used in the analysis	200
C.12	Descriptive statistics	201
C.31	Results of regressions with different samples of countries and years using the growth rate of national GDP per capita as dependent variable	204
C.32	Results of regressions with different fixed effects, time trends, and Conley standard errors	205
D.21	Descriptive statistics	208
D.31	Results of a logistic regression on the adoption of carbon pricing	209
D.32	Results of linear regression on the size of climate policy portfolio at the time of adoption of a national carbon pricing policy	210

D.51	List of countries included in the main analysis and their carbon pricing policies by the end of 2020	212
D.52	List of countries not included in the main analysis because of insufficient data on policy adoption, but which had a national carbon price implemented by the end of 2020	213
D.71	List of instrument categories, their instrument type, and their frequency	216
D.72	List of instrument categories, their instrument type, and their frequency (cont.)	217
D.73	List of instrument categories, their instrument type, and their frequency (cont.)	218
E.21	Results of estimation with different lag times	224

Introduction

The “credibility revolution” in empirical economics, a term sometimes used to refer to the development of new econometric techniques and a new emphasis on research design since the late 1980s, has had a lasting imprint on the economics of climate change. Equipped with new methods and better data, economists have quickly improved our understanding of how weather and climate affect societies. Major advances include the detection of many more channels through which climate influences economic activity, growing evidence on non-linearities in the effects of temperature, and an increasing awareness of the conceptual differences between the response of an economy to weather events and to gradual changes in climate.

This thesis contains empirical work in the spirit of this relatively new field of climate econometrics. The thesis consists of two parts. The first three chapters focus on the economic consequences of climate change. They apply novel strategies to identify the economic costs of climate variability and report new evidence on how climate change might affect economies in the future. The remaining two chapters use empirical methods to study the constraints and benefits of policies to mitigate climate change. They improve our understanding of two specific dimensions of climate policy adoption, the role of climate policy sequencing and the importance of international policy diffusion in countries’ efforts to reduce greenhouse gas emissions.

Temperature fluctuates a lot from year to year, season to season, and day to day. This variability is potentially important because evidence suggests that daily temperature levels have a non-linear effect on many socioeconomic outcomes and because larger variability generally means larger uncertainty. However, most estimates of the costs of future climate change are based only on projections of annual mean temperature. Chapter 1 hence studies the economic costs of temperature variability at interannual, seasonal, and day-to-day time scales. For the first time, the chapter examines the effects of

temperature variability at several time scales in one empirical framework. It is also the first assessment of the costs of temperature variability at seasonal and interannual time scales that examines the effect of the average climate as opposed to the effect of weather events. For identification, I use a novel estimation strategy based on spatial first-differences that reduces concerns about omitted variable biases. My results suggest that larger temperature variability comes at a cost at most time scales in most regions. Importantly, some of those costs are projected to be particularly large under future climate change in relatively warm and relatively poor regions.

Seasonality is common in time-series of GDP, but despite conjectures that annually recurring fluctuations of production are to some extent driven by temperature, economists have so far not been able to attribute seasonal economic cycles to seasonal temperature variability. The relationship between the two has gained additional relevance from projections of future climate change that suggest that in most countries some seasons are projected to warm more than others. In Chapter 2, I re-examine this relationship using data with better geographical coverage and longer time-series than previous work. Most importantly, I propose a new identification strategy that accounts for anticipation of seasonality. Contrary to previous findings, my results suggest that temperature differences between summer and winter can significantly explain differences in GDP between the two seasons. Furthermore, I find that climate change has already affected seasonal economic cycles since 1981. The effect of temperature on production seems to be due to industries in which labour is relatively more exposed to ambient temperatures and is smaller in richer countries. Projections of future climate change suggest additional re-allocation of economic production between the seasons, with larger seasonal economic cycles in the future in many countries.

Most prior studies of the effect of temperature shocks on economic growth have been on the level of countries and without industry dis-aggregation. This is potentially problematic, as both climate and economic production exhibit large heterogeneity at the subnational level and industries might respond differently to the same shock. Chapter 3, co-authored with Ben Groom and Sefi Roth, studies how unusually warm and cold years affect economic production in Europe. To do so, we use geographically granular data on Gross Value Added at the level of industries. We pay particular attention to heterogeneity with respect to regions and industries to improve our understanding of how weather affects economic production in specific contexts.

Contrary to previous work based on global samples of countries or regions, we find that warmer-than-average years are particularly costly in relatively cold regions. We can attribute this effect to agriculture, manufacturing, and mining and utilities. In relatively warm districts, a negative effect of higher annual mean temperatures on GVA in trade and other services is offset by positive effects in other industries. Furthermore, we find evidence for adaptation to days with very cold and with very hot temperatures.

Since the development of new empirical methods, researchers in the field of climate econometrics have increasingly combined annual observations of economic production with annual observations of temperature to identify the effect of weather on economic outcomes (Schlenker et al. (2006); Deschênes and Greenstone (2011); Dell et al. (2012), among others). Identification is obtained from the quasi-random year-to-year fluctuations of weather. This identification strategy is also pursued in Chapter 3. However, while the the marginal effect of fluctuations of weather can be the same as the marginal effect of gradual changes of climate under certain assumptions, the former will generally not adequately account for adaptation and therefore under- or over-estimate the true effect of climate change (Hsiang, 2016). The marginal effect of temperature variability, especially at seasonal and interannual time scales, is particularly difficult to identify from variations of weather. This is because of a mismatch of time scales in the case of interannual variability and a possible conflation of the effect of seasonal variability and the effect of seasonal temperature levels. The first two chapters of this thesis therefore also make important methodological contributions on how one can estimate the effect of climate variability on economic outcomes.

The empirical evidence in the first three chapters is important for adaptation. The results suggest that future changes to climate variability will incur additional costs of climate change (and benefits in some places). This has previously been reported only by Kotz et al. (2021b) for variability at the daily time scale. Chapters 1 and 3 also contribute to the emerging field of the economic geography of climate change (Alvarez and Rossi-Hansberg (2021); Rudik et al. (2021), among others) emphasising the importance of accounting for heterogeneity of local climatic and socioeconomic contexts and for higher order moments of the temperature distribution. Chapter 2 contributes to a series of publications in macro that have attempted to understand the drivers of seasonal economic cycles (Barsky and Miron (1989); Beaulieu et al. (1992); Beaulieu and Miron (1992), among others). The results suggest that

anticipation plays an important role in the effect of temperature on these cycles and that temperature might be a fundamental driver of preference and technology shocks identified in the prior literature. My findings also suggest that climate change will lead to a reallocation of economic production between the seasons.

Economists commonly consider carbon pricing as the most economically efficient policy to reduce greenhouse gas emissions, but in practice many countries have adopted alternative policies such as technology subsidies which are typically associated with higher costs. This discrepancy has been explained with the theory of climate policy sequencing, which suggests that countries' earlier policies, such as subsidies, can help to remove barriers to later policies with higher stringency, such as carbon taxes or tradable emission permits. This theory has received a lot of attention, but it has so far been without comprehensive international evidence except case studies. In Chapter 4, Adil Mohommad, Gregor Schwerhoff, and I examine the sequencing of policies to reduce greenhouse gas emissions. To do so, we use an international database of climate policies including information on seven instrument types and six sectors and focus on countries that have adopted carbon pricing in the last 20 years. We find that carbon pricing has tended to be adopted after the adoption of all other instrument types. Furthermore, countries that adopted carbon pricing in a given year had larger climate policy portfolios than those that did not. We also find that the size of portfolio at the time of adoption declined over time and that the larger the portfolio at the time of adoption, the higher the initial stringency of the pricing policy.

Many countries might be reluctant to adopt carbon pricing because of a perceived limited effectiveness of domestic policies in relatively small economies and because of concerns about lower international competitiveness and leakage. In Chapter 5, co-authored with Adil Mohommad and Gregor Schwerhoff, we study the international diffusion of carbon pricing policies. This diffusion is important, because if a country can increase the probability that other countries adopt a similar policy, this counteracts concerns about competitiveness and leakage. Furthermore, policy adoption and emission reductions in those other countries represent additional benefits of domestic policy adoption. We first use the observed adoption of carbon pricing policies between 1988 and 2020 to identify international policy diffusion in the data. Our empirical estimates then inform Monte Carlo simulations to quantify the indirect emission reductions from diffusion. We find that for

most countries indirect emission reductions due to diffusion can be larger than direct domestic emission reductions. This is especially important for relatively small countries, for which indirect emission reductions can be several times larger than their direct counterparts. Overall, our results however support a nuanced view on policy diffusion. While for individual countries the global benefits from policy diffusion appear substantial, diffusion alone seems to make a limited contribution to the achievement of a high geographical coverage of carbon pricing policies over the coming decades.

The results of Chapter 4 suggest that in situations with multiple market failures and political constraints, climate policy sequencing likely played an important role in facilitating the adoption of carbon pricing in many countries. This insight is generally consistent with prior quantitative work that focused exclusively on the energy sector (Meckling et al., 2015, 2017) and qualitative work on selected countries (Pahle et al., 2018). Our work extends the internal and external validity of these earlier findings with refined statistical analyses and more comprehensive data. Furthermore, we report additional insights, for example on trends in policy sequences over time and the association between policy sequences and the stringency of subsequent pricing policies. Our findings on international policy diffusion in Chapter 5 also reconcile somewhat contradictory results in previous studies on the diffusion of carbon pricing policies (Steinebach et al., 2021; Dolphin and Pollitt, 2021) by considering carbon taxes and emission trading systems as two types of the same policy. Our evidence on climate policy diffusion is generally consistent with prior qualitative (Thisted and Thisted, 2020) and quantitative work (Fankhauser et al., 2016). For the first time, the chapter also quantifies the benefits from policy diffusion.

The last two Chapters contribute to debates about opportunities, benefits, and constraints of carbon pricing policies. The insights from these two chapters are of high relevance for policy makers. Chapter 4 improves our understanding of the objectives, benefits, and constraints of climate policy making. This evidence allows countries to learn from the past, which is relevant especially for the many countries which are just at the start of pathways of rapid GHG emission reductions. The evidence on sequencing can also inform expectations about future policy adoption. The results in Chapter 5 counter possible concerns of policy makers about the international competitiveness of their economies and relatively small effectiveness of adopting carbon pricing at home. The results also contribute to debates about policy

instruments that put in place additional incentives for countries to follow those with relatively stringent climate policies, such as carbon border adjustments and international carbon price floors.

Chapter 1

Temperature variability and long-run economic development

This study estimates causal effects of temperature variability on economic activity. For identification I use a novel research design based on spatial first-differences. Economic activity is proxied by nightlights. I distinguish between day-to-day, seasonal, and interannual variability and find that the type of variability matters. The results suggest an economically large and statistically significant negative effect of day-to-day variability on economic activity at most temperature levels. Regarding seasonal variability, I find a smaller but also negative effect. The estimated effect of interannual variability is positive at low and negative at high temperatures. These effects are robust, they can be identified in urban and rural areas, and they cannot be explained with the spatial distribution of agriculture. The results draw attention to the effect of climate variability, which is projected to change but has so far been mostly overlooked in assessments of the impacts and costs of climate change.

1.1 Introduction

The debate about how climate and climate change influence economic development has a long history, but with few exceptions it has been all about annual mean climate. Fluctuations of temperature around its annual mean have thus been mostly neglected, although temperature variability is generally very common. For example, in many countries temperature frequently changes by several degrees Celsius from one day to the next. Furthermore, in many places temperature differs by more than 10 degrees Celsius between summer and winter. Differences between years are generally smaller, but annual mean temperature can still change by about 1-2 degrees Celsius from one year to the next, comparable in magnitude to global warming over the last 100 years. The current lack of evidence on the effects of this variability at the time scale of days, months, and years means that possible costs of larger variability are not included in most estimates of the costs of future climate change.

Advancing our understanding of the consequences of temperature variability has possibly been hindered by the challenge of identifying its causal effects. In recent years, the marginal effect of climate on economic activity has primarily been estimated with panel regression models using annual observations and unit and year fixed effects (Dell et al., 2012; Burke et al., 2015a). While this approach has generally been regarded as more credibly identifying causal effects than cross-sectional regressions, it cannot be used for variables that need to be measured over periods longer than a year and that change relatively slowly, such as seasonal and interannual temperature variability.

In this paper I estimate the causal effect of temperature variability at different time scales on long-run economic development. For identification I use a novel econometric framework based on differences between geographically proximate observations (Druckenmiller and Hsiang, 2018). This spatial first-differences research design allows me to identify the effect of slow-changing climatic variables under weaker assumptions than a regression on a cross-section of levels. The identification strategy can be interpreted as matching based on geographical proximity with a continuous treatment variable. I apply this method to a global dataset consisting of grid cells with a size of approximately 25 km x 25 km that contain information on economic activity, measured by satellites as the intensity of light at night, and temperature

and its variability from climate reanalysis, as well as several climatic and geographic controls.

Temperature variability at different time scales has different underlying physical mechanisms, predictability, and projected changes under future climate change. I therefore distinguish between temperature variability at the time scales of days, months, and years: *day-to-day*, *seasonal*, and *interannual* variability. To measure these variables, I calculate the intra-monthly standard deviation of daily temperature levels, the intra-annual range of monthly mean temperatures, and the inter-annual standard deviation of annual mean temperatures respectively. For seasonal variability, the range of monthly means is chosen instead of the inter-monthly standard deviation in order to avoid any overlap with the measure of day-to-day variability.

My empirical results suggest economically large and statistically significant negative effects of *day-to-day* and *seasonal* temperature variability on long-run economic development. On average, one sample standard deviation of day-to-day variability (1.44 degrees Celsius) and seasonal variability (14.71 degrees Celsius) reduces nightlights by 17 and 9 percent, respectively. If these effects are benchmarked with the estimated effect of annual mean temperature, they correspond to increases of annual mean temperature from 25 degree Celsius to approximately 30 and 28 degrees Celsius, respectively. Regarding *interannual* variability, I find a positive effect on economic activity below and a negative effect above an annual mean temperature of 20 degrees Celsius.

I discuss several theories about why and how temperature variability can affect economic activity, including non-linear effects of daily temperature levels (ex post effects) and greater uncertainty (ex ante effects). Because I use cross-sectional variation of long-term averages for identification, my estimates capture both ex post and ex ante effects. This contrasts most previous work that used time-series variation with annual frequency and thus captured primarily ex post effects (Pretis et al., 2018; Kotz et al., 2021b; Rudik et al., 2021). In line with the empirical effects that I find, previous microeconomic evidence suggests overall negative effects of temperature variability on economic activity. The fact that day-to-day variability is less predictable than seasonal temperature variability and hence introduces larger uncertainty could explain its more negative effect. Regarding interannual variability, I note that the pattern of estimated coefficients is consistent with

an asymmetry whereby the benefits/costs of colder-than-average years are smaller than the benefits/costs of warmer-than-average years, possibly due to heating being less costly and generally more widespread than cooling (Rode et al., 2021).

Previous research suggests that nightlights are a better proxy for GDP per capita in urban areas than in rural areas and that the spatial distribution of nightlights also reflects the local sectoral composition of the economy (Chen and Nordhaus, 2019; Gibson, 2020). I therefore also examine whether my results are primarily driven by urban areas and whether they can be explained by the spatial distribution of agricultural activity. I find that the estimated coefficients are indeed largest in urban areas, but I also find significant effects with the same sign for less densely populated regions, including the least densely populated areas within countries. Furthermore, the main results are unaffected by controlling for the spatial distribution of agricultural activity.

This is to my knowledge the first study to examine the long-run effect of temperature variability accounting for both *ex ante* and *ex post* effects and examining variability at multiple time scales. The results generally agree with previous studies finding negative effects of day-to-day variability on regional GDP (Kotz et al., 2021b), negative effects of daily and seasonal temperature variability on regional GDP in the US (Rudik et al., 2021), and negative effects of seasonal temperature variability on specific economic outcomes such as in agriculture (Mendelsohn et al., 2007a) and health (Hovdahl, 2020). Furthermore, I find positive marginal effects of annual mean temperature at relatively low temperature levels and negative effects at relatively high temperatures, consistent with previous findings of a negative quadratic relationship between annual mean temperature and economic growth (Burke et al., 2015b; Kalkuhl and Wenz, 2020).

The results also contribute to the debate about the the future costs of anthropogenic climate change. With few exceptions (Calel et al., 2020; Kikstra et al., 2021; Rudik et al., 2021), temperature variability is not accounted for in estimated costs of climate change. My results suggest that the costs of temperature variability should be included in assessments of future costs and deserve a closer look at their geographical distribution. Climate models project that seasonal variability will tend to decrease in cold and increase in relatively warm countries (Dwyer et al., 2012). These projections, together

with my results, suggest that accounting for seasonal variability decreases the costs of climate change in relatively cold countries and increases its costs in relatively warm (and currently poor) countries. The results on interannual temperature variability are generally less robust, but suggest that the benefits or costs of projected changes to interannual variability depend on current annual mean temperature levels. Together with projections of climate models (Bathiany et al., 2018), the results suggest that future changes to interannual temperature variability also tend to increase the costs of climate change in relatively warm (and currently poor) countries. Also observed trends of day-to-day variability over the last decades suggest additional economic costs of climate change in relatively warm countries (Kotz et al., 2021a).

The use of spatial first-differences reduces omitted variable biases (Druckemiller and Hsiang, 2018). However, because identification still relies on cross-sectional variation, I conduct a formal sensitivity analysis and several robustness tests. Specifically, I show that any omitted variable would need to be more strongly associated with both temperature variability and nightlights than any of the climatic control variables (annual mean temperature, precipitation, precipitation variability, relative humidity, solar radiation) or geographic control variables (elevation, terrain ruggedness, distance from coast, distance from inland water body) such that controlling for this confounder could make the estimates insignificant. I also compare estimates obtained from spatial differences in only either the North-South or the West-East direction and with different functional forms for my control variables. Furthermore, I do not find evidence that spatial spillovers or spatial dependence can explain my results.

Another concern regarding the identification strategy is reverse causality: climate can have an effect on local economic development, but local economic development can also influence the local climate. To address this concern I use an alternative source of nightlights data which allows me to examine changes of nightlights over time. This enables me to regress changes of nightlights over time on temperature variability observed over an earlier period. By doing so, any feedbacks from local economic development on local climate are excluded by design, but the main results can still be recovered.

The paper is structured as follows. In the next Section, I briefly explain why temperature variability might matter for economic activity (Section 1.2.1), introduce the three measures of climate variability and explain their geo-

graphical distribution (Section 1.2.2). I then describe the data in Section 1.3. In Section 1.4, I present the research design and identification strategy. All results are presented in Section 5.3: I first present the main results (Section 1.5.1) and then conduct several robustness tests (Section 1.5.2 and 1.5.3). In Section 1.6 I discuss several mechanisms that could explain the results. Finally, I discuss results in light of previous findings and point out implications for future research in Section 5.4.

1.2 Climate variability

1.2.1 How temperature variability can affect economic activity

Annual mean temperature affects economic production in both developing and developed countries (Dell et al., 2012; Burke et al., 2015b; Kalkuhl and Wenz, 2020). This effect appears to be non-linear, with possibly positive marginal effects at low and increasingly negative marginal effects at high temperature levels (Burke et al., 2015b; Kalkuhl and Wenz, 2020). These empirical results have been explained with alternative mechanisms, including effects of daily temperature levels on human cognitive processes (Almås et al., 2019) and effects of daily temperature levels on production unrelated to labour, such as crop failures in agriculture.

Temperature variability can add to the costs of annual mean temperature if the relationship between daily temperature levels and economic production is non-linear. In that case, the net costs of variability in any given location depend on the relative frequencies of different levels. This effect of temperature variability is explored for example by Rudik et al. (2021) and by Calel et al. (2020) and Kikstra et al. (2021) using integrated assessment models with non-linear damage functions, which yield an overall negative effect of larger variability. Such effects are also a possible explanation of the negative effect of diurnal temperature ranges reported in prior work (Mitton, 2016). Alternatively, temperature variability introduces costs if there are heterogeneous locally optimal temperature levels. Such locally optimal temperature levels have been documented for example for the choice of crops in South America (Seo and Mendelsohn, 2008) and can be observed for human physiology (Hanna and Tait, 2015). In both areas, detrimental effects of temperature variability have been reported on respectively crop yields (Wheeler et al., 2000) and temperature-related mortality (Hovdahl, 2020).

These possible costs of temperature variability are associated with realised temperature levels (ex post effects of variability). In addition, temperature variability can affect economic activity through expectations (ex ante effects of variability). To the extent that larger variability implies greater uncertainty about future temperature levels, and assuming there are effects of realised temperature levels on production, larger temperature variability means larger uncertainty of income and returns to investments. This uncertainty can have a negative effect on economic activity by discouraging investment. Such effects of variability have been documented e.g. for exchange rate fluctuations (Aghion et al., 2009) and volatility of government spending (Ramey and Ramey, 1994). Regarding climate, the effect of rainfall variability on output volatility was examined e.g. by Malik and Temple (2009). Economic agents can be expected to respond to greater climate uncertainty through risk diversification (Bellemare et al., 2013; Bezabih and Di Falco, 2012; Ashraf and Michalopoulos, 2015; Colmer, 2021; Buggle and Durante, 2021), but such diversification might not always be possible, be limited in its effectiveness, and come at a cost.

Overall, it can therefore be expected that *ceteris paribus* temperature variability has a negative effect on economic activity, either due to more frequently observed detrimental temperature levels, more frequent deviations from locally optimal temperature levels, or larger uncertainty. Only in situations in which greater variability means more frequent beneficial temperature levels and this positive ex post effect is larger than possible negative ex ante effects, will variability have a net positive effect. Because of non-linear effect of temperature levels, the effects of variability are expected to depend on the average temperature level. Furthermore, the effect of variability is expected to depend on the time scale of variability because more frequent fluctuations can be learned about more easily while they allow for less time for adjustments, as well as on its predictability.

1.2.2 Temperature variability: day-to-day, seasonal, and interannual

Due to human activity, mainly the burning of fossil fuels and the associated emission of greenhouse gases into the atmosphere, temperatures have been increasing since at least the second half of the 20th century. This slow trend has been overlaid by fluctuations on a range of time scales. In this paper I examine variability at the time scale of days, months, and years, that is day-

to-day, seasonal, and interannual variability, respectively. In this Section, I first describe how I isolate variability at different time scales and then explain the global distribution of temperature variability at different time scales with the underlying physical processes of weather and climate.


The construction of the three measures of temperature variability is illustrated in Figure 1.1. All three measures are calculated from timeseries of daily temperature levels in several steps, as explained in the following.

Day-to-day temperature variability is defined as fluctuations of temperature within the same month. These fluctuations are generally observed on top of gradual trends due to seasonality. To isolate the measure of day-to-day temperature variability from such trends, which are considered separately in the measure of seasonal temperature variability that is described further below, I first subtract a smooth average annual cycle from the daily timeseries (Figure 1.1b,c). To estimate a smooth average annual cycle, I follow (Moberg et al., 2000) and fit a smooth curve into the multi-year average daily mean temperatures (daily mean temperatures averaged over my reference period 1985-2014). I use a Hodrick-Prescott filter with a smoothing parameter $\lambda = 10,000$ which subtracts trends extending over multiple years but not fluctuations from one year to the next. I subtract the average cycle because in countries with large annual cycles (large differences of temperature between summer and winter), this cycle means relatively steep trends in spring and fall, which lead to variation of temperature within months. Its subtraction prior to the calculation of intra-monthly standard deviations of daily temperature levels thus ensures that day-to-day variability is isolated from any influence of seasonal variability (Moberg et al., 2000).







Seasonal temperature variability is defined as the annually recurring differences in temperature between the relatively warm and the relatively cold seasons of a year. In this paper, I quantify seasonal temperature variability using the intra-annual range of monthly mean temperatures (Figure 1.1d,e). I choose monthly means instead of daily means to reduce the influence of potentially rare and extreme days. An alternative measure is the standard deviation of monthly mean temperatures. The range of monthly means is chosen to further reduce any overlap with the measure of day-to-day variability and to allow for better comparison of this measure across the globe and thus easier interpretation of the results, because some locations close to the Equator exhibit complex patterns of seasonality with several peaks

Figure 1.1: Calculation of my three measures of temperature variability: day-to-day, seasonal, and interannual variability.

Daily mean levels



Day-to-day variability: Seasonal variability: Interannual variability:



Notes: The top figure shows daily temperature levels for London from 1985 to 2014 using ERA5 reanalysis (see Section 1.3). The first column (Figures b, c) show two steps to calculate the day-to-day variability: after subtraction of a smooth average annual cycle, I calculate the intra-monthly standard deviation of daily temperature anomalies (Figure b), which I then average 1985-2014 (Figure c). The second column (Figures d, e) show two steps to calculate seasonal variability: I first calculate the intra-annual range of monthly mean temperatures (Figure d), which I then average 1985-2014 (Figure e). The third column (Figures f, g) show two steps to calculate inter-annual variability: I first subtract a smooth trend from annual mean temperatures (Figure f) and then calculate the inter-annual standard deviation of annual temperature anomalies 1985-2014 (Figure g). Figures b, c, d, e do not show the full time period 1985-2014 for readability.

during the year which are all included in the standard deviation of monthly means, but not in their range (see also Chapter 2). In a robustness test, I find that the main results are however very similar if I measure seasonal variability using the inter-monthly standard deviation of monthly mean temperatures, with a slightly higher significance of the effect of seasonal variability (Appendix A.4). For both day-to-day and seasonal temperature variability, I average the monthly values of intra-monthly standard deviations and the annual values of intra-annual ranges respectively over the period 1985-2014, which is the 30 years period preceding the year of the nightlights data.

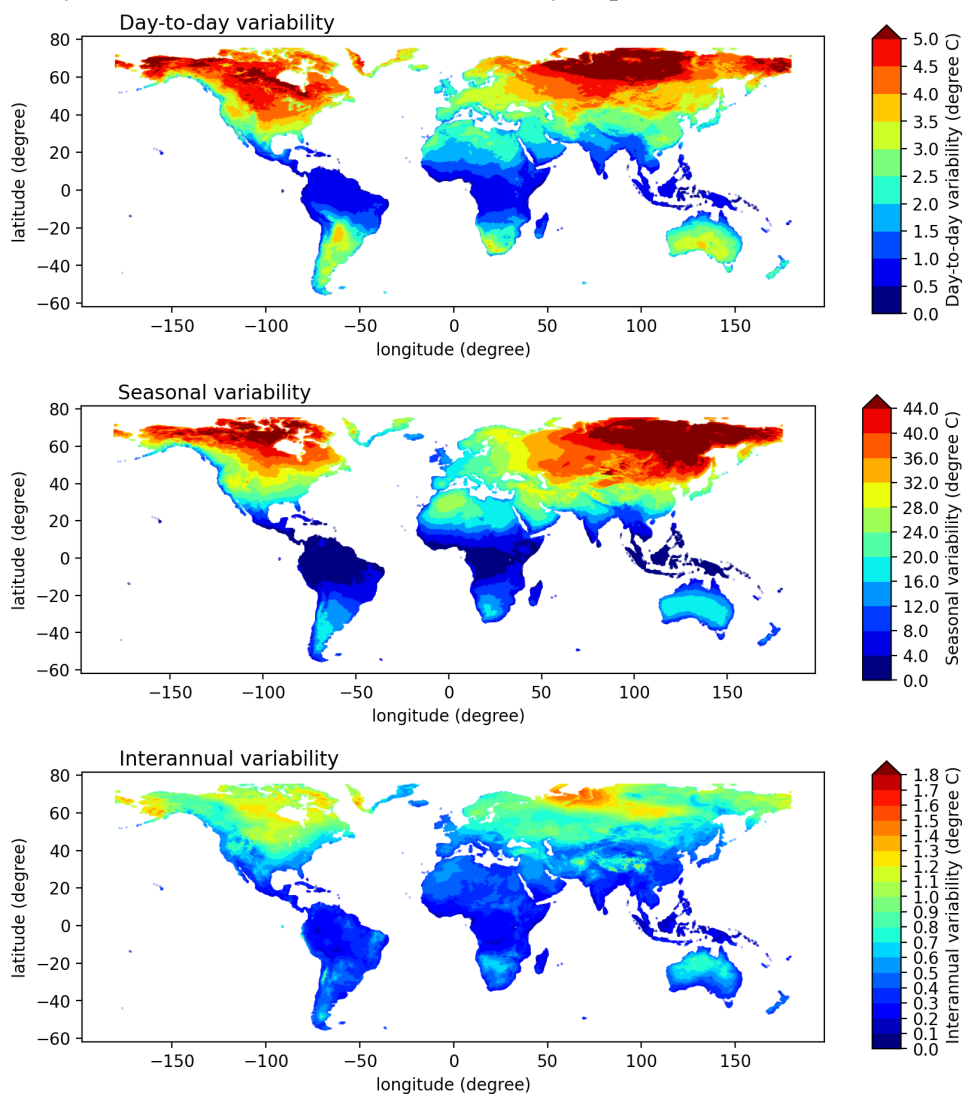
Interannual variability is calculated as the between-year standard deviation of annual mean temperatures over the same 30 year period (Figure 1.1c,f). Before I calculate the standard deviation, I remove a slow trend in order to isolate interannual variability from any warming (or cooling) trends due to anthropogenic climate change. The omission of trends from anthropogenic warming is motivated by the focus of this paper on variability (fluctuations) and not on gradual changes (trends). In some contexts, these two concepts are not always separated and the term climate variability is more loosely used to refer also to gradual trends. In contrast, the terminology of this paper distinguishes between variability and trends as typically done in the climate science community based on whether or not observed changes show a repeatedly recurring pattern.

The global maps of temperature variability reflect the influence of astronomy, geography, and climate dynamics (Figure 1.2). While the maps of day-to-day, seasonal, and interannual variability resemble each other and suggest positive correlations between the variables, I explain in the following how the relative importance of several physical processes differs. I look into the econometric implications of the high degree of spatial correlation in Section 1.4.

Day-to-day variability is generally larger at higher latitudes. This is primarily due to the influence of high and low pressure systems travelling eastwards at these latitudes, which cause frequent changes between local advection of cold (polar) and warm (tropical) air. Furthermore, day-to-day variability is larger further away from the coastline as land responds faster than water to changes in air temperature between days.

Seasonal variability of temperature is generally larger at high latitudes than at low latitudes due to the tilt of Earth's axis (Figure 1.2b). Furthermore, be-

Figure 1.2: Geographical distribution of temperature and its variability: day-to-day, seasonal, and interannual variability (top to bottom).



Source: ERA-5 reanalysis (see Section 1.3.2).

cause land responds faster to changes in solar radiation than oceans and the land areas are larger in the Northern hemisphere than in the Southern hemisphere, seasonal variability is generally larger in the Northern hemisphere and smaller closer to the coast and large inland water bodies (Legates and Willmott, 1990). Because at mid and high latitudes the wind tends to flow from West to East and the temperature of a parcel of air is influenced by the temperature of the surface over which it has been transported (McKinnon et al., 2013; Stine and Huybers, 2012), seasonal variability also tends to be larger on the Eastern parts of large continents (America, Eurasia).

Interannual temperature variability is partly driven by external astronomical influences, such as solar cycles of about 11 years, but primarily due to internal climate variability (Mann and Park, 1994). Internal climate variability results from oscillations in the climate system, which are often related to interactions between different components of the climate system, such as the atmosphere and the ocean. Examples are the El-Nino Southern Oscillation (ENSO) and the North-Atlantic Oscillation (NAO) (IPCC, 2013). Interannual variability is generally larger further away from the coasts because the oceans have a larger heat storage capacity and thus a larger year-to-year inertia than land (Figure 1.2c). Interannual variability is largest in high Northern latitudes and at high altitude due to its amplification by the snow/ice-albedo feedback.

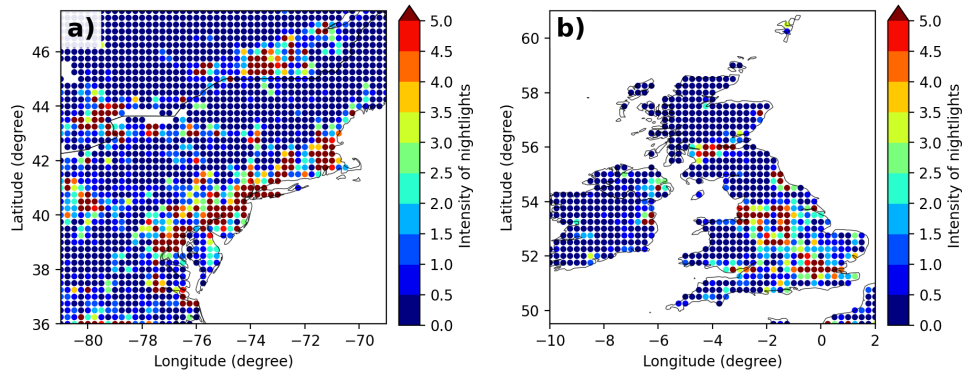
1.3 Data and descriptive statistics

1.3.1 Economic variables

I proxy economic activity by the intensity of lights at night (Chen and Nordhaus, 2011; Henderson et al., 2012; Nordhaus and Chen, 2015). Nightlights are measured by satellites and come with a resolution that outcompetes census-based measures of economic activity. This granularity of the data is particularly important in my research design, as identification rests on the comparability of neighbouring observations. Another advantage of using nightlights instead of population or GDP is that nightlights are consistently measured with the same quality and the same resolution worldwide. I take data on the intensity of lights at night from the satellites of the Visible Infrared Imaging Radiometer Suite (VIIRS) (Elvidge et al., 2017). The VIIRS is a relatively new satellite product which can be regarded as a successor of the popular DMSP data. As compared to the DMSP data, the VIIRS data suffer less from blur-

ring, a lack of sensor calibration, and a limited range of sensitivity (Chen and Nordhaus, 2019; Gibson et al., 2021).

Figure 1.3: Geographical distribution of VIIRS nightlights: a) North-Eastern coast of the USA and b) the British Isles.



Source: VIIRS nightlights.

I use annual average radiance values which have undergone some post-processing to remove the effect of clouds and to filter out fires and other ephemeral lights. I use nightlights for the year 2015 with a resolution of 15 arc-seconds, which I aggregate to a resolution of 0.25 degrees. The year 2015 is the earliest year for which VIIRS nightlights were available at the time of the analysis. Because interest is in the effect of climate, typically defined over 30 years, on the level of economic development, climatic averages of 1985-2014 are combined with nightlights for 2015. Robustness tests with an average of nightlights 2015-2019 and with a very recent version of the VIIRS nightlights data yield the same results (Appendix A.7).

As most economic activity occurs on land rather than on water, the average radiance tends to be larger in grid cells with a higher share of land area. This could bias my results if the share of land area correlates with my climatic variables. I address this concern by multiplying the average radiance of a grid cell by the total area of the grid cell and dividing it by its total land area. All grid cells without land area are dropped from the data. Because the distribution of normalised nightlights is highly skewed, I log-transform the data.

At a global scale, the spatial distribution of VIIRS nightlights primarily shows the location of large metropolitan areas. At the regional scale, the spatial distribution of nightlights also shows variation outside metropolitan areas (Figure 1.3).

Although I prefer the VIIRS data to the DMSP data due to technological improvements (Gibson et al., 2021), I also download DMSP data for an additional robustness check for the years 1992 and 2012. The data are processed with the same steps as the VIIRS data. Furthermore, I use population data from the Gridded Population of the World (GPW) dataset version 4.0 (Center For International Earth Science Information Network-CIESIN-Columbia University, 2018). I choose this dataset as it is based on official censuses only and thus independent of my nightlights data. I also use data on the global distribution of cropland and pasture lands (Ramankutty et al., 2008) provided by NASA (Ramankutty et al., 2010), which I aggregate from its native resolution to a resolution of 0.25 degrees.

1.3.2 Climate variables

I use climate data from the global reanalysis ERA-5. Reanalysis data are produced by feeding an adjusted weather forecast model with the full global record of observational data, including weather station records and satellite data (Parker, 2016). ERA-5 belongs to the newest generation of reanalysis datasets and is provided with a resolution of 0.25 degrees. I choose reanalysis data instead of station-based weather data because of the physical consistency of reanalysis data. Furthermore, meteorological measurements are globally unevenly distributed and I expect that processing with a dynamic model evens out some of the heterogeneity in data quality.

The reanalysis data also has the advantage that they include climate variables in addition to temperature and precipitation. I include several additional variables in my model to reduce potential biases due to omitted climate variables. These biases could lead to a misattribution of empirically observed causal effects, but are not necessarily problematic as long as the physical relationships between the variables can be expected to be constant over time or if the estimated relationships are not used for future projections. To avoid misattribution, I also include relative humidity and solar radiation in my regressions. I use daily mean values of all climate variables for the period 1985 to 2014, which is the 30 years period prior to the VIIRS nightlights data, and for the period 1982-1991 (the period before the DMSP nightlights data). In another robustness check, I find that the results are qualitatively very similar if I use an earlier period for the climate data (1955-1984) (Appendix A.6).

1.3.3 Geographical covariates

The use of the spatial-first differences research design reduces omitted variable biases from all variables whose spatial gradients do not systematically correlate with the gradients of temperature variability at the spatial scale of my observations (about 25 km). For example, I expect that any differences in institutions between countries cannot bias my results. Furthermore, in order to reduce biases from specific variables, I also include several geographic controls.

Table 1.1: Descriptive statistics. Number of observations: 233,362.

Variable	Unit	Mean	Std.	Min.	Max.
log Nightlight intensity VIIRS		0.11	0.35	0.00	7.19
log Nightlight intensity DMSP		0.46	0.87	0.00	13.72
Elevation	km	0.62	0.80	-0.24	6.31
Terrain ruggedness	-	102.42	146.06	0.00	1355.07
Distance from nearest coast	1000 km	0.55	0.52	0.00	2.50
Distance from nearest lake/river	1000 km	0.28	0.50	0.00	6.33
Annual mean temperature	deg C	25.29	13.39	-5.94	48.82
Day-to-day var. of temperature	deg C	2.96	1.44	0.31	6.07
Seasonal var. of temperature	deg C	24.37	14.71	0.74	65.08
Interannual var. of temperature	deg C	0.64	0.30	0.10	1.56
Annual total precipitation	mm	69.50	67.73	0.05	2499.60
Seasonal var. of precipitation	mm	49.79	42.37	0.16	1037.65
Interannual var. of precipitation	mm	0.01	0.01	0.00	0.38
Annual mean rel. humidity	%	89.34	7.03	64.94	98.62
Annual mean solar radiation	W m ⁻²	179.87	57.63	76.34	309.41
Share of cropland	%	10.66	19.64	0.00	100.00
Share of pasture land	%	18.47	27.06	0.00	100.00

Notes: Climate variables computed over period 1985-2014. VIIRS nightlights annual composite for 2015. DMSP nightlights annual composite 2012.

Elevation increases transport costs and is hence a major geographic factor for economic development. Furthermore, elevation is one of them main determinants of local climate. I take data on elevation from the Global Land One-kilometer Base Elevation (GLOBE) dataset in version 1 provided by the National Oceanic and Atmospheric Administration (NOAA) (Hastings et al., 1999). The dataset has global coverage with a horizontal resolution of 0.0083°. I download the data as tiles, merge them, and then aggregate it to 0.25° by averaging.

Previous research has revealed a statistically significant association between

terrain ruggedness and economic development in Africa (Nunn and Puga, 2012). Furthermore, terrain ruggedness influences the horizontal and vertical exchange of air, which in turn affect the local climate at the surface. I therefore also include terrain ruggedness as a control variable. Data on terrain ruggedness are taken from a global dataset with a resolution of 1 km (Shaver et al., 2018), which I aggregate to 0.25 degrees.

Economic activity tends to be clustered at the coasts in many countries (Henderson et al., 2018). Furthermore, seasonal variability of temperature tends to be smaller closer to the coast (Section 1.3.2). I therefore also include distances from the nearest coast and distance from inland water bodies as control variables. Distances from the nearest coast are taken from a dataset provided by the NASA. The dataset covers the whole globe with a uniform horizontal resolution of 0.04° . I also use data on distance from inland water bodies (GloboLakes dataset provided by the CEDA archive) (Carrea et al., 2015). The data were created from ENVISAT satellite images. The data are provided with a 300 m resolution. I aggregate both datasets to a resolution of 0.25 degrees using mean values.

1.3.4 Descriptive statistics

The final dataset consists of 233,362 complete observations (Table 1.1). Each observation corresponds to a grid cell of 0.25 degrees width in both latitudinal and longitudinal direction, which corresponds to about 28 km at the equator, about 23 km at 45 degrees latitude, and about 20 km at 60 degrees latitude. The final dataset excludes grid cells that are not located on land and grid cells on land that are covered by water or ice. Furthermore, due to the spatial coverage of the nightlights data, the dataset is bounded by the latitudes 75 N and 60 S. For the main analysis, nightlights in the year 2015 are combined with time-invariant geographical covariates and climate variables averaged over the period 1985-2014. The exclusion of the year 2015 in the climate data and the averaging over multiple years reduces the influence of (contemporaneous) extreme events, and a 30-years period corresponds to the conventional definition of climate.

1.4 Econometric strategy

I estimate the model using a spatial first-differences research design. The spatial first-differences (SFD) estimator has recently been proposed as an

econometric estimation method that can reduce omitted variable bias for cross-sectional data (Druckenmiller and Hsiang, 2018). This is especially useful for the analysis of climate variability, which is a characteristic of climate and not weather and for which identification therefore arguably requires the “Ricardian” approach of comparing locations with different climate (Hsiang, 2016; Auffhammer, 2018). The SFD estimator uses only variation between spatially adjacent units of observations. Identification hence relies on the local conditional independence assumption

$$E [Y_i|(D_{i-1}, X_{i-1})] = E [Y_{i-1}|(D_{i-1}, X_{i-1})] \forall i \quad (1.1)$$

whereby observations are indexed with i along a spatial dimension, Y is the outcome variable (log nightlights in the main model of this paper), D is the treatment variable (temperature variability), and X are control variables (climatic and geographic covariates). Equation 1.1 means that the SFD estimator requires that, conditional on all covariates, *spatially adjacent units of observation* with the same treatment have the same expected outcome. This is a weaker assumption than the assumption underlying a conventional cross-sectional regression of levels, for which conditional on all covariates *all units of observation* with the same treatment need to have the same expected outcome.

The OLS estimator of the SFD design can then be written as

$$\hat{\beta}_{SFD} = (\Delta X' \Delta X)^{-1} (\Delta X' \Delta Y) \quad (1.2)$$

where Δ refers to the first difference between adjacent units of observations. If the local conditional independence assumption (Equation 1.1) is satisfied, it implies that

$$E [\Delta X' \Delta C] = 0. \quad (1.3)$$

for any potentially omitted variable C . The SFD estimator thus eliminates biases due to omitted variables if the spatial differences of the treatment variable and the spatial differences of a potential confounder are not systematically correlated (Druckenmiller and Hsiang, 2018). Another strength of the SFD research design is a unique robustness test. This robustness test exploits

the fact that the estimator can be used with spatial differences in any direction, including North to South (NS) and West to East (WE). If the identifying assumption of SFD is satisfied, the regression coefficients obtained from differences in different directions should be statistically the same (Druckemiller and Hsiang, 2018). I conduct this robustness test in Section 5.3.

The SFD framework can also be compared with a spatial regression-discontinuity (RD) research design. In contrast to an estimation with RD, the SFD estimator does not require a discontinuity of the treatment variable. Instead, the marginal effect is recovered from all changes in the outcome (nightlights) and treatment variable (temperature variability) along the North-South or East-West direction. This reduces the risk that estimates primarily reflect correlations of gradients in temperature variability and nightlights in places with sharp gradients of temperature variability because of extraordinary geographical features, where the identifying assumption for a model with polynomial terms of all control variables might be violated. This concern is addressed with a robustness check in which the top 5% and bottom 5% of observations in terms of temperature variability are excluded and which yields very similar results as the main estimation (Appendix A.5).

Grid cells at a distance of about 20-30 km can generally be expected to have relatively similar climates. This might raise the question whether differences in temperature variability as observed over 30 years are due to systematic differences in climate or instead due to singular events in this time period that affected one location more than the other. If the latter concern was correct, the SFD estimator would identify a short-term effect of weather rather than a long-term effect of climate (Hsiang, 2016). To address this concern, I vary the time periods over which temperature variability (1955-1984 and 1985-2014) and nightlights (2015 and 2015-2019) are observed and find similar results (Appendices A.6 and A.7). This suggests that my estimates can be interpreted as the marginal effects of climate rather than weather.

In the presence of spatial spillovers, the estimates obtained from spatial first-differences will be biased due to a violation of the SUTVA assumption (Druckemiller and Hsiang, 2018). Such spatial spillovers can result from flows of people, capital, or goods between adjacent locations. For example, economic activity in a specific location might be negatively affected by outward-migration to neighbouring locations. If this migration was partly influenced by local temperature variability, the empirical estimates would be

biased because the treatment (i.e. temperature variability) of one unit of observation (i.e. location) would affect the outcome (i.e. economic activity) of another unit of observation (i.e. location). The presence of spatial spillovers can be examined with a model that includes spatial lags of the treatment variable (Druckenmiller and Hsiang, 2018). In a robustness test, I therefore also include spatial lags of annual mean temperature and the different measures of temperature variability before taking spatial first differences. The results suggest that spatial spillovers are not important in this setting (Appendix Table A.6). Furthermore, estimation with spatial first-differences does generally not account for possible spatial dependence of the outcome variable. To examine whether there is evidence that spatial dependence influences the results, I estimate my main model but pair grid cells before taking spatial differences not with their immediate neighbour, but with their neighbour's neighbour. In other words, I always skip one observation in space when taking spatial differences. The results are similar to my main results, suggesting that also spatial dependence is no major issue in this setting (Appendix Table A.7).

To illustrate the strengths of the SFD framework, I estimate a simple model in which I explain variation of nightlights by day-to-day temperature variability (and annual mean temperature). The exercise focuses on day-to-day variability as I can use recent estimates of its effect on regional GDP per capita using variation across time for identification as a benchmark (Kotz et al., 2021b). Using the sign of this previously reported effect as a reference, the results suggest that the SFD estimator reduces omitted variable biases as compared to a regression with levels. While this first evidence is reassuring, possible omitted variable biases are in more depth discussed in the next Section. Furthermore, what has not been reported before but is important for this paper, the SFD estimator also reduces multicollinearity in the model (Appendix A.2). The reason for the latter insight is that annual mean temperature and temperature variability at different time scales are influenced by latitude and thus strongly correlated with each other (and with other climate variables such as solar radiation). These correlations are substantially reduced when one uses spatial first differences instead of levels as the influence of latitude is smaller (relative to other variables) if one compares only neighbouring observations (Appendix A.2). Taken together, the SFD estimator thus seems to be a promising tool for navigating concerns of omitted variable biases on the one hand and multicollinearity on the other

hand, which have been identified as key challenges of empirical work on the effect of weather and climate on socioeconomic outcomes (Auffhammer et al., 2013).

1.5 Results

1.5.1 Main results

Previous authors have found non-linear relationships between annual mean temperature and GDP per capita (Burke et al., 2015a). Similarly, non-linear associations have been reported between daily temperature levels and many different socio-economic outcomes including labour productivity and health (Carleton and Hsiang, 2016). This suggests that the effect of temperature variability on long-run economic outcomes might be moderated by the effect of annual mean temperature (Section 1.2.1). I explore this hypothesis by estimating a flexible model in which I interact temperature variability with dummy variables for bins of annual mean temperature that are 4 degrees Celsius wide:

$$\Delta \log \left(1 + \frac{\text{nightlight}_i}{\text{landarea}_i} \right) = \sum_k \delta_i^k \left(1 + \beta_1^k \Delta \sigma_i^d + \beta_2^k \Delta \sigma_i^m + \beta_3^k \Delta \sigma_i^y + \beta_4^k \Delta \bar{T}_i \right) + \lambda \Delta \tilde{\mathbf{C}}_i + \gamma \Delta \mathbf{G}_i + \epsilon_i \quad (1.4)$$

where observations are indexed by i and Δ is the spatial first difference operator. Units of observations are grid cells with 0.25 degrees width, corresponding to about 25 km at the Equator. The dependent variable is the annual mean luminosity of nightlights divided by the landarea of a grid cell. Day-to-day, seasonal variability and interannual variability of temperature are denoted by σ^d , σ^m , and σ^y , respectively. δ_i^k is an indicator variable for temperature bin k that takes on values 0 and 1, \bar{T} is annual mean temperature, and $\tilde{\mathbf{C}}$ is a matrix of climate controls including terms for annual total precipitation, relative humidity, solar radiation, and the same three measures of variability of precipitation. The matrix of geographic controls \mathbf{G} includes grid cell averages of the distance to the nearest coast, the distance to the nearest water body, elevation, and terrain ruggedness. I estimate models with quadratic polynomials for all control variables. Standard errors are clustered at the country level to account for heteroskedasticity and spatial autocorrelation.

I also estimate models with standard errors clustered at the level of subnational administrative units, which yields smaller standard errors. This suggests that unexplained factors that determine the intensity of lights at night tend to be correlated within countries (e.g. electrification).

The results suggest that day-to-day variability and seasonal variability tend to reduce economic activity at most levels of annual mean temperature (Appendix A.3). Regarding annual mean temperature, I find a positive marginal effect at annual mean temperatures between 4-16 degrees Celsius and a negative effect at all other temperature levels. This pattern of marginal effects is consistent with results of previous findings indicating a negative quadratic relationship between annual mean temperature and economic growth (Burke et al., 2015b), except the negative marginal effect at very low temperature levels.

Furthermore, the analysis with the binned-model yields negative coefficients of day-to-day and seasonal variability across most levels of annual mean temperature, but an coefficients of interannual variability whose sign is positive at low and negative at high levels of annual mean temperature. For parsimony I thus also estimate a model that is as simple as possible but still able to produce these main findings. The model includes linear terms for day-to-day seasonal variability and an interaction between a linear term for interannual variability and a dummy variable for annual mean temperature:

$$\begin{aligned} \Delta \log \left(1 + \frac{\text{nightlight}_i}{\text{landarea}_i} \right) &= \beta_1 \Delta \sigma_i^d + \beta_2 \Delta \sigma_i^m \\ &+ \delta(\bar{T}_i < 20) (1 + \beta_3^A \Delta \sigma_i^y) \\ &+ \delta(\bar{T}_i \geq 20) (1 + \beta_3^B \Delta \sigma_i^y) \\ &+ \lambda \Delta C_i + \gamma \Delta G_i + \epsilon_i \end{aligned} \quad (1.5)$$

where $\delta(\bar{T}_i \geq 20)$ and $\delta(\bar{T}_i < 20)$ are indicator variables that take on the value 1 if annual mean temperature \bar{T}_i is larger or equal/smaller than 20 degrees Celsius and 0 otherwise. As expected from the patterns in Figure A.2, the estimation yields negative coefficients of day-to-day and seasonal temperature variability (Column 1 in Table 5.2). For interannual variability, I find a positive coefficient below 20 degree Celsius and a negative coefficient above this temperature level.

Table 1.2: Results of a linear model estimated with SFD.

Dependent variable:	<i>log Nightlight density</i>		
Spatial first differences:	Pooled	WE	NS
Column:	1	2	3
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.69073*** (0.15170)	-0.41483*** (0.11675)
Seasonal variab. of T	-0.28016 (0.17127)	-0.14383 (0.17498)	-0.32325* (0.17247)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.17134*** (0.05661)	0.16404*** (0.04454)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.23005** (0.09922)	-0.25865** (0.11796)
<i>Effect of increase by 1 deg. C on log nightlights</i>			
Day-to-day variab. of T	-0.11631	-0.15926	-0.09564
Seasonal variab. of T	-0.00632	-0.00325	-0.00729
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.19185	0.18926	0.18119
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.27857	-0.25410	-0.28569
Climate controls (linear)	x	x	x
Climate controls (quadratic)	x	x	x
Geographic controls (linear)	x	x	x
Geographic controls (quadratic)	x	x	x
R2	0.0249	0.0250	0.0254
df	448877	224426	224425

Notes: The table shows the results of a linear model (Equation 1.5) estimated with spatial first-differences. Standard errors in parentheses. WE = West-East, NS = North-South. Pooled = pooling differences in WE and NS.

The magnitude of the estimated coefficients is substantial. An increase of day-to-day and seasonal variability by one standard deviation (1.44 and 14.71 degrees Celsius, respectively) is associated with a reduction of nightlights by about 17 percent and 9 percent respectively (Column 1 in Table 5.2). For interannual variability, the magnitude is about 6 percent below and 8 percent above 20 degree Celsius of annual mean temperature.

As a first robustness test, I compare the results obtained by estimating the model with spatial first-differences in the West-East (WE) direction (Column 2) with the results obtained from North-South (NS) differences (Column 3), as well as with differences in these two directions pooled (Column 1). If my estimated coefficients of temperature variability could be explained with an omitted variable whose association with temperature variability or nightlights were not similar in both directions, I would expect to obtain different

estimates. However, for both directions I find that all coefficients have the same sign, similar magnitude, and similar significance.

As another robustness test I change the model specification for all control variables. I test models with linear terms, linear and quadratic terms, and linear terms interacted with bins of annual mean temperature. The results are presented in Table A.11 in Appendix A.9. Overall, the estimated coefficients of temperature variability are similar across specifications (Columns 4, 5, 6). I also do not find evidence for a violation of the SUTVA assumption due to spatial spillovers (Appendix Table A.6) and no evidence that the results are significantly affected by spatial dependence (Appendix Table A.7).

In the next paragraphs I conduct two additional robustness tests using this model, before turning to a discussion of mechanisms in Section 1.6.

1.5.2 Sensitivity analysis

In order to quantify the robustness of my key results to omitted variable bias I also conduct a formal sensitivity analysis. Specifically, I calculate how much of the residual variation of temperature variability and the residual variation of nightlights of the model in Equation 1.5 an omitted variable would need to explain such that including this additional variable could make the estimated coefficients of temperature variability insignificant or even reduce them to zero. Following Cinelli and Hazlett (2020), I quantify the robustness using partial R^2 , which means that my results on robustness are not specific to any assumed functional form of the omitted variable in the model, but instead provide an upper bound on the sensitivity to any set of omitted variables including non-linear terms and interactions. The full results are presented in Table A.10 in Appendix A.8.

To make the estimated coefficients insignificant at $\alpha = 0.05$, I find that an omitted variable would need to explain at least 2.71, 0.74, 0.92, and 0.45 percent of the residual variation of nightlights and of the residual variation of day-to-day, seasonal, and interannual variability below and above an annual mean temperature of 20 degrees Celsius, respectively. To reduce the estimates to zero, these robustness values are respectively 2.99, 1.03, 1.28, and 0.94 percent. These values can be put into perspective by comparing them with the partial R^2 of variables included in the model. This benchmarking shows that none of the included climatic and geographic control variables are comparably strongly associated with both nightlights and temperature

variability. While a few variables are sufficiently strongly associated with temperature variability, none of these explains enough of the residual variation of nightlights. This means that no potentially omitted variable that is similarly strongly associated with temperature variability and nightlights as any of the included control variables could make my results insignificant if it were additionally included in the model.

1.5.3 Reverse causality

It is well established that air temperature tends to be higher in the center of a city than in its surroundings due to what is referred to as the urban heat island. If the spatial distribution of economic activity affected also the variability of temperature, for example through human land use that changes the heat capacity of the surface, the statistically significant association between temperature variability and nightlights could generally also be explained with reverse causality. I address this concern by regressing changes of nightlights over time (between 1992 and 2012) on temperature variability measured over an earlier period (1982 to 1991). This means that any effects of nightlights on temperature variability are intentionally excluded by the design of the regression. For this analysis I use the older DMSP nightlights data, as the VIIRS data are only available since 2015.

I first regress DMSP nightlights in 2012 on climate over the period 1982-2011 (Table 1.3, Column 1), similar to my main regression with VIIRS data. I find the same key results as for the VIIRS data: a (significantly) negative coefficient of day-to-day and seasonal variability and a negative and positive coefficient of interannual variability of temperature respectively below and above an annual mean temperature of 20 degrees Celsius. The size of the coefficients is more difficult to compare as the two datasets measure the intensity of nightlights with different technological devices and on different scales.

To address the concern of reverse causality, I regress the growth of nightlights between 1992 and 2012 on the mean climate of the period 1982 to 1991 (Column 2). Reassuringly, I find the same sign and significance of the coefficients as in Column 1. I take this as evidence that my results are robust to possible confounding effects due to reverse causality. This result holds true also if I include nightlight density in 1992, which is likely associated with temperature variability from 1982-1991 and, as the results confirm (Column

3), also with subsequent changes in nightlights.

Table 1.3: Results of regressions addressing concerns of reverse causality using the DMSP nightlights data.

Dependent variable:	$\log NL (2012)$		$\Delta \log NL (1992 \text{ vs. } 2012)$
	1982 - 2011	1982 - 1991	1982 - 1991
Time period for climate variables:	1982 - 2011	1982 - 1991	1982 - 1991
Column:	1	2	3
Day-to-day variab. of T	-0.31131*** (0.08156)	-0.04719* (0.02665)	-0.07712*** (0.02819)
Seasonal variab. of T	-0.19933 (0.12123)	-0.06371** (0.02934)	-0.07539** (0.03635)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.12011*** (0.02794)	0.02603* (0.01419)	0.02898* (0.01541)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.11410 (0.07603)	-0.01148 (0.02808)	-0.01724 (0.02979)
Nightlights in 1992			-0.12585*** (0.01775)
Climate controls (linear)	x	x	x
Climate controls (quadratic)	x	x	x
Geographic controls (linear)	x	x	x
Geographic controls (quadratic)	x	x	x
R2	0.0473	0.0102	0.0374
df	448877	448877	448876

Notes: The table shows the results of a model with linear terms for day-to-day and seasonal variability and an interaction term for interannual variability (Equation 1.5) estimated with spatial first-differences. Standard errors in parentheses. NL = nightlight density based on DMSP data.

1.6 Mechanisms

1.6.1 Urban areas

It is well known that nightlights are a better proxy for GDP per capita in urban areas than in rural areas (Chen and Nordhaus, 2019; Gibson et al., 2021). I therefore examine whether my results are primarily driven by urban areas. To do so, I first categorise all grid cells based on their population density relative to all other grid cells of the same country and then estimate my model on subsets of the data. I find that the magnitude of my estimated effects is largest in urban areas defined as the 5 percent most densely populated grid cells of every country (Appendix A.10). At the same time I find significant effects of the same sign also in less densely populated areas, including the 50

percent least densely populated areas. The effect of temperature variability thus seems to be geographically widespread and not limited to urban areas.

1.6.2 Agriculture

A possible explanation for my empirical effects is that temperature variability affects the local sectoral composition of economic activity. For instance, regions with higher seasonal variability could be relatively more or less suitable for agriculture than regions with lower variability. Because agricultural activity tends to be associated with lower levels of nightlights than other economic activity for a similar total economic output (Chen and Nordhaus, 2019; Gibson et al., 2021), these climatically induced relative advantages could be reflected in the spatial distribution of nightlights and thus explain the estimated coefficients.

I thus use satellite data on agricultural land use to examine heterogeneity of the estimated effects across regions with different types of economic activity. I include both cropland and land used for pasture in my model. I find that land used for pasture has indeed a significant effect on nightlights, but my estimates of temperature variability remain unaffected by including either one or both of these variables (Table A.13 in Appendix A.11). These results suggest that the estimated effect of temperature variability on nightlights cannot be explained with the spatial distribution of agricultural activity.

1.7 Conclusion

In this paper I combine a global high-resolution satellite-derived dataset on nightlights with climate reanalysis data and additional geographical datasets to examine how day-to-day, seasonal, and interannual temperature variability affect economic activity. I use a spatial first-differences research design (Druckenmiller and Hsiang, 2018), which reduces potential omitted variable biases and multicollinearity of climate and geographical variables. This allows me to study how temperature variability at different time scales influences aggregate economic activity. Furthermore, compared to previous work on the short-run effect of annual weather fluctuations (Burke et al., 2015a; Dell et al., 2012), I focus on the long-run effects of climate including the potential effect of adaptation (Waldinger, 2022). Compared to previous work on the long-run effect of climate (Nordhaus, 2006; Mendelsohn and Massetti, 2017), I use a recently developed econometric framework which allows for a

more plausible identification of causal effects.

This approach allows me to identify the total effect of temperature variability including possibly non-linear effects of daily temperature levels (ex post effects) and larger uncertainty (ex ante effects). Evidence from micro-econometric studies suggest this total effect to be predominantly negative. Furthermore, I expect these effects to be more negative for variability at larger time scale due to more difficult learning about variability and to be more negative for variability that is less predictable. Because of the underlying physical processes of climate day-to-day and interannual variability are less predictable than seasonal variability.

I find a statistically significantly negative effect of day-to-day variability on economic activity across the range of observed annual mean temperatures. On average, one additional degree Celsius of the average within-month standard deviation of daily temperature levels reduces economic activity by about 11 log points (approximately 11 percent). Regarding seasonal variability, I also find a negative but smaller and less significant effect on economic activity. On average, one degree Celsius of the average within-year range of monthly mean temperatures reduces nightlights by about 0.6 percent. My results on interannual variability suggest that it has a positive effect at low temperature levels (about 19 percent per degree Celsius of the between-year standard deviation of annual mean temperatures) and a negative effect at high temperature levels (about 28 percent per degree Celsius).

While I am to my knowledge the first to empirically analyse the effect of seasonal and interannual variability on aggregate economic activity, the results align with previous work finding a negative short-term effect of day-to-day variability on regional GDP (Kotz et al., 2021b) and existing literature reporting negative effects of temperature variability on agriculture (Wheeler et al., 2000; Mendelsohn et al., 2007b) and health (Hovdahl, 2020). Furthermore, consistent with previous findings (Burke et al., 2015b; Kalkuhl and Wenz, 2020), I find a positive marginal effect of annual mean temperature at relatively low temperatures and a negative marginal effect at high temperatures, with a globally optimum annual mean temperature of about 20 degrees Celsius.

My methodology allows me to compare the estimated coefficients of annual mean temperature with the estimated effects of temperature variability. On average, one sample standard deviation of seasonal and day-to-day variabil-

ity reduces nightlights by 9 and 17 percent, respectively. If these effects are benchmarked with the estimated effect of annual mean temperature, they correspond to increases of annual mean temperature from 25 degree Celsius to approximately 28 and 30 degrees Celsius, respectively.

I explore several explanations for my findings. Regarding my estimated effect of interannual variability I find one possible explanation consistent with my results: Below the optimal temperature, the positive effect of unexpectedly warmer-than-average temperatures could be larger than the negative effect of colder-than-average temperatures, whereas above the optimal temperature, the negative effect of unexpectedly warmer-than-average temperatures could be relatively larger. This could be the case, for example, if responding to a colder-than-average year was generally easier or less costly than responding to a warmer-than-average year.

It is well known that nightlights are a better proxy for GDP per capita in urban areas than in rural areas and to some extent also reflect the local sectoral composition of the economy (Chen and Nordhaus, 2019; Gibson et al., 2021). I therefore examine whether my results are primarily driven by urban areas and whether my results can be explained by the spatial distribution of agricultural activity. I find that the estimated effects of temperature variability are indeed strongest in urban areas, but can also be recovered from less densely populated regions, including the least densely populated areas within countries. Furthermore, my results are unaffected by controlling for the spatial distribution of agricultural activity.

My results are robust to a variety of robustness tests. I also recover my main results also for models that include no control variables and for models for which all control variables are included with flexible functional forms. Furthermore, my results pass a robustness test unique to the spatial first-differences research design, namely comparing estimates obtained by using spatial differences in orthogonal geographical directions. Furthermore, I conduct a sensitivity analysis which indicates that an omitted variable would need to be more strongly associated with nightlights and temperature variability as any of my geographic and climate control variables to render the estimated coefficients of temperature variability insignificant. In an additional robustness test, I show that my results cannot be explained by reverse causality.

My results suggest that more research should be devoted to temperature

variability. Several avenues could be pursued. For example, temperature variability might have influenced the initial spatial allocation of economic activity but could also have shaped subsequent local economic development including sectoral specialisation (Emerick, 2018; Henderson et al., 2017) and migration (Cattaneo and Peri, 2016). My methodology does not allow me to separate these two effects as nightlights are available only for recent decades. Furthermore, research is needed to investigate how the influence of temperature variability changes as economies develop, and to examine specific mechanisms in more detail.

Further research seems especially important given that climate models project changes to the seasonal cycle (Dwyer et al., 2012) and interannual temperature variability (Bathiany et al., 2018). Also day-to-day temperature variability seems to be influenced by anthropogenic climate change (Kotz et al., 2021a; Wan et al., 2021). Given the geographical patterns of these trends and projections, my results suggest that future increases of day-to-day, seasonal, and interannual temperature variability might add to the economic costs of climate change especially in currently relatively warm regions. I thus conclude that temperature variability at different time scales deserves more attention in economics research to improve our understanding of its influence on human societies and to get better estimates of the expected costs of future climate change and their geographical distribution.

Chapter 2

Seasonal temperature variability and economic cycles

In this paper, I examine to what extent seasonal temperature variability can explain seasonal economic cycles. To this aim, I first construct a novel dataset of seasonal temperature and seasonal GDP for a sample of 81 countries. This dataset reveals a much larger diversity of seasonal economic cycles around the world than previously reported. I then attribute these economic cycles to variation in temperature. For identification, I propose and apply a novel econometric approach that accounts for expectations and is based on seasonal differences. The results suggest that seasonal temperature has a statistically significant positive effect on seasonal GDP. The effect appears large, as seasonal temperature can explain a substantial share of the variation in seasonal GDP. Using data on GVA for different industry groups I can attribute this effect to industries that are relatively more exposed to ambient temperature. Furthermore, the results suggest that economic development makes countries more resilient to temperature fluctuations. Regarding future anthropogenic climate change, the results suggest that changes to seasonal temperatures will lead to a reallocation of economic activity from one season to another of up to several percentage points of annual GDP, pointing to a channel through which climate change will affect economic production that has so far been overlooked.

2.1 Introduction

A large part of the variation of time-series of macroeconomic variables is due to seasonality (Hylleberg et al., 1993). Understanding the causes of this seasonality has been an active area of macroeconomic research. While it has long been conjectured that some of the seasonality can be attributed to weather, research has come to the conclusion that observed quarterly variation of Gross Domestic Product (GDP) can mostly be explained by recurring shifts in preferences and technologies due to high consumption around Christmas and mid-year vacations (Beaulieu et al., 1992; Barsky and Miron, 1989; Cubadda et al., 2002; Beaulieu and Miron, 1992; Braun, 1995; Chatterjee and Ravikumar, 1992; Miron and Beaulieu, 1996; Franses, 1996). Furthermore, it has been pointed out that an important role of temperature seems to be in contradiction with similar economic cycles observed in countries in different hemispheres experiencing opposite seasons (Beaulieu et al., 1992). However, these conclusions were based on small samples of mostly OECD countries and little attention was paid to causal identification and to attribution of observed fluctuations to fundamental rather than proximate drivers. Given that anthropogenic climate change is projected to change seasonal cycles of temperature (Dwyer et al., 2012), the role of temperature for fluctuations of GDP appears to be an important question.

In this study, I empirically examine the influence of temperature on seasonal economic cycles. To do so, I construct a new dataset covering the period 1981-2020 using a global dataset of quarterly GDP covering 81 countries, a dataset on quarterly Gross Value Added (GVA) for 35 European economies, and climate reanalysis. Using information on quarterly temperature, I define seasons in a consistent way across countries in different hemispheres. For causal identification, I propose and apply a novel estimation strategy that is based on variation across countries in the differences in temperature and GDP between summer and winter.

I first use this novel dataset to identify stylised facts about quarterly fluctuations of GDP around the world. Previous studies were based on fewer economies mostly located in the Northern hemisphere and reported relatively similar cycles across countries with a primary peak of production in the fourth quarter and a trough in the first quarter. In contrast to these results, I find a large diversity of quarterly economic cycles. Of the 24 possible quarterly patterns, 15 are observed by at least one country in the sample.

Economic cycles also seem to systematically differ between countries in the Northern and in the Southern hemisphere. Next I aggregate quarterly production to production in two seasons (Q1+Q4 and Q2+Q3), to which I refer as summer and winter depending on which season tends to be warmer. I find that production is larger in summer in 44 countries and larger in winter in 37 countries.

I next examine the contribution of temperature to these cycles and find that seasonal differences in temperature between summer and winter have a statistically significant positive association with seasonal differences in GDP. The estimated coefficient is robust to the inclusion of a variety of control variables, including annual mean temperature and the level of GDP per capita. The estimates are also robust to the choice between nominal and real GDP and to changing the time period from 1991-2020 to 2011-2020. Furthermore, I find similar effects if I consider the warmest and coldest quarter as summer and winter, respectively. Overall, the effect of temperature appears large, similar in size to the average observed seasonal economic cycle.

These results could be explained by several mechanisms through which temperature affects economic activity. I first use the global sample of countries and explore the role of agriculture, tourism, and international trade. I do not find evidence suggesting that any of these possible channels is important. I next use data on GVA for different industry groups for European economies and find a statistically significant effect of seasonal temperature on GVA only for industries in which production is relatively exposed to ambient temperature. At the level of industries, I can attribute this effect primarily to Construction, Industry, and Manufacturing. The results hence appear consistent with an effect of temperature on the supply side of the economy.

In the last part of the paper, I examine possible consequences of climate change. To do so, I also estimate a long differences version of the seasonal differences model. Despite the different identifying assumptions, the results are qualitatively and quantitatively similar to the results of the cross-sectional seasonal differences estimation. Specifically, I find that between 1981-2000 and 2001-2020, seasonal GDP increased relatively more if one season warmed more than the other season. This seems to be primarily due to a reallocation of economic activity between the two seasons, as I do not find evidence that seasonal warming had an effect on annual GDP.

I then combine my estimates with projections of seasonal temperature from

climate models for alternative scenarios of future climate change. The results suggest that changes to the seasonal temperature cycle will cause a reallocation of economic activity across seasons of up to several percentage points of annual GDP. The results indicate a lot of variation between countries in terms of the projected reallocation. For a scenario of strong climate change (RCP8.5), we find that on average seasonal economic cycles are projected to increase. The magnitude of these effects is relatively uncertain, but central estimates suggest that in some countries seasonal economic cycles are projected to even double in size.

This paper contributes to prior work on seasonal economic cycles which has so far explained them primarily with recurrent shifts of preferences and technologies (Beaulieu et al., 1992; Barsky and Miron, 1989; Cubadda et al., 2002; Beaulieu and Miron, 1992; Braun, 1995; Chatterjee and Ravikumar, 1992; Miron and Beaulieu, 1996; Franses, 1996; Lumsdaine and Prasad, 2003). In contrast to this prior work, I find that some of the previously identified stylised facts can be observed only in about half of all countries because of large heterogeneity of seasonal economic cycles across countries. Furthermore, I find that the average effect of seasonal temperature is of a similar magnitude as the average seasonal economic cycle. This result does not rule out that preference and technology shocks are important channels, but points to the possibility that temperature is one fundamental driver of those shifts.

This paper also contributes to previous work on the effect of temperature on economic production. Previous work suggests a positive effect of annual mean temperature on economic production in relatively cold (and rich) and a negative effect in relatively warm (and poor) countries (Dell et al., 2012; Burke et al., 2015a; Kalkuhl and Wenz, 2020). In this paper, I find evidence for an overall positive effect of seasonal temperature on seasonal production in countries with larger production in summer than in winter, but I also explain how this estimated effect is conceptually different from the effect of annual temperature on annual GDP estimated in previous studies.

The paper is structured as follows. In the next Section, I present the theoretical framework, explain the identification strategy, and describe the data used in this study. In Section 3, I first present stylised facts of seasonal economic cycles for my global sample of countries. I then discuss results obtained from my econometric estimation, before showing stylised facts and econometric

results for the data on industry groups for countries in Europe. Furthermore, I combine my empirical estimates with results from climate models to quantify the order of magnitude of future possible seasonal reallocation of economic production. Conclusions are drawn in Section 5.

2.2 Methods

2.2.1 Theoretical framework

Identifying the causal effect of temperature on economic production requires an empirical framework that takes into account expectations. This is especially important for seasonal changes of temperature which are recurring every year and thus likely to be anticipated. In essence, seasonal cycles of temperature can be considered as a characteristic of the climate of a location, rather than its weather. To illustrate the challenge of causal identification in the presence of expectations and to explain the solution proposed in this paper, I start by formulating a simple conceptual model of economic production Y as a function of climate C and other factors X . I follow Hsiang (2016) and assume that climate influences production through two channels: through the actually realised weather c and through beliefs about climate b :

$$Y(C, X) = Y[c(C), b(C), X] \quad (2.1)$$

In this framework, both climate C and weather c are characterised by meteorological variables that describe the state of the atmosphere, such as temperature, precipitation, and humidity. The difference between the two concepts is that climate C refers to the (theoretical) probability distribution of these variables, while weather c refers to the (empirical) frequency distribution of their actually realised values. In other words, climate refers to the population of possible events, whereas weather refers to a sample drawn from that population. Weather can affect economic production directly for example through effects of precipitation on agricultural output or effects of temperature on the productivity of labour. Beliefs b are based on climate and affect economic production through actions of economic agents that are influenced by the expected future weather, such as the choice of production technology.

Climate and weather are specific to a location and a specific time period. Climate is typically defined for a period of 30 years, whereas weather is defined for shorter periods (hours, days, maybe a year). The term climate is com-

monly also used to refer to the statistics of weather of only certain parts of a year. For the purpose of this paper I use the term *seasonal climate* to refer to the climate of specific months. For example, seasonal climate can refer to the average weather of the months January, February, and March in London over the time period 1981-2010.

Given Equation 2.1 the marginal effect of (seasonal) climate on production can be written as

$$\frac{\partial Y(\mathbf{C})}{\partial \mathbf{C}} = \sum_{k=1}^K \frac{\partial Y(\mathbf{C})}{\partial \mathbf{c}_k} \frac{d\mathbf{c}_k}{d\mathbf{C}} + \sum_{n=1}^N \frac{\partial Y(\mathbf{C})}{\partial \mathbf{b}_n} \frac{d\mathbf{b}_n}{d\mathbf{C}} \quad (2.2)$$

The (marginal) effect of climate on production can hence be considered as the sum of direct effects (first term of Equation 2.2) and belief effects (the second term of Equation 2.2). For simplicity, it is assumed here that agents form their beliefs based on only climate and not weather.

2.2.2 Identification strategy

The decomposition of the marginal effect of climate on economic production into two channels has implications for its identification in empirical research. This identification can generally be based on variation across time or across units of observation. Depending on this choice, the two channels in Equation 2.2 will be captured to a greater or lesser extent by empirical estimates. Generally, variation of output across units of observations includes both direct and belief effects of climate, but cross-sectional estimates are prone to omitted variable biases. Exploiting variation of temperature and output over time at a frequency of days, months, or years removes possible biases of unobserved time-invariant effects, but is unlikely to recover belief effects of climate.

This trade-off between a plausible identification of causal effects of climate and the credible identification of both direct and beliefs effects of climate is a thread throughout the climate econometrics literature (Hsiang, 2016). For the purpose of this paper, I propose a new empirical strategy for navigating this trade-off. The strategy relies on temperature differences between two seasons of the same year. It can be considered a hybrid approach, exploiting variation across time and across units of observations for identification. In this respect, it resembles the long differences approach of panel data analysis

(Hsiang, 2016). In mathematical terms, I propose to estimate an Equation:

$$Y_{i\tau_1} - Y_{i\tau_2} = \alpha_{SD} + (\mathbf{c}_{i\tau_1} - \mathbf{c}_{i\tau_2})\beta_{SD} + (\mathbf{x}_{i\tau_1} - \mathbf{x}_{i\tau_2})\gamma_{SD} + \tilde{\mathbf{x}}_i\delta + \epsilon_i \quad (2.3)$$

where seasonal weather over a time period of several years is indexed by τ_1 and τ_2 , with a vector of time-varying controls \mathbf{x} , and with a vector of season-invariant controls $\tilde{\mathbf{x}}$. The two seasons can be considered as any two time periods within a year for which both temperature and production are observed. In the empirical part of the paper, I distinguish two seasons summer and winter and use two alternative ways of assigning the four quarters of a year to these two seasons (Section 2.2.3).

Identification of a causal effect of seasonal climate using Equation 2.3 relies on a special form of the *unit homogeneity assumption*:

$$E[Y_{i\tau_1} - Y_{i\tau_2} | \mathbf{c}_{\tau_1} - \mathbf{c}_{\tau_2}, \mathbf{x}_{i\tau_1} - \mathbf{x}_{i\tau_2}, \tilde{\mathbf{x}}_i] = E[Y_{j\tau_1} - Y_{j\tau_2} | \mathbf{c}_{\tau_1} - \mathbf{c}_{\tau_2}, \mathbf{x}_{j\tau_1} - \mathbf{x}_{j\tau_2}, \tilde{\mathbf{x}}_j] \quad (2.4)$$

or, using the greek letter Δ to denote seasonal differences,

$$E[\Delta Y_i | \Delta \mathbf{c}, \Delta \mathbf{x}_i, \tilde{\mathbf{x}}_i] = E[\Delta Y_j | \Delta \mathbf{c}, \Delta \mathbf{x}_j, \tilde{\mathbf{x}}_j] \quad (2.5)$$

This assumption differs from the unit homogeneity assumption of a conventional cross-sectional regression in that it does not require that the expected *levels* of production are the same for two units of observation conditional on the level of climate and on observables, but that expected *seasonal differences* of production are the same for two units of observation conditional on the same seasonal differences in climate and conditional on observables. This means that the effect of any time-invariant variables that affect production in both seasons in the same way, such as the level of education of the workforce, cannot confound the estimated relationship.

Furthermore, the identification of both direct and belief effects of differences in the seasonal climate on differences in economic production relies on a *treatment comparability assumption*

$$E[Y_i | \mathbf{c}_{\tau_1}] - E[Y_i | \mathbf{c}_{\tau_2}] = E[Y_i | C_{\tau_1}] - E[Y_i | C_{\tau_2}] \quad (2.6)$$

This assumption is more credibly satisfied the longer the time period used to characterise seasonal weather c.

While seasonal differences can be used to estimate the effect of any climate variable, the focus of this paper is on temperature. Temperature differences between summer and winter are primarily determined by the amplitude of the annual cycle of the intensity of the Sun's radiation at the surface and are thus larger at higher latitudes. Seasonality is also larger on land than over the oceans due to the smaller heat capacity of land surface materials than water. For the same reason, seasonal temperature differences tend to be larger in the East than in the West on large continents at mid-latitudes (North-America, Eurasia) because the wind blowing from West to East leads to more continental climate in the East.

Empirical work on the factors underlying seasonal economic cycles point to a preference shift around Christmas and a technology shifts around July and August due to vacations. To address omitted variable biases, I include seasonal differences in rainfall and annual mean temperature as control variables in all regressions. This is important because geography affects climate in several ways and previous work suggests that also annual mean temperature and rainfall affect economic production.

To explore possible channels through which seasonal temperature cycles might affect seasonal economic cycles, I examine the influence of GDP per capita, the share of agriculture of GDP, import shares, export shares, and international tourism expenses and receipts as a share of GDP. For each of these variables, I first regress that variable on seasonal temperature differences and then include it as a control variable in the main regression. Furthermore, I examine the effect of temperature on GVA for 11 industry groups for 35 European economies, which sheds further lights on the sensitivity of specific sectors and possible underlying mechanisms.

In additional robustness tests whose results are shown in the Appendix, I include additional control variables. Specifically, I use the share of the Christian population and the share of Muslim population as proxies of the consumption boom around Christmas and to control for cultural differences (including public holidays) across countries more broadly. Other controls are the latitude of a country and the average real interest rate. Furthermore, I check robustness to using either real or nominal quarterly GDP and to changing the time period from 1991-2020 to 2011-2020.

The seasonal differences approach proposed here resembles the long differences approach because both can be considered hybrid approaches that use variation over time and over space for identification (Hsiang, 2016). The interpretation of the obtained estimates is however different. An important difference between the two is that the seasonal difference estimator allows one to identify the effects of beliefs about differences between summer and winter, whereas the long differences estimator accounts for beliefs about long-term trends. In terms of the requirement of data, the seasonal differences approach requires observations at sub-annual frequency (e.g. quarters or months) of at least the treatment and the outcome variable, but it does not require observations going as far back in time as the long differences approach. Furthermore, the seasonal differences approach benefits from a larger temperature treatment at least in countries at higher latitudes for which temperature differences between summer and winter exceed gradual temperature trends by about one order of magnitude.

Estimates obtained from seasonal differences are potentially more prone to omitted variable biases than results based on long differences and are less indicative of the effects of future changes to the climate of one season or to the annual climate. However, they cannot be biased due to confounding long-term trends of unobservables that might affect long difference estimation.

The two approaches are not necessarily exclusive. In the last part of the paper I therefore combine the two approaches. This addresses some remaining concerns about omitted variable biases. Furthermore, it is more likely that the obtained estimates are indicative of the consequences of future climate change. In mathematical terms, I estimate an Equation:

$$\begin{aligned}
(Y_{i\tau_1^B} - Y_{i\tau_2^B}) - (Y_{i\tau_1^A} - Y_{i\tau_2^A}) = & \alpha_{LD} + ((\mathbf{c}_{i\tau_1^B} - \mathbf{c}_{i\tau_2^B}) - (\mathbf{c}_{i\tau_1^A} - \mathbf{c}_{i\tau_2^A})) \beta_{LD} \\
& + ((\mathbf{x}_{i\tau_1^B} - \mathbf{x}_{i\tau_2^B}) - (\mathbf{x}_{i\tau_1^A} - \mathbf{x}_{i\tau_2^A})) \gamma_{LD} \\
& + ((\tilde{\mathbf{x}}_{i,B} - \tilde{\mathbf{x}}_{i,A}) \delta_{LD} + \epsilon_i \tag{2.7}
\end{aligned}$$

where A and B index two time periods, an earlier and a later time period. For example, in the main specification in the results section, $Y_{i\tau_1^A}$ is the average of log GDP of country i in winter over the time period 1981-2000, while $Y_{i\tau_1^B}$ is the same average over the period 2001-2020.

In the remainder of this paper, I will denote seasonal differences by Δ and long differences by Δ_{LD} . With this notation, Equation 2.7 can be written as:

$$\begin{aligned}\Delta_{LD}\Delta Y_i &= \alpha_{LD} + \Delta_{LD}\Delta \mathbf{c}_i\beta_{LD} \\ &+ \Delta_{LD}\Delta \mathbf{x}_i\gamma_{LD} \\ &+ \Delta_{LD}\tilde{\mathbf{x}}_i\delta_{LD} + \epsilon_i\end{aligned}\tag{2.8}$$

This combined approach has the advantage that the estimate β_{LD} cannot be biased by any country characteristics that are either stationary or have parallel trends over time. This includes, for example, any geographical characteristics of countries which affect both seasonal differences in temperature and seasonal differences in GDP. Another advantage is that the estimates obtained from long differences are based on changes of temperature and economic production over time scales similar to those of anthropogenic climate change.

2.2.3 Data


I use data on quarterly Gross Domestic Product (GDP) in USD provided by the International Monetary Fund (IMF). The data cover 81 countries with different temporal coverage across countries. The data are provided in nominal and real terms and the temporal coverage differs between the two products for some countries. I restrict the data to the time period 1980-2020. The data include at least 7 years of observations for every country (the first year in which data are available for Honduras is 2014 and for the Maldives 2012; for all other countries I have at least 10 years of data). In order to improve the balance of the panel data and informed by the definition of climate as an average over 30 years, I reduce the sample to the years 1991-2020 for the main estimation. As reported in the Appendix, the results are robust to using data only for the years 2011-2020. I combine this economic data with the climate reanalysis ERA5 provided by the European Center for Medium Range Weather Forecast (ECMWF). I use monthly mean temperature levels and monthly mean daily precipitation which I aggregate to quarterly frequency. The data have a spatial resolution of 0.25 degrees (approximately 25 km at the Equator) which I aggregate to the level of countries using grid-cell population from the Gridded Population of the World (GPW) dataset as

weights.

I also use data on quarterly Gross Value Added (GVA) for 11 industry groups provided by EUROSTAT. The data cover 35 countries in Europe. The data cover again different time periods across countries, with all countries reporting for at least 10 years (since 2009).

To identify seasonal patterns in the time-series, I first detrend the data. This also means that differences between nominal and real GDP are restricted to changes of prices between the seasons. In robustness tests shown in the Appendix, I find that the results are robust to using either of the two. For detrending I use a Hodrick-Prescott Filter with $\lambda = 50$. This filter has received criticism regarding its use in timeseries econometrics (Hamilton, 2018), but these issues are of minor concern here because the filter is only used to subtract a gradual trend. The detrended data are then averaged over time and identification is obtained from cross-sectional variation. After applying the filter and removing deterministic trends, I add the mean value of the last year in the time-series. The process is illustrated for time-series of the USA in Figure 2.1.

Figure 2.1: Time-series of quarterly real production for the USA before (left) and after detrending (right).



The identification strategy requires to define two seasons consistently across locations. The seasonal cycle of temperature is due to the tilt of the Earth's rotation axis and driven by the movement of the Earth around the Sun. From an Earth-centric perspective, the seasonal cycle of temperature arises from a perpetual oscillation between the time period with the maximum and the time period with the minimum of the amount of Solar radiation received at the top of the atmosphere. Except for locations close to the Equator, where variation in the distance between Earth and Sun dominates the oscillation of received Solar radiation, the time periods of minimum and maximum ir-

radiation are around mid of December and mid of June respectively in the Northern hemisphere. In the Southern hemisphere, the pattern is the opposite.

For data with quarterly frequency a natural choice is thus to aggregate quarterly data to two time periods summer and winter. For a country in the Northern hemisphere the quarters 2 and 3 (months 4-9) can generally be considered as summer ($\tau_j, j \in \{1, 2\}$) and the quarters 1 and 4 (months 1-3 and 10-12, respectively) as winter (τ_{3-j}). For countries not too close to the Equator, winter and summer defined this way will result in warmer and colder six months periods, respectively. Countries close to the Equator can experience more complex seasonal cycles with several peaks and troughs over the course of a year. For the empirical part of the paper I thus aggregate the four quarters to two seasons and then categorise the two six months periods as summer (S) and winter (W) for every country based on their average temperature. This ensures that the period referred to as summer is everywhere the warmer of the two six months periods of the year. I then use this assignment of the four quarters of a year to summer and winter to sum the detrended quarterly GDP for every country to seasonal GDP, that is GDP in summer and GDP in winter (illustrated for the USA and Brazil in Figure B.11 in the Appendix).

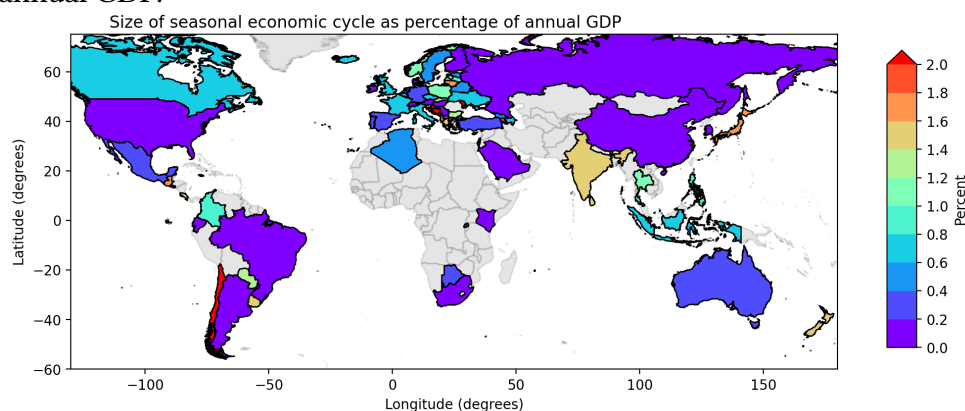
This definition of seasons has the advantage that economic production in all four quarters of the year is taken into account. An alternative choice is considering the warmest quarter as summer and the coldest quarter as winter. Results obtained with this definition of seasons are very similar and shown as a robustness check in the Appendix.

The magnitude of the seasonal cycle, defined as the difference between GDP in summer and winter divided by the annual GDP, is shown in Figure 2.2. The map reveals some geographical heterogeneity but no apparent pattern associated with latitude.

For control variables I use several sources. Data on GDP per capita, land area, and the share of agriculture and manufacturing are taken from the World Development Indicator database of the World Bank. For data on trade, tourism, and interest rate I use the TC360 database of the World Bank. Information on religion is obtained from the Pew Research Center.

Projections of future climate change are taken from the CMIP6 ensemble as

Figure 2.2: Size of the seasonal economic cycle (difference between production in the summer half-year and in the winter half-year) as percentage of annual GDP.



provided by the ECMWF. The model MPI-ESM1.2 is chosen as previous studies have shown relatively small biases for historical seasonal temperatures (Xu et al., 2021). Reassuringly, results for Europe also suggest that future warming of seasonal mean temperatures is robust across the model ensemble (Carvalho et al., 2021). I download monthly mean values for the historical period 1990-2014 and for the future periods 2041-2070 and 2071-2100. The monthly means are then used to calculate seasonal means. The seasons are defined as for the empirical analysis described above.

The analysis of future projections is based on future changes instead of future absolute values. This has the advantage that no bias correction is required, as future changes are calculated from simulations of past and future climate with the same climate model. This approach is also referred to as the delta method and very common in climate impact research. To calculate future changes I first compute mean values for both periods, 2041-2070 and 2071-2100, and then subtract the mean value of the historical period 1990-2014. All variables are aggregated from grid cells to the country level using the same population weights as for the ERA5 reanalysis data.

Descriptive statistics are shown in Table 2.1.

Table 2.1: Descriptive statistics.

Variable	Unit	Mean	Std.	Min.	Max.	No. obs.
$\Delta \log \text{GDP}$	USD 2010	-0.005	0.03	-0.17	0.04	81
ΔT	deg. C	-8.832	4.94	-18.75	-0.04	81
ΔP	mm day ⁻¹	-0.020	0.05	-0.18	0.10	81
$\Delta_{LD} \Delta \log \text{GDP}$, 2001-2020 minus 1981-2000	deg. C	-0.004	0.02	-0.05	0.03	60
$\Delta_{LD} \Delta T$, 2001-2020 minus 1981-2000	deg. C	-0.074	0.39	-1.30	1.05	60
Annual mean temperature	deg. C	15.394	6.74	4.11	27.87	81
Change in ΔT for RCP4.5, 2041-2070 minus 1990-2014	deg. C	0.058	0.33	-0.98	0.69	81
Change in ΔT for RCP4.5, 2071-2100 minus 1990-2014	deg. C	0.080	0.36	-0.78	0.91	81
Change in ΔT for RCP8.5, 2041-2070 minus 1990-2014	deg. C	-0.082	0.37	-1.00	0.50	81
Change in ΔT for RCP8.5, 2071-2100 minus 1990-2014	deg. C	-0.319	0.88	-2.25	1.30	81
Share of agriculture in GDP	percent	6.131	5.88	0.07	33.54	81
Share of exports of GDP	percent	44.203	32.43	11.13	187.44	81
Share of imports of GDP	percent	46.342	28.79	12.44	165.54	81
Share of tourism receipts of GDP	percent	12.269	12.07	0.40	62.81	81
Share of tourism expenditures of GDP	percent	6.571	3.68	1.07	25.50	81
Real interest rate	percent	5.914	6.61	-21.13	41.14	77
Share of Christian population	percent	62.823	33.08	0.17	100.00	81
Share of Muslim population	percent	14.330	27.63	0.01	98.05	81
log GDP per capita	USD 2010	9.805	0.83	7.34	11.65	81
Land area	1E6 km ²	11.735	2.28	5.77	16.61	81
Latitude	degrees	28.449	27.72	-41.00	65.00	81

Notes: Δ denotes seasonal differences, calculated as winter (W) minus summer (S). Δ_{LD} denotes long differences. Unless otherwise stated, statistics are based on averages over the period 1991-2020 for years in which there is quarterly GDP data for a given country.

2.3 Results

2.3.1 Stylised facts on seasonal economic cycles

Research over the last 35 years has mostly come to the conclusion that recurring shifts in preferences and technologies can explain most of the seasonal variation of economic activity (Beaulieu et al., 1992; Barsky and Miron, 1989; Cubadda et al., 2002; Beaulieu and Miron, 1992; Braun, 1995; Chatterjee and Ravikumar, 1992; Miron and Beaulieu, 1996; Franses, 1996). These shifts have been explained especially with high consumption during Christmas and vacations in June, July and August, leading to potentially very similar seasonal economic cycles across countries and industries. These conclusions were however based on small samples of countries, exclusively developed economies, and mostly located in the Northern hemisphere.


I hence first identify stylised facts about seasonal cycles in my sample of 81 economies. For every country I first regress trend-adjusted quarterly production on quarterly dummy variables. I then use the estimated coefficients of the four dummy variables to identify the pattern of the seasonal cycle. I distinguish 24 possible patterns. For example, the first of the 24 patterns corresponds to production tending to be largest in the first quarter, followed by the second, third, and fourth quarter.

Furthermore, to account for opposite seasonal cycles of temperature, I split the sample into countries located in the Northern hemisphere (NH) and countries in the Southern hemisphere (SH). To do so, I compare the average temperature of the months 10-12 and 1-3 with the average temperature of the months 4-9. A country is then assigned to the Northern Hemisphere if the months 4-9 are warmer than the months 10-12 and 1-3.

Overall, seasonal economic cycles around the world appear quite diverse, with 15 of the 24 possible patterns being exhibited by at least one country. The most common pattern in the sample (22 of 81 countries) is a peak of production in the fourth quarter, followed by the third, second, and first quarter (Figure 2.3). The second most frequent pattern (11 countries) is a peak in the third quarter, followed by the fourth, second, and first quarter.

This new evidence on quarterly cycles also reveals that some stylised facts identified by previous work are not as widespread as that work might suggest. One of these facts is a peak of production in the fourth quarter, possibly due to consumption boom around Christmas (Beaulieu and Miron, 1992). I

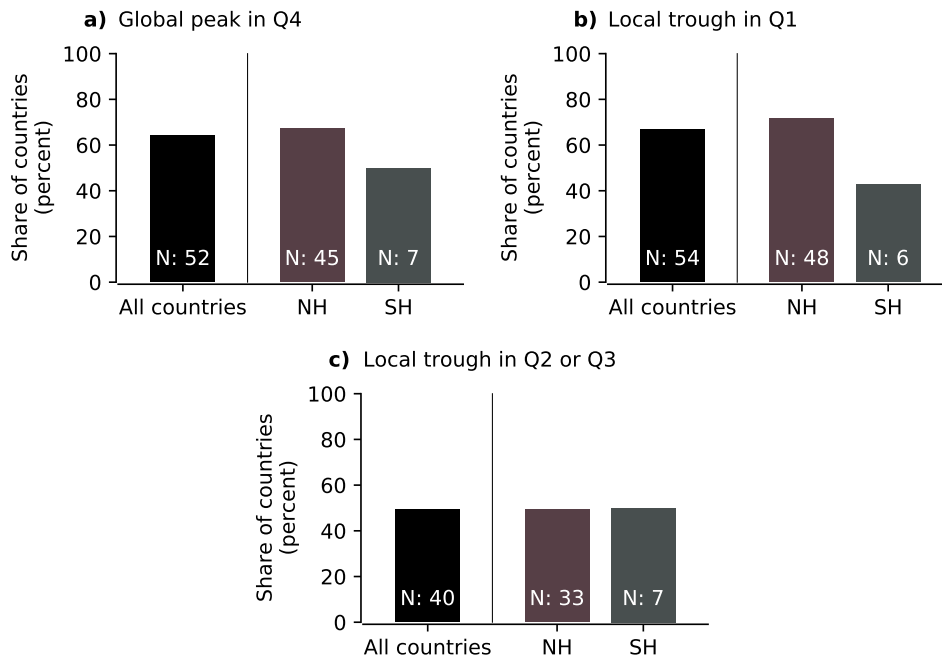
Figure 2.3: Frequency of patterns of quarterly economic production.



Notes: The figure shows the frequency of the 24 possible patterns of quarterly economic production among the countries in the sample. Every row corresponds to one pattern, which is shown with the blue bars on the left. The number in each cell is the number of countries that exhibit the corresponding pattern. Columns correspond to different samples of countries: full sample, countries in the Northern hemisphere, countries in the Southern hemisphere, and countries in Europe. Colors indicate relative frequency based on the size of the corresponding sample.

find that this is primarily a phenomenon of countries in the Northern hemisphere (Figure 2.4). In the full sample, about 64 % of countries (52 out of 81 countries) have the maximum production in the fourth quarter. These represent 67% of countries (45 of 67 countries) in the Northern hemisphere and 50% of countries (7 of 14 countries) in the Southern hemisphere.

Figure 2.4: Stylised facts identified in previous studies. Relative frequencies shown in percentages for different groups of countries.



Another stylised fact reported previously is a trough of production in the first quarter of the year, possibly due to reorganisation of production and generally economic activities at the beginning of the calendar year that result in less measurable economic output. Again I find that this can be found in countries in the Northern hemisphere (72% of countries) more frequently than among countries in the Southern hemisphere (43% of countries), but also that this fact is even in the Northern hemisphere only exhibited by about two thirds of all countries.

A third stylised fact reported previously is a slowdown of economic activity around June, July, and August, possibly due to school holidays in many countries and mid-year vacations. Such a local minimum of production in either the second or the third quarter can be found in 49% of countries in the Northern hemisphere and in 57% in the Southern hemisphere.

In sum, all three stylised facts about seasonal economic cycles seem to be found in only slightly more than half of all countries. Furthermore, two of these facts seem to be more frequently observed in the Northern hemisphere. One of these two, the peak of economic production in the fourth quarter, has been used to question the influence of temperature on seasonal economic cycles as countries in both hemispheres appeared to exhibit this feature in small earlier samples (Beaulieu and Miron, 1992). In contrast, the evidence of my larger sample suggests substantial differences between countries in the two hemispheres.

To reduce the dimensionality of the analysis and prepare the data for a model estimation based on seasonal differences, I next sum trend-adjusted production of the quarters 1 and 4 and 2 and 3 to semi-annual production. Based on the assignment of countries to the SH or NH, I refer to the quarters 1 and 4 as winter (summer) and to the quarters 2 and 3 as summer (winter) respectively. I refer to countries with larger production in summer as summer-peak countries and to all other countries as winter-peak countries.


In the full sample, summer-peak countries are slightly more frequent than winter-peak countries (54%, or 44 of 81 countries) (Figure 2.5). In the Northern hemisphere, the share of winter-peak countries is slightly larger (58%, or 39 of 67 countries). In the Southern hemisphere, countries tend to have larger production in winter than in summer (64%, or 9 of 14 countries).

Relating these findings to the cycles at quarterly frequency, in the Northern hemisphere the smaller production in winter than in summer can partly be explained with the small production in the first quarter, which seems to dominate the large production in the fourth quarter and the small production in the third or second quarter. In the Southern hemisphere, the relatively small production in summer is in line with only few countries exhibiting a peak in the fourth quarter and most countries exhibiting a local minimum of production in the second or third quarter.

Overall, countries in the Northern hemisphere thus tend to have larger production in winter than in summer. The fact that production in winter tends to be the sum of the quarter with the maximum production and the quarter with the minimum production suggests to examine the effect of temperature not only at semi-annual but also at quarterly frequency.

The geographical distribution of summer-peak and winter-peak countries

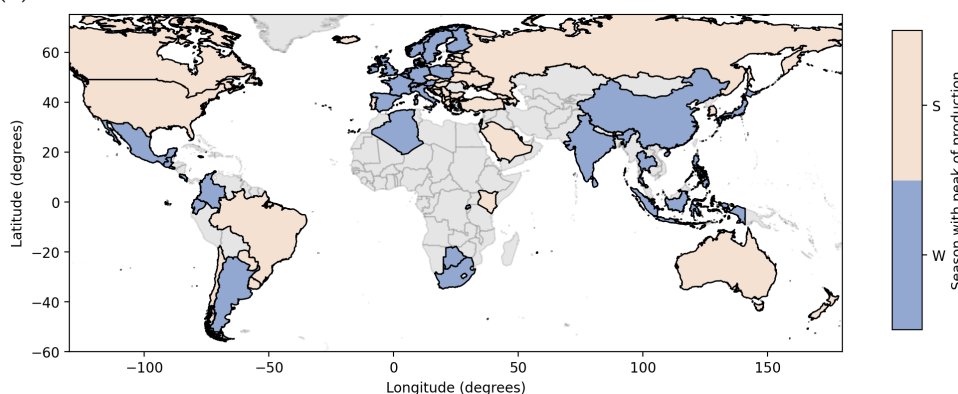
Figure 2.5: Frequency of patterns of economic production in summer and winter.



Notes: The figure shows the frequency of the 2 possible patterns of semi-annual economic production among the countries in the sample. Every row corresponds to one pattern, which is shown with the blue bars on the left. The number in each cell is the number of countries that exhibit the corresponding pattern. Columns correspond to different samples: full sample, countries in the Northern hemisphere, countries in the Southern hemisphere, and countries in Europe. Colors indicate relative frequency based on the size of the corresponding sample.

suggests some geographical clustering (Figure 2.6). For example, most countries in Northern and Western Europe are winter-peak countries, while most countries in Eastern Europe and the Middle East are summer-peak countries. In the Southern hemisphere, winter-peak countries are more common than summer-peak countries. There is however no clear effect of absolute latitude, with winter-peak countries being relatively frequent at relatively high and at low latitudes.

Figure 2.6: Geographical distribution of winter-peak (W) and summer-peak (S) countries.



I examine the balance of the two subsamples also more formally using statistical tests (Table B.21 in the Appendix). I find that winter-peak countries tend to have a smaller (less positive) temperature difference between summer and winter ($p < 0.05$). They also tend to be richer ($p < 0.05$) and have a smaller share of tourism receipts of GDP ($p < 0.05$). I further quantify the contribution of different factors to the observed economic cycle, including the role of temperature, in the next Section.

2.3.2 The contribution of seasonal temperature variability

In order to examine the contribution of temperature to seasonal economic cycles, I regress seasonal differences in GDP on seasonal differences in temperature (Equation 2.3) using the data on 81 countries illustrated in the previous section. I find a significant positive association between temperature and GDP (Column 1 in Table 2.2). This estimate is robust to including a variety of control variables including seasonal differences in rainfall, annual mean temperature, and GDP per capita (Column 2). The results are also robust to including a large number of additional control variables, to using only data from the years 2011-2020, and to using data on nominal or real quarterly GDP

(Figure B.51 in the Appendix). The results are also qualitatively the same for differences between the quarters with maximum and minimum temperature (Table B.81). As expected, using differences between six-months periods instead of quarters attenuates the estimated effect of temperature because in many countries the quarter with the maximum and the quarter with the minimum GDP fall within the same six-months period.

I next study heterogeneity in the effect of seasonal temperature. Including interaction terms in the model, I find that the effect of temperature is smaller in countries with higher level of GDP capita and turns negative for countries with a GDP per capita higher than about 36.000 USD (Column 3). I do not find evidence that other control variables in the model moderate the effect of seasonal temperature (Column 5).

Table 2.2: Results of regressions using a global sample of GDP of 81 countries.

Dependent variable:	$\Delta \log GDP$			
	1	2	3	4
ΔT	0.0018*** (0.0004)	0.0014** (0.0007)	0.0105** (0.0050)	0.0135** (0.0062)
$\Delta T \cdot \log GDP\ pc$			-0.0010* (0.0005)	-0.0009* (0.0005)
$\Delta T \cdot \text{Annual mean temperature}$				-0.0000 (0.0001)
$\Delta T \cdot \log Landarea$				-0.0003 (0.0002)
$\Delta \text{Precipitation}$		-0.1817** (0.0849)	-0.1787** (0.0835)	-0.1802* (0.0906)
Annual mean temperature		0.0004 (0.0005)	0.0005 (0.0004)	0.0004 (0.0010)
$\log GDP\ pc$		0.0107*** (0.0036)	0.0031 (0.0058)	0.0029 (0.0053)
$\log Landarea$		0.0024** (0.0010)	0.0020* (0.0011)	-0.0000 (0.0018)
R2	0.10	0.34	0.36	0.37
R2 adj.	0.09	0.29	0.30	0.30
N	81	81	81	81

Notes: Sample period is 1991-2020. Seasonal differences Δ calculated as winter minus summer. Significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The magnitude of the estimated effect of temperature is large. The sample mean of the seasonal difference in temperatures of about 8.8 degree Celsius (Table 2.1) is associated with a seasonal difference in GDP of about 1.2 per-

cent. This effect is of the same order of magnitude as the sample mean of the seasonal difference in GDP, which is about 0.5 percent (Table 2.1). The effect seems not to be driven by outliers and there is no apparent difference in the magnitude of the effect between countries in the Northern and in the Southern hemisphere (Figure B.31 in the Appendix). Temperature thus appears to be an important factor contributing to the larger production in summer than in winter that is observed in many countries (Figure 2.6).

To study possible mechanisms, I examine the role of agriculture, international trade, GDP per capita, and tourism. I first regress variables such as the share of agriculture of GDP, import shares, GDP per capita, and the share of international tourism expenses of GDP on the seasonal difference on temperature. I do not find any significant associations (Table B.41 in the Appendix). This suggests that these sectors are not the primary channels through which seasonal temperature cycles affect economies. As an additional test, I also include those variables as additional explanatory variables in my main specification (as possibly “bad controls”) and find that my main estimates barely change (Figure B.51 in the Appendix). In sum, the results suggest that other channels are primarily responsible for the estimated effect.

2.3.3 Effects by industry groups for European economies

The results in the previous section suggest that temperature has a positive effect on production in some countries and no effect or a negative effect in others. In this section, I explore to what extent this finding can be explained by differences in sectoral composition. To this aim, I use data on gross value added (GVA) by industry group for 35 countries in Europe. Focusing on Europe has the advantage that reporting quality is probably more homogeneous across countries than for the global sample and that also the climate and especially seasonal temperature cycles are more similar. Furthermore, EUROSTAT provides to my knowledge the most comprehensive homogeneous database of quarterly production by industry group.

Reassuringly, I find a similar significantly positive effect of seasonal temperature on seasonal GDP as for the global sample, with about twice the magnitude. I next estimate the seasonal differences model in Equation 2.3 for each of the industry groups. I follow previous literature and group industries according to whether labour is relatively more or less exposed to outdoor temperature (Behrer and Park, 2019; Acevedo et al., 2020). I accord-

ingly classify agriculture, construction, manufacturing, and other industries as relatively exposed. I find a significantly positive effect of seasonal differences in temperature on seasonal differences in GVA for total GVA and for GVA in exposed industries. For all other, non-exposed industries I find an insignificantly positive effect (Table 2.3). I conduct the same exercise at the level of individual industries and find that the positive coefficients for all exposed industries can be explained primarily with positive coefficients for Construction and other Industry, and possibly also Manufacturing, but not Agriculture (Table B.61 in the Appendix).

Table 2.3: Results of regressions using a sample of GVA by industry groups of 35 European economies.

Dependent variable:	$\Delta \log GDP$		
	All	Exposed	Non-Exposed
Column:	1	2	3
ΔT	0.0037** (0.0018)	0.0071** (0.0033)	0.0014 (0.0016)
Δ Precipitation	-0.0807 (0.1565)	-0.3897 (0.3410)	0.1322 (0.1485)
Annual mean temperature	-0.0022* (0.0012)	0.0001 (0.0024)	-0.0027** (0.0011)
log GDP pc	0.0179* (0.0088)	0.0339* (0.0175)	0.0082 (0.0075)
log Landarea	0.0052*** (0.0016)	0.0078** (0.0029)	0.0040*** (0.0014)
R2	0.54	0.49	0.32
R2 adj.	0.46	0.40	0.21
N	35	35	35

Notes: Seasonal differences Δ calculated as winter minus summer. Exposed industries: Agriculture, Construction, Manufacturing, other Industry. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In sum, the results suggest that exposure to ambient temperature in economic production is an important moderator, pointing to effects of temperature on the supply side of economies as a possible mechanisms.

2.3.4 Results from long differences

These results obtained from seasonal differences estimation are based on cross-sectional variation in seasonal differences and therefore have two caveats. The first is that there is still a risk of omitted variable biases from

omitted but important country characteristics. These could include geographical characteristics that influence both seasonal temperature variability and seasonal economic cycles. The second caveat is that the results are not necessarily indicative of the effects of changes over time to seasonal temperature variability in a specific country, as discussed in Section 5.2.

To overcome these limitations, I also estimate a model based on long differences of seasonal differences. Because of limited data availability at the quarterly frequency, I use the two time periods 1981-2000 and 2001-2020. This reduces the sample to 60 countries for which GDP data from the earlier period are available. For all these countries, both winters and summers became warmer between the two time periods. In about half of all countries winters warmed more strongly than summers (Figure B.91 in the Appendix).

The results are qualitatively similar to the results obtained from cross-sectional variation in seasonal differences (Table 2.4). Specifically, I find that seasonal temperature has a positive effect on seasonal GDP which becomes smaller and even negative for richer countries. The magnitude of the effect is larger. Furthermore, as for the cross-sectional estimates, seasonal rainfall has a negative effect on seasonal GDP. These results are robust to including trends in annual mean temperature and annual total precipitation (Column 5) and similar for countries for which winters warmed more than summers and countries with opposite trends (Column 6).

The magnitude of the effect estimated from long differences is larger than the magnitude of the effect estimated from the cross-section 1991-2020. This difference can primarily be explained with the different samples that are used in the two estimations due to a lack of data for the earlier period for some countries (Table B.71 in the Appendix). Once the differences between the two samples are taken into account, the coefficients have a similar magnitude. For example, for a country with a GDP per capita of about 22,000 USD, which corresponds to the mean value of the larger sample, the estimated coefficients are 0.011 and 0.015 log points of GDP per degree Celsius based on the cross-sectional and the long differences estimation respectively.

Table 2.4: Results of regressions with long differences using a global sample of GDP of 60 countries.

Dependent variable: Column:	$\Delta_{LD} \Delta \log GDP$					
	1	2	3	4	5	6
$\Delta_{LD}\Delta T$	0.0059 (0.0065)	0.0058 (0.0063)	0.1053** (0.0405)	0.0821* (0.0469)	0.1075** (0.0440)	0.1177*** (0.0421)
$\Delta_{LD}\Delta T \cdot \log GDP$ pc			-0.0105** (0.0040)	-0.0091** (0.0043)	-0.0106** (0.0043)	-0.0121*** (0.0042)
$\Delta_{LD}\Delta T \cdot$ Annual mean temperature				0.0009 (0.0008)		
$\Delta_{LD}\Delta T \cdot (\Delta_{LD}\Delta T > 0)$						0.0099 (0.0124)
$\Delta_{LD}\Delta$ Precipitation		-0.1122*** (0.0337)	-0.1595*** (0.0338)	-0.1259*** (0.0358)	-0.1826*** (0.0378)	-0.1918*** (0.0434)
Annual mean temperature				0.0006** (0.0003)		
log GDP pc			0.0088*** (0.0024)	0.0100*** (0.0021)	0.0105*** (0.0026)	0.0109*** (0.0026)
Δ_{LD} Annual mean temperature					-0.0094 (0.0065)	-0.0087 (0.0066)
Δ_{LD} Precipitation					0.1148 (0.1645)	0.0822 (0.1675)
$(\Delta_{LD}\Delta T > 0)$						0.0009 (0.0048)
R2	0.02	0.13	0.36	0.41	0.39	0.40
R2 adj.	0.00	0.10	0.32	0.35	0.32	0.30
N	60	60	60	60	60	60

Notes: Long differences Δ_{LD} calculated by subtracting mean over 1981-2000 from mean over 2001-2020. Seasonal differences Δ calculated as winter minus summer. For the purpose of the analysis, annual mean temperature and log GDP per capita (the 5th and 6th variable in this table, respectively) are considered as time-invariant variables. Their average values over 1991-2020 are included as a moderator variable in the same way as in the cross-sectional regression (Table 2.2). Significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

These results suggest that if one season warmed more than another between 1981-2000 and 2001-2020, it also witnessed a relatively stronger growth in GDP. In additional analysis, I do not find evidence for significant effects of changes in seasonal temperatures on changes in annual GDP (Table 2.4 in the Appendix), suggesting that the results in Table 2.4 are primarily driven by a reallocation of economic activity between summer and winter and not due to temperature increasing or decreasing economic production in only one of the two seasons.

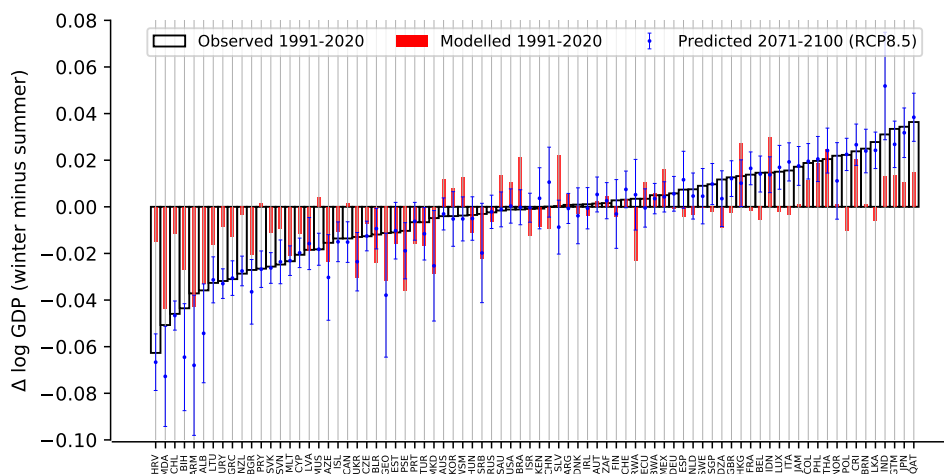
2.3.5 Scenarios of future climate change

These results suggest that economic production will be reallocated between winter and summer if under future climate change one season warms more strongly than the other. Such changes are indeed projected by global climate models. In a few countries, winters are projected to warm more quickly than summers because of reductions in snow cover in winter and accelerated warming due to the snow-albedo feedback (Carvalho et al., 2021). In most other countries outside the tropics, summers are projected to warm more quickly than winters because of increased dryness in summer and thus less surface humidity that can reduce the projected warming through latent heat transfer (Byrne, 2021). The following analysis aims to quantify the approximate order of magnitude of possible reallocations of economic activity across the seasons due to these changes.

To this aim, I combine climate model projections with the empirically estimated effect of seasonal differences in temperature on seasonal economic production from the long difference estimation with interaction term (Column 5 in Table 2.4). I first focus on the RCP8.5 scenario and the time period 2071-2100. I find that for some countries, production will shift more towards summer and for other countries towards winter (Figure 2.7). The magnitude of the projected changes varies greatly among countries, depending on their GDP per capita levels and projected changes to seasonal temperatures, and can be up to several percentage points of annual GDP large.

Reallocation of economic production between the seasons can increase or decrease seasonal economic cycles. For the RCP8.5 scenario, more countries experience a reallocation towards summer than towards winter, and the effect is particularly large, with a long tail, for countries that have their peak of production in summer (Figure 2.7 and Figure B.121 in the Appendix). On

Figure 2.7: Projections of Δ GDP by country for the global sample.



Notes: The plot shows the projected changes of Δ log GDP for individual countries for the RCP8.5 scenario based on the results from the long difference estimation. Seasonal differences Δ calculated as winter minus summer. Positive values mean that for the given scenario GDP will be reallocated from summer to winter.

average, seasonal economic cycles are projected to increase. This average effect masks however substantial heterogeneity across countries (Figure 2.7 and Figure B.131 in the Appendix). These results are for the far future and a scenario of strong climate change (RCP8.5). The projected changes are qualitatively similar for an earlier period (2041-2070) and an alternative scenario (RCP4.5) (Figure B.121 in the Appendix).

2.4 Conclusion

In this paper I study the effect of seasonal temperature on seasonal economic production. For causal identification I propose a novel econometric approach using variation of differences between seasons across countries. This seasonal differences estimator is applied to a global sample of 81 countries using quarterly data on GDP and climate reanalysis. The results suggest that differences in temperature between summer and winter can explain a major part of the observed differences in GDP between summer and winter. This finding is in contrast to previous work which concluded that temperature plays at most a minor role for seasonal cycles of GDP. This discrepancy can partly be explained with limited evidence available at the time of earlier studies, inappropriate methods to infer causality that neglected expectations, and possibly a focus on proximate (technology shocks, preference shocks) rather than fundamental drivers of economic fluctuations.

The analysis also reveals a large diversity of seasonal economic cycles and systematic differences between countries in the Northern and in the Southern hemisphere. Given that previous work focused on small subsets of countries, this global heterogeneity can potentially explain some of the differences from earlier studies. Furthermore, I find that in the majority of countries economic activity is larger in summer than in winter. Somewhat consistent with this finding, my results suggest an overall positive association between seasonal temperature and seasonal economic production.

This effect of seasonal temperature is both significant and large. On average it is of the same magnitude as observed differences in seasonal GDP. To address concerns about causal inference from cross-sectional variation, I conduct extensive robustness tests with a wide range of control variables, including seasonal differences in rainfall, annual mean temperature, GDP per capita levels, religious composition, geographical size of a country, latitude, and variables related to the sectoral composition of an economy, international trade, and international tourism. The results are also robust to considering the quarter with maximum and minimum temperature as summer and winter respectively and to shortening the time period to 2011-2020. Regarding possible mechanisms, results on the industry level for a subsample of European countries points to an important role of industries that are relatively exposed to ambient temperature, including Construction, Industry, and Manufacturing.

Regarding future climate change, the results suggest that economic activity will be reallocated between the seasons. However, the results do not allow conclusions about the extent to which annual GDP will be affected by future changes to annual mean temperature. That question has been the focus of a large body of prior literature which tends to agree that an increase in annual mean temperature has a negative effect on GDP in very warm countries and a possibly positive effect in relatively cold countries (Dell et al. (2012); Burke et al. (2015a); Acevedo Mejia et al. (2018); Colacito et al. (2019), among others).

To quantify the possible seasonal reallocation of economic production under future climate change, I combine empirical estimates obtained from a long differences specification of the model with the projections of climate models. The results point to substantial future changes to seasonal economic cycles, which are projected to more than double in some countries by 2071-2100

for the scenario RCP8.5. For the time period 2041-2070 and for the scenario RCP4.5, the results suggest a reallocation of economic production of up to one percentage point of seasonal GDP.

The results overall suggest that temperature should be taken into account in seasonal forecasts of economic production. While this is already the case in some countries (see e.g. Bundesbank (2012, 2014)), the results point to an influence of weather on seasonal economic cycles across a wide range of socio-economic and climatic contexts. Given that climate change will increase seasonal economic cycles in some countries, the results also suggest a future increase in demand for fiscal, monetary, and structural policies that help to smoothen quarterly fluctuations of production and employment. Furthermore, my results suggest that economic development can make economies generally more resilient to the influence of seasonal temperature variability, pointing to possible adaptation.

The quarterly GDP data used in this paper cover 81 countries around the world representing all continents and a large range of socioeconomic contexts and climates. The analysis of heterogeneity suggests that the effect of seasonal temperature on seasonal GDP decreases with the level of GDP per capita of a country, but this pattern is based on relatively few economies in Africa, demanding caution when extrapolating from the global sample to other countries.

The results also point to a new avenue of macroeconomic research on the fundamental drivers of fluctuations of GDP, employment, and prices accounting for the deterministic and the stochastic part of temperature variability. The evidence presented here suggests that temperature affects production through productivity shocks, but does not exclude that part of the estimated effects is also due to seasonal shifts in preferences. Disentangling the two with a structural model appears to be one promising research perspective. Furthermore, the analysis revealed that some countries have largest production in winter and others in summer. Future research could examine to what extent these opposite patterns can be explained with economic specialisation and trade.

Previous research has found negative effects of seasonal temperature variability on economic activity (Chapter 1). The results in this paper corroborate an influence of seasonal temperature variability on economic production. Furthermore, the results suggest that larger seasonal variability is asso-

ciated with larger seasonal differences in GDP. While previous research has found that fluctuations of GDP between years have a negative effect on GDP (Ramey and Ramey, 1994), this possible mechanism has not been studied in the context of quarterly or seasonal fluctuations and seems to deserve the attention of future research. Given that future climate change is projected to change seasonal temperature differences, this points to yet another channel through which climate change will affect economic production in the future.

Chapter 3

Some like it cold: The persistent cost of higher temperatures in the economic sectors of Europe

Prior econometric analyses of the global impact of temperature fluctuations on aggregate economic performance (GDP) suggest that higher temperatures are costly to warm countries but potentially beneficial to cooler ones. However, aggregate temperature-GDP relationships reflect the net effect of temperature on productivity in the constituent sectors of the economy and across different spatial scales, potentially masking the heterogeneous sectoral and local effects that could inform efficient adaptation policies. Focusing on Europe, we use administrative district level data on the growth rate of Gross-Value Added (GVA) to estimate the impact of temperature fluctuations on GVA growth at the district level and for individual industry groups. Unlike previous studies with a global focus, for Europe we find persistently negative effects of warmer-than-average years on total GVA in relatively cold districts (annual mean temperatures < 13 degrees C). At an annual mean temperature of 11 degrees Celsius, one additional degree lowers the growth rate of GVA in the same year by -0.37 percentage points (SE = 0.2). Dis-aggregating by economic sector we find that the negative aggregate impact in cold districts stems from costs to agriculture, manufacturing, and mining and utilities. In relatively warm districts, the negative effect of higher annual mean temperatures on GVA in trade and other services but it is offset by positive effects in other sectors. Finally, we find that productivity impacts are also subject to spatial spillovers from neighbouring districts, reflecting how temperature effects dissipate through the regional economy.

3.1 Introduction

Economic analysis on the effect of annual mean temperature fluctuations on the economy suggests that high temperatures are negatively linked with GDP growth (Burke et al., 2015b; Dell et al., 2012; Kalkuhl and Wenz, 2020). These empirical studies point to a non-linear relationship between temperature and economic performance, whereby aggregate impacts on growth become more severe in regions with higher average temperatures. These aggregate effects are relatively similar across countries, and also reflect studies of certain individual economic behaviours and activities, such as crime or arable agriculture (Burke et al., 2015b; Hsiang, 2010). However, aggregate temperature-GDP relationships reflect the net effect of temperature on productivity in the constituent sectors of the economy and across different sub-national regions (Dingel et al., 2020), potentially masking the important heterogeneous sectoral and local level effects that could inform efficient adaptation policies. Just as economic stabilisation policies are tailored to the vulnerabilities of the different sectors of the economy (Guellec, 2009), policies geared towards adapting to, and decoupling economic activity from climate change should also be informed by the vulnerabilities of specific sectors in different regions (Bowen and Hepburn, 2014; Hepburn et al., 2020). Disentangling the pattern of economic costs, both spatially and across the economy, can also inform efforts to mitigate climate change by providing more accurate measures of economic damages. A more detailed understanding of the link between temperature and economic activity across different sectors and sub-national geographies: the economic geography of climate change, is therefore required (Cruz and Rossi-Hansberg, 2021).

We provide the first comprehensive analysis of the pattern of temperature-economy relationships for Europe using exhaustive national accounting data on economic productivity at the sectoral and sub-national level, coupled with granular weather data. More specifically, this study examines the effects of temperature fluctuations on economic growth across different constituent sectors of the economy at the European district level. Economic performance is measured using Gross Value Added (GVA): Gross Domestic Product (GDP at basic prices) minus intermediate consumption (inputs at producer prices) and is strongly correlated with GDP. Unlike GDP, GVA data are reported at the district and sectoral level in Europe, providing us with a unique spatial coverage of economic activity dis-aggregated by sector. Our

study overcomes several shortcomings of previous studies which have either focused on individual sectors or activities, with no unifying national accounting framework (Hsiang, 2010), or have been severely constrained in the temperature - economy relationships that could be estimated due to limitations to the climatic data (Colacito et al., 2019).

Our data on gross-value added comes from EUROSTAT and the OECD and is provided at the level of nuts-3 administrative districts in Europe (nuts-2 for Turkey) and the tl-2 level of the OECD. Importantly, this data are available in aggregate and by industry. For information on temperature and precipitation, we use high-resolution reanalysis data, which are based on the European Centre for Medium-Range Weather Forecasts (ECMWF) reanalysis and spatially refined, using the model COSMO. The data have a resolution of about 6 km (Bollmeyer et al., 2015) and we aggregate it to administrative districts using gridded population data from the Gridded Population of the World data set (Center For International Earth Science Information Network-CIESIN-Columbia University, 2018). Overall, our final dataset combines year-to-year fluctuations of annual mean and seasonal temperatures with district-level annual total GVA, and GVA in six exhaustive industry groups and covers over 1000 individual districts across 31 European countries between the years 1997 and 2020.

To identify the link between temperature and economic output, we use a Fixed Effects (FE) estimator which controls for unobservable characteristics at the district-level and for each year. The temperature-economy relationship for growth aggregate and industry-level GVA is estimated using a polynomial functional form on mean average temperature. We prefer growth as our outcome variable rather than GVA levels firstly since it is a more useful outcome variable for measuring changes in productivity, and secondly because differencing of the logarithm of GVA addresses concerns about non-stationarity of the data series, which would otherwise invalidate our estimates (Kalkuhl and Wenz, 2020; Burke and Emerick, 2016). To test for the persistence of temperature shocks on aggregate and industry-level GVA we also use lags of temperature and to test for seasonality of the temperature-economy effect we interact the temperature variables with indicator variables for the seasons. Furthermore, we test for adaptation, using a degree days model with polynomials within bins following Deryugina and Hsiang (2017). Finally, we use lagged dependent variables to check for dynamic growth effects, and spatial lags for temperature to evaluate whether regional

temperature changes are transmitted to a particular district.

3.2 Methods

3.2.1 Econometric framework

We estimate fixed-effects models with the growth rate of gross-value added (either total or for a specific industry; in this paper also referred to as economic output) as dependent variable. As independent variables we use polynomials of annual mean temperature, seasonal mean temperature, or degree-days. The temporal frequency of our variables is years. This means that we exploit exogenous fluctuations of temperature from year-to-year for the identification of causal effects of temperature on economic activity.

Our main model can be written as

$$\log(y_{i,t}) - \log(y_{i,t-1}) = \sum_{j=1}^k \beta_j (\bar{T}_{i,t})^j + \gamma_1 \bar{P}_{i,t} + \gamma_2 (\bar{P}_{i,t})^2 + \alpha_i + \theta_t + \epsilon_{i,t} \quad (3.1)$$

where observations are indexed by administrative districts i and years t , with gross value added y (aggregate or for a specific industry group), annual mean temperature \bar{T} , annual total precipitation \bar{P} , and district and time fixed effects α_i and θ_t respectively. We conduct robustness tests for which we also include district-specific linear time trends and country-year fixed effects.

If the effect of temperature lowered economic output persistently, we would expect to find significant effects of temperature on economic growth with the same sign also for lagged temperature. If temperature had only an instantaneous effect on output, we would expect to find effects of lagged temperature with opposite sign. To examine the persistence of temperature effects, we estimate distributed lags models (Dell et al., 2012):

$$\log(y_{i,t}) - \log(y_{i,t-1}) = \sum_{l=0}^s \sum_{j=1}^k \beta_j (\bar{T}_{i,t-l})^j + \gamma_1 \bar{P}_{i,t} + \gamma_2 (\bar{P}_{i,t})^2 + \alpha_i + \theta_t + \epsilon_{i,t} \quad (3.2)$$

with lags up to the order s . We then compute cumulative effects by summing over the coefficients of the coefficients of the individual lags. If temperature

lowered output in the same year but had no persistent effect on economic growth, we would expect a cumulative effect of zero.

To explore differences between seasons, we replace polynomials of annual mean temperature by an interaction between the climatic annual mean temperature and seasonal mean temperatures:

$$\begin{aligned}
\log(y_{i,t}) - \log(y_{i,t-1}) = & \sum_{j=1}^k \left[\beta_j^{\text{DJF}} \bar{\bar{T}}_i \left(\bar{T}_{i,t}^{\text{DJF}} \right)^j + \beta_j^{\text{MAM}} \bar{\bar{T}}_i \left(\bar{T}_{i,t}^{\text{MAM}} \right)^j \right] + \\
& + \sum_{j=1}^k \left[\beta_j^{\text{JJA}} \bar{\bar{T}}_i \left(\bar{T}_{i,t}^{\text{JJA}} \right)^j + \beta_j^{\text{SON}} \bar{\bar{T}}_i \left(\bar{T}_{i,t}^{\text{SON}} \right)^j \right] \\
& + \xi \bar{\bar{T}}_i + \psi^{\text{DJF}} \bar{T}_{i,t}^{\text{DJF}} + \psi^{\text{MAM}} \bar{T}_{i,t}^{\text{MAM}} + \psi^{\text{JJA}} \bar{T}_{i,t}^{\text{JJA}} + \psi^{\text{SON}} \bar{T}_{i,t}^{\text{SON}} + \\
& + \gamma_1 \bar{P}_{i,t} + \gamma_2 \left(\bar{P}_{i,t} \right)^2 + \alpha_i + \theta_t + \epsilon_{i,t} \tag{3.3}
\end{aligned}$$

where e.g. $\bar{T}_{i,t}^{\text{MAM}}$ is the seasonal mean temperature for the meteorological spring (March, April, May) in district i and year t and $\bar{\bar{T}}_i$ is the mean temperature (mean over all years) of district i .

To examine the effect of adaptation we follow Deryugina and Hsiang (2017) and estimate degree day models with polynomial terms for every bin:

$$\log(y_{i,t}) - \log(y_{i,t-1}) = \sum_b \sum_{j=1}^k \beta_{j,b} (n_{b,i,t})^j + \gamma_1 \bar{P}_{i,t} + \gamma_2 \left(\bar{P}_{i,t} \right)^2 + \alpha_i + \theta_t + \epsilon_{i,t} \tag{3.4}$$

where b indexes different bins of daily temperatures and $n_{b,i,t}$ is the number of days within bin b for district i in year t .

Spatial spillovers are examined with a model that includes a district i 's own temperature and in addition the average temperature of other districts that belong to the same region $r(i)$ or the same country $c(i)$, $\bar{T}_{r(i),t}$ or $\bar{T}_{c(i),t}$, which we also refer to as spatial lags. In mathematical terms, we estimate a model

$$\log(y_{i,t}) - \log(y_{i,t-1}) = \sum_{j=1}^k \beta_j \left(\bar{T}_{i,t} \right)^j + \sum_{j=1}^k \beta_j^r \left(\bar{T}_{r(i),t} \right)^j + \gamma_1 \bar{P}_{i,t} + \gamma_2 \left(\bar{P}_{i,t} \right)^2 + \alpha_i + \theta_t + \epsilon_{i,t} \tag{3.5}$$

To account for heteroskedasticity, serial autocorrelation, and spatial autocorrelation of the error term, we cluster standard errors at the level of countries. We also try clustering at the second administrative level and using Conley HAC standard errors, but since these tend to yield smaller standard errors we decide to choose the most conservative clustering method.

3.2.2 Data

We use data on gross-value added by industry from EUROSTAT and the OECD. The data are provided at the level of nuts-3 administrative districts in Europe (nuts-2 for Turkey) and the tl-2 level of the OECD. We complement these data with data on population from the same sources. The data use the NACE v2 industry classification with a breakdown of total GVA into up to 11 industry groups (Table C.11 in the Appendix). Not all countries report GVA for all of these groups. We further aggregate some of these industry groups. Our choice of aggregation is informed by the individual consideration of specific industries with a large share of economic activities occurring outdoors and well known influences of weather (agriculture, construction) and those with relatively large share of total GVA (manufacturing, industry, trade). The use of GVA means that our data differ conceptually from GDP data used in related previous studies (Burke et al., 2015b; Kalkuhl and Wenz, 2020) in that GVA excludes intermediate consumption.

For data on temperature and precipitation we use high-resolution reanalysis data. The data are based on reanalysis of the ECMWF, which was spatially refined using the model COSMO. The data have a resolution of about 6 km (Bollmeyer et al., 2015). We aggregate it to administrative districts using gridded population data from the Gridded Population of the World dataset in version 4 (Center For International Earth Science Information Network-CIESIN-Columbia University, 2018). To improve the balance of our final panel dataset, we drop spatial units with less than 5 observations in time. Descriptive statistics are provided in Table C.12 in the Appendix.

3.3 Results

3.3.1 Main findings

Our baseline estimates show that aggregate GVA growth has an approximately quadratic relationship with temperature, with positive effects on economic growth of very cold and very warm annual mean temperatures as evi-

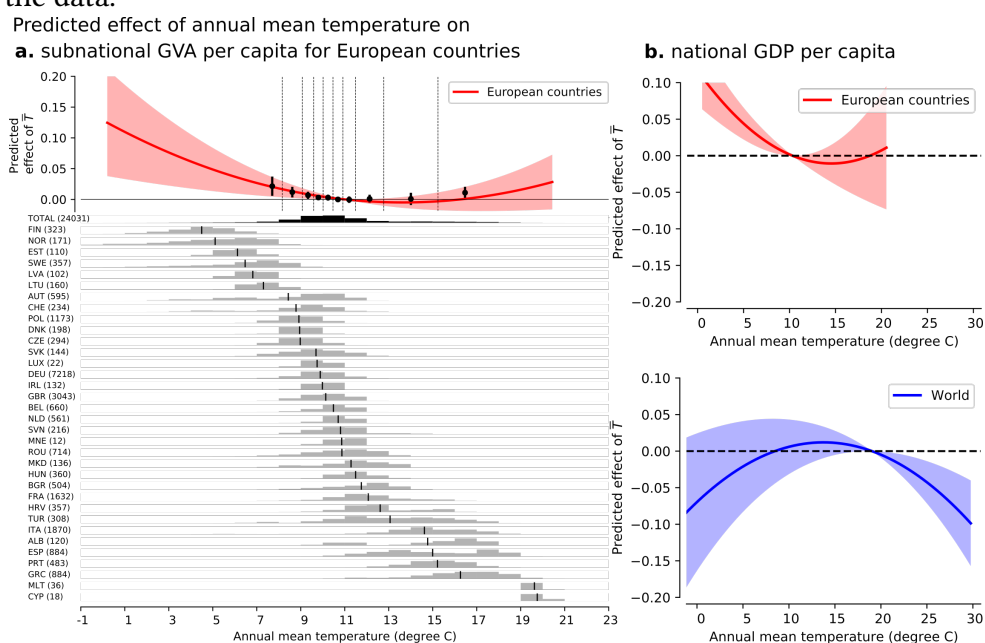
dent in Figure 3.1a, top. This overall convex relationship is constructed from marginal effects estimated at different levels of annual mean temperature, whereby different countries provide the support at different ranges of annual mean temperature. For example, support of average temperatures for districts in Finland ranges from about -1 to about 9 degrees Celsius, while in Greece the range is approximately from 10 to 20 degrees Celsius. Figure 3.1a, bottom, shows how this temperature variation underpins the estimation of the quadratic curve.

We also estimate a more flexible, binned model and present the results in Figure 3.1a top, in black. In this model, the two tails of the temperature distribution are examined independently of each other which is in contrast to the polynomial model which pools all regions and hence cannot assess whether it is the warm regions or the cold regions that are primarily responsible for the positive quadratic shape of the curve. The results of the binned model are reassuring in the sense that they support a positive parabola in both the warm and cold tail of regions in Europe.

For ease of interpretation and to illustrate statistical significance, the marginal effects of temperature change are depicted in Figure A.2, which plots the slope of the quadratic curve in Figure 3.1a. Overall, we find negative effects of warmer-than-average years on total economic output in cold districts (annual mean temperatures < 13 degrees C). In particular, Figure A.2a shows the instantaneous marginal effect as a dashed line, with the 95% confidence interval in blue. These marginal effects are significant and negative until the regional average temperature is about 13C. Figure A.2b depicts with a dashed line and green confidence interval the marginal effects of a model with lagged temperature variables for 6 periods. As evident from the figure, the cumulative negative effects are even more pronounced for cooler areas (< 13C), and GVA is reduced in several consecutive years after an initial temperature shock. More specifically, we find increasingly negative and significant cumulative effects up to six-time periods into the future. Beyond six time periods, the uncertainty of the point estimates substantially increases and we are unable to reject the null hypothesis of zero effect.

In terms of magnitudes, these estimates are substantial. For example, at an annual mean temperature of about 11 degrees Celsius, one additional degree lowers the growth rate of GVA in the same year by -0.37 percentage points (SE = 0.2) and by -1.8 percentage (SE = 0.5) points over six years. However,

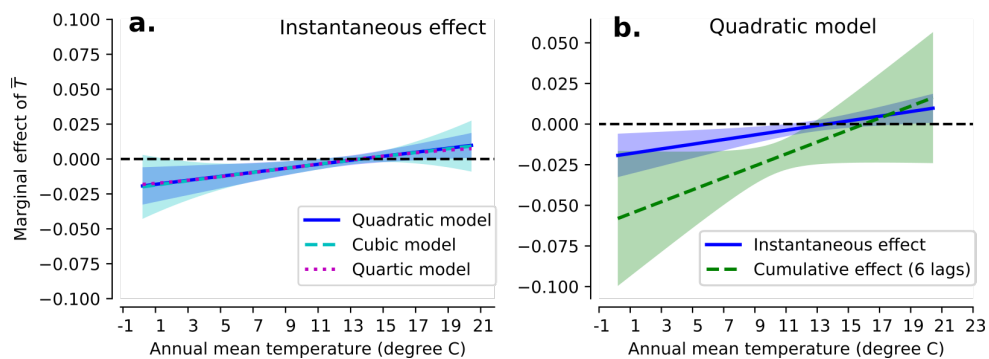
Figure 3.1: Main estimate (quadratic model) and the underlying variation in the data.



Notes: **a.** Effect of annual mean temperature on subnational GVA per capita for a sample of European countries (1997-2019). In brackets: number of observations. Dots indicate estimated coefficients of a model with dummies for bins, with bins based on deciles of the temperature distribution of the sample. Deciles are denoted by dashed vertical lines. **b.** Effect of annual mean temperature on national GDP per capita for a global sample of countries (1950-2019) and a sample of European countries (1997-2019). Results are qualitatively similar if the same time period is selected (Table C.31 in the Appendix).

there is little evidence for any effect of higher annual mean temperature on total economic output in warm districts (annual mean temperatures above 13 degrees Celsius). Importantly, these findings, that productivity is higher if regions remain cold, is in contrast to previous work using global samples of countries (Burke et al., 2015b) and sub-national regions (Kalkuhl and Wenz, 2020; Conte et al., 2020) which suggested that colder regions are likely to economically benefit from a warmer climate. Furthermore, we find little evidence for the previously reported negative effect of higher annual mean temperature on total economic output at annual mean temperatures above 13 degrees Celsius. As such, our analysis suggests that the appropriate geographical focus of policies designed to ameliorate the costs of climate change in Europe, which are partly based on such prior studies and this general notion, should focus more on colder regions.

Figure 3.2: Marginal effects of an increase in annual mean temperature obtained with alternative model specifications, obtained from the sample of European regions.



Notes: : **a.** Higher-order polynomials. **b.** Cumulative effect from a distributed lags model with six lags.

3.3.2 The impact of temperature change on GVA across regions and sectors

The above temperature-GDP relationship reflect the net effect of temperature on aggregate productivity but is silent regarding the important heterogeneous effects that could inform efficient climate policies. In this section, we unpack these aggregate level impacts on GDP spatially and by the component sectors of the economy. We start with the former, by combining the estimated relationship between annual mean temperature and the annual growth rate of GVA per capita with the distribution of annual mean temperatures in Europe. This heterogeneity is shown on a map of the pre-

dicted marginal effect by administrative districts in Figure 3.3. Overall positive marginal effects of higher-than-average temperatures are found in most of Southern Europe up to a latitude of about 45 degrees North. Further North, we find mostly negative marginal effects, especially further away from the Atlantic coast and at higher altitudes where annual mean temperatures tend to be lower (see Appendix for further information on temperatures by region in Figure C.11).

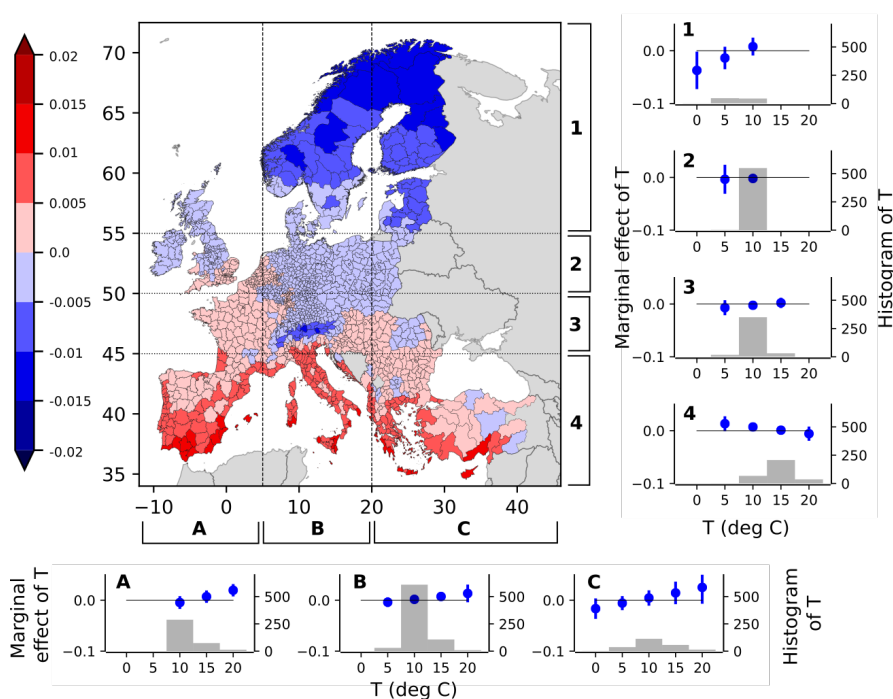
To further analyse the spatial heterogeneity in the temperature-economy relationship, we also estimate our main model using only observations that fall within certain bands of latitude and longitude. Reassuringly, we find similar responses to higher-than-average temperatures in all three bands of longitude, mirroring east to west the overall relationship in Figure A.2. Greater heterogeneity in the response to warmer-than-average years is found across different bands of latitude. For regions north of 55 degrees latitude, which includes Scotland, all Nordic countries, Latvia, and Estonia, we find large and significant negative marginal effects at low temperatures (around and below 5 degrees Celsius) and positive effects at higher temperatures. We find no significant marginal effects in moderate latitudes. Further to the South, specifically South of 45 degrees latitude, which includes Portugal, Spain, Italy, Bulgaria, Romania, Greece, and Turkey, regions exhibit small but significant positive marginal effects around and below 10 degree Celsius and negative effects at higher temperatures. Hence, the marginal effects that we find for regions in Southern Europe are more closely mirror the effects found for a global sample of countries.

We also analyse spatial spillover effects and find that temperature shocks to neighbouring regions lower economic output similarly to shocks to a region itself (Figure C.21 in the Appendix). We cannot disentangle whether the initial shock is transmitted along supply chains (Pankratz and Schiller, 2021), via demand shocks, or the capital and labour market, but it suggests that in sum negative spillover effects are larger than potentially positive spillover effects associated with inter-regional competition, and other changes in demand and supply. We note that these general and spatial equilibrium effects are relevant also from a methodological point of view, as in their presence country-by-year fixed effects will absorb some of the local effects of a temperature shock.

We continue our heterogeneity analysis by examining the temperature-GDP

Figure 3.3: Geographical distribution of estimated marginal effects of annual mean temperature across Europe.

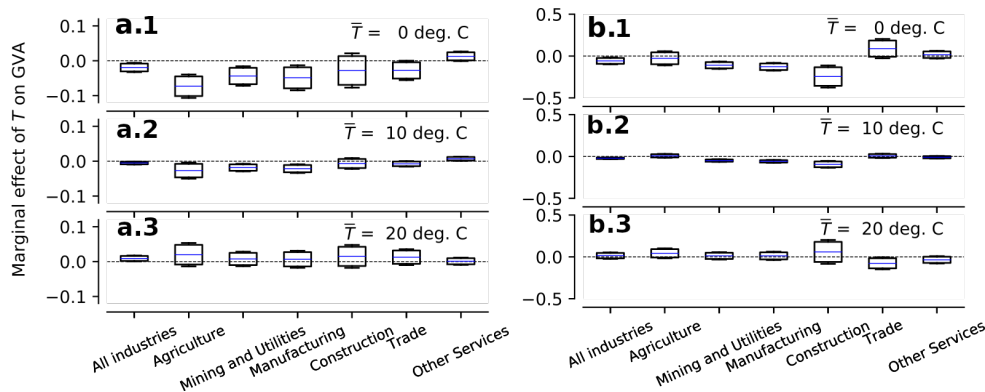
Marginal effect of annual mean temperature on GVA per capita



Notes: The map shows the marginal effect of a quadratic model fitted to the whole sample of European regions. The panels 1-4 and a-c show marginal effects of quadratic models fitted to subsamples, with histograms showing the frequency of temperature levels in these samples: 1-4. Estimated marginal effects for bands of latitudes. a-c. Estimated marginal effects for bands of longitudes.

relationships by sector. In particular, we study how yearly temperature fluctuation affects GVA growth in six broad industrial sectors that together make up total GVA: Agriculture, Construction, Manufacturing, Mining and Utilities, Trade and Other Services. Figure 3.4 shows the resulting decomposition of marginal effects. We find that for five industry groups: Agriculture, Construction, Manufacturing, Mining and Utilities, and Trade, warmer years have a negative instantaneous effect on GVA at low temperature levels (i.e. 0 degrees Celsius) (Figure 3.4a1). These negative marginal effects are smaller at moderate temperatures (10 degree Celsius) (Figure 3.4a2) and become insignificant at high temperatures (20 degree Celsius) (Figure 3.4a3). The cumulative effects over six years are qualitatively very similar, and tend to be larger in magnitude (Figure 3.4b). There are two exceptions where the instantaneous and the cumulative effects are qualitatively different: Agriculture, for which we find an insignificant cumulative effect at low temperatures (Figure 3.4b1) and a positive effect at high temperatures (Figure 3.4b3), and Trade, for which we find a positive cumulative effect at low temperatures (Figure 3.4b1) and a negative cumulative effect at high temperatures (Figure 3.4b3).

Figure 3.4: Estimated marginal effects by industry at three different levels of temperature, based on quadratic model fitted to the sample of European regions.



Notes: **a.** Instantaneous effect. **b.** Cumulative effect over six years. Estimates weighted by industry shares are shown in Figure C.22 in the Appendix.

While the nature of our data prevents us from causally identifying the underlying mechanisms for our heterogeneous findings, we posit a few possible explanations which can be broadly classified into two categories: physical and behavioural responses to warming. On the physical side, in cold regions of Europe, warmer temperatures in winter have been found to be detrimen-

tal to growth (e.g. crops) in agriculture (Van Passel et al., 2017). Furthermore, warmer-than-average temperatures have generally been found to negatively affect fisheries (Clark et al., 2020; Kelly et al., 2020). In construction and mining, we hypothesise that warmer-than-average temperatures around the freezing point detrimentally affect economic activity through increased costs of, e.g. pumping melt-water or subsidence (Hjort et al., 2018, 2022). Furthermore, we note that warmer-than-average temperatures can also cause damages to transport infrastructure in Northern regions (Gädeke et al., 2021). On the behavioural side, warmer-than-average temperatures in cold regions are likely to reduce the energy demand of the economy, thereby reducing revenues in Utilities. Furthermore, in industries with most work taking place indoors, such as Trade and Other Services, warmer-than-average temperatures in cold regions could affect working hours similarly to rainfall and very hot days by changing the opportunity costs of work versus leisure (Connolly, 2008; Graff Zivin and Neidell, 2014).

To gain some additional insights into mechanisms, we examine heterogeneity across seasons (Figure C.23 in the Appendix). We first focus on summer (JJA) and autumn (SON), the two warmest of the four seasons. For these seasons, cold regions appear to experience especially large drops in GVA from warmer-than-average temperatures for agriculture, mining and utilities, manufacturing, and construction (construction only in summer). These industries have previously been identified as those in which labour is relatively exposed to ambient temperature (Behrer and Park, 2019). In warm regions, the responses of the same industries are mostly insignificant, with possibly positive effects on GVA in agriculture (autumn) and construction (summer). These industry-level effects result in overall negative responses of GDP in cold and insignificant (but positive) responses in warm regions, generally consistent with evidence for some adaptation to days with high temperature, which we examine further below.

For winter (DJF) and spring (MAM), we find that GVA drops in response to warmer-than-average seasons in cold regions in mining and utilities, consistent with lower demand for heating and possibly higher maintenance costs of mining infrastructure as discussed above. Furthermore, we find negative responses for manufacturing. In warm regions, we find negative responses of GVA in Construction and Trade to higher-than-average temperatures in winter and positive responses to temperature in spring.

Some of these effects are likely to be more transitory than others. Persistent effects can be expected, for example, if consumption is actually reduced rather than delayed in response to a weather shock, if a recovery in response to a shock is hindered by rigid capacity constraints, if there are damages to the capital stock, or if investment is negatively affected. In construction, for example, the relatively large and persistent negative effect of warmer-than-average temperatures in cold regions could be explained with physical damages to capital and capacity constraints. In contrast, in Trade the negative instantaneous but positive cumulative effect could be due to a delay of consumption or an inter-temporal re-allocation of work and leisure (Connolly, 2008; Graff Zivin and Neidell, 2014), both of which reduce economic production in a transitory manner. These dynamic effects on Trade could also reflect Trade as a medium-term response to temperature induced supply constraints. Since such mechanisms change the pattern of economic growth, we think that it is worthy of further investigation using micro-data.

Finally, we look at adaptation. Specifically, we interpret adaptation as non-linear effects of the number of days in specific bins of daily mean temperature (degree-days). For example, in the presence of adaptation we expect that very hot days (≥ 30 degrees C) have a more negative effect on GVA the more rarely they tend to be observed in a specific region. Our results suggest some adaptation to days with very low temperatures (< -5 degrees C) and to very hot days (≥ 30 degrees C) (Figure C.24a in the Appendix). For these days, the estimated effect exhibits a positive curvature, suggesting that the marginal effect of degree-days becomes more and more positive the more of these days are observed. For intermediate temperatures (0-25 degrees C) we find no clear evidence for adaptation.

3.3.3 Robustness

The choice of a second-order polynomial in our main analysis is informed by the relationship found by Burke et al. (2015b) and Kalkuhl and Wenz (2020). This quadratic specification is also supported by results obtained with a flexible model using bins of annual mean temperature (black dots in Figure 3.1). Furthermore, as robustness check, additional models with higher-order polynomials are estimated, which yield very similar results (Figure A.2a). We also validate our results by estimating our empirical model on datasets from previous studies and reproducing their results using the period used in this study (Figure 3.1b, bottom) and using GDP data only for Europe (Figure 3.1b, top).

The results clearly illustrate that differences between our results and previous results obtained with a global sample of countries (Burke et al., 2015b) arise from our geographical focus on Europe, rather than the period under scrutiny or the use of GVA data instead of GDP data (see also Table C.31 in the Appendix).

The results are also robust to clustered standard errors at the country rather than district level, Conley HAC standard errors (Table C.32 in the Appendix) and the inclusion of time trends. Somewhat consistent with our results above that suggest relatively large spatial spill-over effects, we find that including country-by-year fixed effects makes (which absorb the spatial effects) our estimates insignificant (Table C.32 in the Appendix, Columns 3 and 4). Moreover, regarding our analysis of adaptation based on degree-day models we find qualitatively very similar results for models with cubic polynomials suggesting that the results are robust to alternative choices of model specification (Figure C.24b in the Appendix). Finally, the results are very similar for alternative weather data (Figure C.31 in the Appendix).

3.4 Conclusion

The empirical evidence on the global relationship between economic activity and temperature variation typically suggests that positive temperature shocks in warmer than average locations are costly, but the same shocks on colder than average areas (mean temperature below 13 Celsius) temperature can be economically advantageous (Burke et al., 2015b). Such findings underpin a common narrative on the expected costs of climate change and their spatial distribution, and inform national and regional adaptation strategies. The EU strategy on the adaptation to climate change (Commission, 2021) largely reflects this narrative, using examples from climate change ‘hotspot’ areas in Southern Europe and the Mediterranean as exemplars of the geographical incidence of climate damages and the need for strategic interventions on climate change adaptation. The results of this paper turn this narrative on its head in the European context, and show that positive temperature shocks are more costly to cold regions of Europe, such as the Nordic countries, Alpine regions and Scotland, than to the Southern and Mediterranean regions.

These findings are of general interest, and show the importance of detailed spatial and sectoral analysis, and a regional focus, for estimating hetero-

geneous temperature-economy relationships. The findings can inform the broader efforts to assess the Social Cost of Carbon and the desirability of meeting the 1.5C Paris target using Integrated Assessment Models, via updated estimates of damage functions (Hänsel et al., 2020). In this regard, our results on regional and sectoral heterogeneity also make the unequal geographical distribution of the costs of climate change explicit. More specifically, the results can inform the EU strategy to forge a climate resilient Europe by 2050 (Commission, 2021). We provide an up to date picture of the potential socio-economic impacts of climate change, fulfilling the EU strategy's desire to 'anchor smarter adaptation in the latest science'. Our illustration of the spatial pattern of the temperature-economy nexus could also inform EU wide adaptation policies, such as the coordination of genetic diversity sharing in agriculture, or the coordination of regional programs of investment in, e.g., adaptation innovation. The finding that the costs of temperature shocks are higher in colder regions can also help to prioritise scarce regional development, infrastructure and adaptation funds to more vulnerable locations and sectors where the benefits will be higher. More broadly, our spatially explicit results help to identify synergies and trade-offs between adaptation to climate change and the objectives of regional development in Europe. Our evidence on decentralised adaptation to temperature extremes can also inform the extent to which intervention is required at all. Ultimately, the higher costs of temperature shocks in cold regions, where warming associated with climate change is expected to be greatest, reflects a socio-economic specialisation, both behavioural and structural, around their colder climates. Contrary to previous findings, from the perspective of income, in colder regions some like it cold.

Future research can further illuminate the costs of climate change in cold regions of the world outside Europe. This will add more evidence on the mechanisms through which warmer-than-average years can reduce economic production in certain industries. Other research questions arising from our results concern the possibilities to adapt to gradual changes of temperature in the long-run, which might be underestimated or overestimated in our results. We hope that as more and more granular data have become available, new insights into the challenges of climate change can be gained from geographically explicit analysis as in this paper, ultimately leading to a better scientific understanding and supporting better policies for more effective adaptation.

Chapter 4

Policy sequencing towards carbon pricing: Empirical evidence from G20 economies and other major emitters

Carbon pricing is considered the most efficient policy to reduce greenhouse gas emissions but it has also been conjectured that other policies need to be implemented first to remove certain economic and political barriers to stringent climate policy. Here, we examine empirical evidence on the the sequence of policy adoption and climate policy portfolios of G20 economies and other major emitters that eventually implemented a national carbon price. We find that all countries adopted carbon pricing late in their instrument sequence after the adoption of (almost) all other instrument types. Furthermore, we find that countries that adopted carbon pricing in a given year had significantly larger climate policy portfolios than those that did not. In the last part of the paper, we examine heterogeneity among countries that eventually adopted a carbon price. We find large variation in the size of policy portfolios of adopters of carbon pricing, with more recent adopters appearing to have introduced carbon pricing with smaller portfolios. Furthermore, countries that adopted carbon pricing with larger policy portfolios tended to implement a higher carbon price. Overall, our results thus suggest that policy sequencing played an important role in climate policy, specifically the adoption of carbon pricing, over the last 20 years.

4.1 Introduction

Carbon pricing has been suggested as an economically efficient instrument to reduce greenhouse gas emissions and a growing body of literature confirms its effectiveness¹ but many countries including some of the world's largest emitters appear reluctant to implement it. By the end of 2020, only about 40 countries had implemented either a carbon tax or a national emission trading scheme, leaving about 150 countries and about 85 percent of global greenhouse gas emissions without an explicit price on carbon ([World Bank], 2022). This slow progress in the adoption of carbon pricing has been associated with several barriers including concerns about high energy prices hurting households and reducing the industrial competitiveness of the economy (Klenert et al., 2018; Dolphin et al., 2019; Levi et al., 2020). Does the presence of these barriers suggest that countries need to give up on the idea of carbon pricing and instead resort to more feasible, yet second-best climate policies? A more optimistic view, based on a careful reading of the experiences of Germany and California, suggests that other climate policies can be used to lower or remove some or all of the barriers, thus paving the way for a subsequent adoption of carbon pricing (Meckling et al., 2015, 2017; Pahle et al., 2018).

This idea of policy sequencing is one theory that can explain the observed diversity and combinations of policy instruments, in addition to the simultaneous presence of several market failures (Bertram et al., 2015; Bataille et al., 2018; Stiglitz, 2019) and theories of second-best substitutes of first-best policies (Benjamin and Stavins, 2007; Fischer et al., 2021). The idea of policy sequencing suggests that climate policies can be used iteratively to address specific barriers to higher stringency. To this aim, government policies can address specific market failures, such as public good properties and asymmetric information. For example, positive externalities of green technological innovation can be addressed with public funding for research and development. This innovation can result in affordable alternatives to high-carbon goods and services lowering the impact of carbon pricing on household expenses, thereby making it both more effective and more acceptable. Likewise, asymmetric information is commonly addressed through support for education and labelling, which can in turn increase the market for less

¹For example, Andersson (2019) examines the effectiveness of carbon pricing in Sweden; for a recent review, see Green (2021).

emission-intensive products and lead to positive returns to scale. Political opposition to stringent climate policies due to lobbying of powerful industry groups can be addressed, for example, through grants and subsidies that support the growth of a green sector, broadening the support base for additional policies. Furthermore, standards on environmental performance can provide long-term orientation and help coordinating private investments into the development of new innovative green technologies.

In this study, we present empirical evidence on sequencing of climate policies focusing on G20 economies and other large emitters that had adopted either a carbon tax or a national emission trading system (ETS) by the end of 2020. To our knowledge, we are the first to provide quantitative and international empirical evidence on how countries have built up their climate policy portfolios over time before eventually adopting a carbon tax or permit system. To this aim we combine a comprehensive dataset on carbon pricing (World Bank Carbon Pricing Dashboard) with a large international dataset on climate policies (den Elzen et al., 2019; Roelfsema et al., 2020; Fekete et al., 2021). For the purpose of our analysis, we aggregate 72 instrument categories to eight different instrument types and distinguish between six sectors. We then derive policy sequences based on pairwise conditional empirical frequencies. Furthermore, we use matching and linear regression to identify significant statistical associations between climate policy portfolios and the adoption and stringency of carbon pricing policies.

We find similar sequences of policy instruments across sectors and countries among countries that have adopted carbon pricing. Carbon pricing tends to be implemented last, after the adoption of all (or almost all) other instrument types. Examining the temporal evolution of countries' climate policy portfolios, we find that countries that adopted carbon pricing in a specific year tended to have larger policy portfolios than other countries. Furthermore, examining individual countries' policy portfolios in greater detail, we find large variation in the overall size of countries' policy portfolios at the time of adoption of carbon pricing, with possibly smaller portfolios among more recent adopters. We discuss several explanations for this finding including variation in institutional capacity but also increasing public support for climate policy, decreasing abatement costs, and policy diffusion between countries. Furthermore, our results suggest that countries with larger policy portfolios tended to implement carbon pricing policies with higher average carbon prices, consistent with the idea that earlier policies remove barriers

to higher stringency.

Our analysis contributes to the debate about an optimal climate policy mix, including more normative work on the benefits of alternative instrument types (Peñasco et al., 2021), policy mixes (Bertram et al., 2015; Fankhauser et al., 2010; van den Bergh et al., 2021), and second-best policies (Fischer et al., 2021). Our paper adds to this debate another layer of complexity, the temporal sequence of policy adoption. In principle, policies that might be considered second-best for a specific market failure, such as the negative externalities from GHG emissions, can also be considered as temporary remedies that facilitate a later adoption of the first-best policy (Pahle et al., 2018). This idea is generally consistent with the empirical evidence on the temporal sequence of policy adoption that we report in this paper. Indeed, our results suggest that earlier policies do not only facilitate the adoption of carbon pricing, but that they are also positively associated with its stringency, pointing to additional benefits of policy sequencing.

We also contribute to a growing debate about the determinants of political support for carbon pricing (for relatively recent empirical work see e.g. Anderson et al. (2021); Douenne and Fabre (2022); Mildemberger et al. (2022)). Our paper takes a macroscopic perspective and focuses on relatively long time scales over which early policies such as technology subsidies might in the past have slowly increased support for more stringent climate policies as they helped to transform the energy system, to reduce the emission-intensity of the economy, and to build pro-environmental interest groups. Indeed, our results suggest that it took countries on average between 5 and 18 years to move from other policies to carbon pricing. At the same, by reporting evidence consistent with the existence of substantial barriers to carbon pricing, we provide additional support for attempts to examine how the design of pricing policies can be used to increase political support and facilitate implementation (see, for example, Baranzini and Carattini (2017); Bechtel et al. (2020); Carattini et al. (2018); Klenert et al. (2018); Kotchen et al. (2017)).

The paper is structured as follows. In Section 5.2, we explain our empirical framework, describe the dataset, and explain the statistical methods of the analysis. Results are presented in several steps in Section 5.3. Finally we discuss our main findings and conclude in Section 4.4.

4.2 Methods

4.2.1 Empirical framework

Building on prior work on climate policy sequencing we expect there to be barriers to the adoption of carbon pricing. Possible barriers include concerns about high mitigation costs for firms in energy-intensive industries who compete internationally, high costs for consumers and possibly a regressive distribution of these costs, political opposition to climate change policies, opposition to pricing policies, and concerns about the effect of a carbon price on employment in industries that rely on fossil fuels. Some of these barriers can generally be addressed with climate policies other than carbon pricing, such as technology subsidies, which can lower mitigation costs through technological innovation or more generally remove opposition to a pricing policy through a gradual transformation of the economy away from carbon-intensive activities.

We therefore expect there to be a positive statistical association between the number of climate policies other than carbon pricing, in the following referred to as the size of the climate policy portfolio, and the adoption and intensity of a pricing policy. This positive association can generally result from an ex-post effect, whereby earlier climate policies increase the probability of the subsequent adoption of a pricing policy, or from an ex-ante effect, whereby an anticipated later adoption of a pricing policy motivates the prior adoption of other policies (Figure 4.1). The latter direction of causation is especially plausible if high mitigation costs are a major concern, as those can reliably be reduced with other climate policies implemented prior to the planned adoption of carbon pricing, albeit at possibly higher costs in terms of welfare. The larger those barriers to adoption, the stronger we expect the statistical association between the size of the policy portfolio and the adoption of the pricing policy to be.

In the first part of the analysis, we examine the sequence of policies to identify the temporal order in which climate policies with different instruments types tend to be adopted. We then briefly describe the temporal evolution of climate policy portfolios. After this more descriptive analysis, we use a matching methodology to establish whether countries that adopted carbon pricing in a given year had larger policy portfolios than those that did not. We find that adopters indeed had larger portfolios. To address concerns




Figure 4.1: **Causal diagram.** The arrows represent possible causal relationships between the adoption of a carbon pricing policies, (unobserved) barriers to carbon pricing, and the climate policy portfolio. The ex-post channel illustrate how early climate policies can incrementally reduce barriers to carbon pricing, eventually allowing for its adoption without this last step being the strategic objective at the time of adoption of the earlier policies. The ex-ante channel illustrates how early climate policies can alternatively result from a strategic objective of implementing carbon pricing after the removal of certain barriers.

about possibly confounding country characteristics (Figure D.11 in the Appendix) we use a linear regression in which we include a few such characteristics. Furthermore, we use the regression analysis to examine the association between the size of the policy portfolio and the level of the pricing policy at the time of implementation.

4.2.2 Data

We use data on climate policies from the website climatepolicydatabase.org, which provides to our knowledge the most comprehensive international dataset on climate policies. The dataset has been gradually composed over the recent years (Nascimento et al., 2021) and been analysed in a number of academic publications (den Elzen et al., 2019; Roelfsema et al., 2020; Fekete et al., 2021; Yao and Zhao, 2022). The dataset is based on other international datasets, reports, and country specific documents, and incorporates a variety of other popular datasets on climate (or in some cases more broadly environmental) policies such as the Climate Change Laws of the World (Eskander et al., 2020) and OECD policy instruments database².

Despite the variety of sources used for the construction of the dataset it can

²<https://www.oecd.org/env/indicators-modelling-outlooks/policy-instrument-database/>

generally not be expected to include all climate policies of every country. Information on data comprehensiveness for individual countries was obtained from the NewClimate Institute. The dataset can generally be considered comprehensive for G20 economies (including EU member countries that are individual members of the G20, but not other EU members) and 18 additional countries to which we loosely refer as other major emitters. These additional countries are mostly advanced and emerging economies in Europe, Asia and Latin-America but also encompass some less developed countries and two countries in Africa (Figure 4.2). For all these countries, climate policies have been collected with the aim of completeness and the dataset has gone through a validation with national stakeholders and experts. For all other countries, the data can generally not be considered comprehensive. We hence drop all those countries from our sample. This includes some of the early pioneers of carbon pricing (Norway, Sweden, Finland, Poland, Denmark) (see Tables D.51 and D.52 in the Appendix for a detailed list of countries).

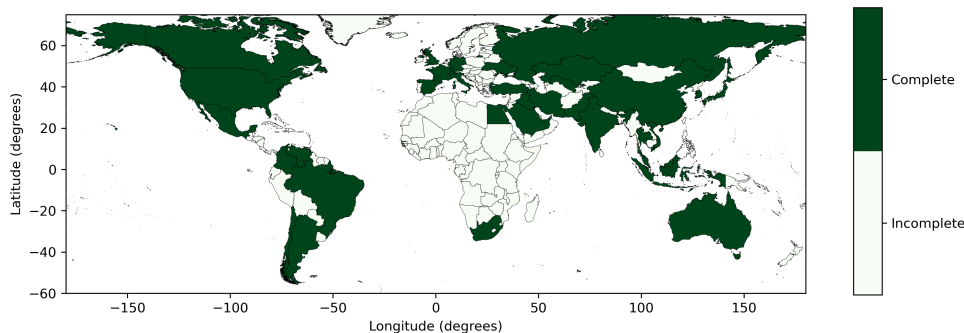


Figure 4.2: **Sample of countries included in the analysis.** Map shows whether the data on climate policies can be considered comprehensive. See also Table D.51 in the Appendix.

Every policy in the dataset carries information on policy objectives, administrative level, instrument types, targeted sectors, and more. To prepare the data for our analysis, we focus only on policies that have climate change mitigation as one of their objectives. Furthermore, we neglect any policies at the subnational level and apply all EU policies to the member countries' portfolio. If a country became member after the policy was decided in the EU, we use the year of joining the EU as the date of policy adoption.

The dataset distinguishes 72 instrument categories, which we aggregate to seven different instrument types based on the instrument typology of the

IEA. Furthermore, we distinguish between five sectors based on the sector definition of the IPCC AR5 WGIII (Electricity and heat production; Transport; Buildings; Industry; Agriculture, forestry, and other land use) and one additional sector General for policies that do not target specific sectors. We combine the final climate policy dataset with data from the Carbon Pricing Dashboard of the World Bank to consistently distinguish an additional instrument type carbon pricing, which can otherwise also be considered a subtype of financial incentives. The final eight instrument types are: Regulatory instruments; Grants, subsidies, and other financial incentives; Information and education; Policy support; Research, development, and deployment; Voluntary agreements; Procurement and investment; and Carbon pricing. A list of the corresponding instrument categories together with their frequency in the dataset is shown in Table D.71 in the Appendix.

We find that every instrument type has been used in every sector in at least one country (Figure D.61 in the Appendix). The number of times we observe a specific instrument-sector combination in our database ranges from 18 (Research, development and deployment in Agriculture, forestry, and other land use) to 1902 (Policy support introduced without targeting a specific sector).³ The latter pattern includes national climate change strategies and emission reduction targets including the NDC. Other frequent combinations are also well known from the climate change mitigation policy research and practice. This includes a frequent use of financial incentives in the energy sector (e.g. feed-in-tariffs, emission permits), and frequent use of regulatory instruments in the buildings and transport sector (e.g. efficiency standards for household appliances, energy efficiency standards for buildings, and emission standards for road transport vehicles).

We complement the climate policy data with country characteristics that we obtain from several sources. This includes GDP per capita data in purchasing power parity from the World Bank, an index of education from the Human Development Indicators provided by the United Nations Development Program, an index of the control of corruption from the World Governance Indicators of the World Bank, and information on fossil fuel reserves from the US Energy Information Administration. Descriptive statistics are provided in Table D.21 in the Appendix.

³In total, we observe 14,540 instrument-sector combinations.

4.2.3 Statistical methods

In the first part of the analysis, we identify policy sequences in terms of their instrument type. The eight instrument types result in 40320 possible sequences. To identify these sequences, we first consider all possible pairs of instrument types. For each of these 28 pairs, we examine which of the two instrument types tends to be adopted first across sectors and countries. We then use the relative timing of these pairs to construct the overall sequence.

Formally, we consider the adoption of two instrument types as events X and Y respectively. Using this terminology, we examine the conditional frequency that event X is preceded by event Y across countries and sectors. In mathematical terms, we examine the conditional frequency $f(Y_{t-1}|X_t)$ whereby X_t and Y_{t-1} are binary variables indicating whether the two policies have been decided up to the year t and $t - 1$ respectively:

$$f(Y_{t-1}|X_t) = \frac{n(Y_{t-1} \wedge X_t)}{n(X_t)} \quad (4.1)$$

with the number of times an event is observed in the data denoted as $n(\cdot)$. We then derive the relative order of all possible pairs of instrument types by comparing $f(Y_{t-1}|X_t)$ and $f(X_{t-1}|Y_t)$. Because we are interested in existing policies at the time of decision of a new policy, we exclude all observations after an event is observed for the first time (i.e. after the first time a specific instrument is adopted in a specific sector in a specific country).

The data used for the identification of policy sequences is illustrated in Figure 4.3. For example, we find that in the USA regulatory instruments X precede voluntary approaches Y in three out of six sectors. In the remaining three sectors, both instrument types are implemented for the first time in the same year. This yields $f(X_{t-1}|Y_t) = 0.5 > 0 = f(Y_{t-1}|X_t)$. For the USA, we hence consider regulatory instruments as preceding voluntary approaches.

In the second part of the paper, we examine differences in the adoption of carbon pricing and in policy sequencing across countries. To this aim, we quantify the size of countries' climate policy portfolios. We do so by counting how many of the eight instrument types have already been implemented in the six sectors mentioned above. This yields a score between 0 and 48 for every country and every year. We complement this information with a range of control variables: GDP per capita, education, control of corruption,




Figure 4.3: **Adoption of policies with different instrument types and sectors over time in different countries.** Shown are only policies that are the first of their kind in terms of their country, instrument type, and sector combination. The figure illustrates all information used for the derivation of policy sequences. See text for explanation and an example.

political globalisation, and the prevalence of fossil fuels in a country. The choice of explanatory variables is informed by the results of comprehensive international analysis of the factors that determine the adoption of pricing policies (Dolphin et al., 2019; Best and Zhang, 2020; Levi et al., 2020).

We first examine the statistical association between the size of countries' climate policy portfolios and whether countries adopted carbon pricing in a given year. To this aim, we compare the climate policy portfolios of countries that adopted a national carbon price in a given year with the portfolios of countries that did not adopt a carbon price neither earlier nor in the same year. We refer to the first group of countries as treated countries and to the second group as control countries. For the statistical analysis, we match every treated country with one control country. To this aim, we assign every treated country a randomly chosen control country. We iterate this random assignment 1000 times and then compare the average size of the policy portfolios of treated countries with the average size of portfolios of the control countries. For inference, we calculate bootstrapped confidence intervals. In addition, we estimate a logit model with the adoption of carbon pricing as binary dependent variable, which allows us to include certain country characteristics as control variables.

We next focus on heterogeneity among countries that eventually adopted a

national carbon price (treated countries). We first examine the size of policy portfolios at the time of adoption of carbon pricing. To this aim, we estimate a linear regression model with the size of the policy portfolio at the time of adoption as dependent variable and the same control variables as above. To examine trends over time we also include the year of adoption as an explanatory variable. Because the number of observations is small relatively to the number of explanatory variables, we also estimate a more parsimonious model with only selected explanatory variables. For this model we choose GDP per capita and the reserves of fossil fuels. Because reserves of oil and gas are highly correlated, we include only reserves of coal and reserves of oil in this model. Furthermore, we use Lasso model selection to identify the most important explanatory variables. Lasso estimation optimises a model that strikes a balance between the explained variation and model complexity as measured by the number of explanatory variables. As a popular method for model shrinkage, it is particularly suitable for the detection of influential variables among several correlated variables. In addition, we also examine the association between the size of policy portfolios at the time of implementation of a pricing policy and the economy-wide average carbon price. To this aim, we estimate a similar linear regression model with the average carbon price as dependent variable.

4.3 Results

4.3.1 The temporal sequence of climate policies

We focus on policy sequences of countries that eventually adopted a carbon pricing policy. Overall we find similar sequences of policy instruments across sectors and countries (Figure 4.4). This is especially true for the relative position of carbon pricing. Pooling policy adoption in all countries and sectors, carbon pricing tends to be the last instrument type. If we pool policies only across countries but keep sectors separate, we find that carbon pricing is the last instrument type in every sector. Furthermore, in 12 out of the 15 countries, carbon pricing tended to be used for the first time in a specific sector after the use of all other instrument types. Overall, we hence find that carbon pricing tends to be implemented last, after the adoption of all or almost all the other seven instrument types.

The results also reveal some recurrent patterns of sequencing for the other seven instrument types. Focusing on the results by sector, we find two

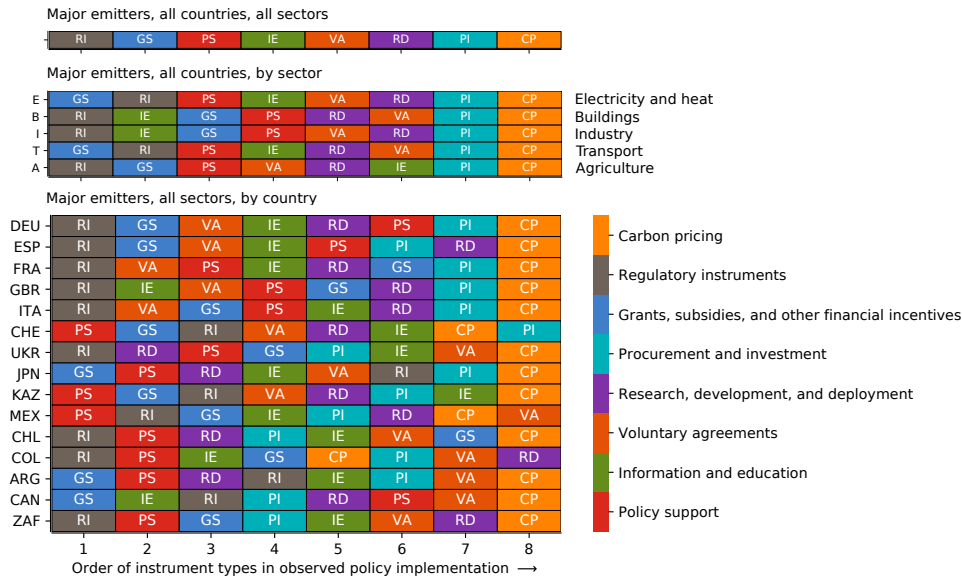


Figure 4.4: **Policy sequences of countries with a carbon price.** Sampling criteria is adoption of carbon price by the end of 2020 and availability of comprehensive data on climate policy adoption. See Figure 4.2 and Tables D.51 and D.52 in the Appendix for further information on sampling.

groups of instrument types. The first group consists of four early instrument types: Regulatory instruments, Grants, subsidies, and other financial incentives, Information and education, and Policy support. The second group of four late instrument types includes Research, development, and deployment, Voluntary agreements, Procurement and investment, and Carbon pricing. With the exception of voluntary agreements and financial incentives in Agriculture, policies of the first group tend to be implemented before policies of the second group in all sectors (Figure 4.4).

We find more variation of sequencing at the level of individual countries (Figure 4.4). The most frequent patterns are a relatively early adoption of Regulatory instruments and Policy support instruments and a relatively late adoption of Procurement and investment and Carbon pricing. The relative positions of the remaining four instrument types (Grants, subsidies, and other financial incentives; Research, development, and deployment; Voluntary agreements; Information and education) show greater variation across countries.

4.3.2 The build-up of climate policy portfolios over time

The results presented in the previous Section suggest that carbon pricing tends to be adopted after the adoption of climate policies with all or almost all other instrument types. We use this insight as motivation to examine whether the climate policy portfolios of countries that adopted carbon pricing in a specific year systematically differ from the portfolios of countries that did not adopt it.

To do so, we quantify the size of countries' policy portfolios as the number of instrument type-sector combinations that a country has already used prior to a given year. The temporal evolution of the portfolios of countries that eventually adopted carbon pricing is shown in Figure 4.5a. The visualisation reveals some interesting patterns. There appear to be at least three different kinds of trajectory of how countries built up their policy portfolios over time. Countries of the first group (blue colors in Figure 4.5a), including Canada, Japan, and South Africa, exhibit a relatively rapid expansion of their portfolio followed by a slow further expansion over several years that eventually includes the adoption of carbon pricing. Countries of the second group, including Argentina and Switzerland (green colors in Figure 4.5a), show a steady gradual expansion of their portfolios up until the introduction of carbon pricing. Countries of the third group (red colors in Figure 4.5a), including the current EU members in the sample, show a rapid expansion of policies almost immediately followed by the introduction of carbon pricing.




Figure 4.5: **Development of countries' climate policy portfolios over time.** Shown is the number of instrument type-sector combinations used in countries' policy portfolios. Colors on the left correspond to groups of countries with similar trajectories; see text for explanation.

This diversity of trajectories also means that the average time between the adoption of new instrument type-sector combinations and the adoption of

carbon pricing systematically differs in the sample. For Canada, Japan, and South Africa, this average time is about 18, 13, and 14 years respectively. For Argentina and Switzerland, the corresponding values are 11 and 10 years respectively. For the current EU countries, the average time is about 5 years.

4.3.3 Policy portfolios of adopters versus non-adopters of carbon pricing

We next examine to what extent the size of countries' policy portfolios can be considered a good predictor of the adoption of carbon pricing. To do so, we first compare the policy portfolios of countries that had adopted a carbon price by the end of 2020 with those that had not. At the time countries of the first group adopted carbon pricing, they had used on average 29.5 instrument type-sector combinations (Figure 4.5a). We then contrast this number with the size of portfolios of countries without a national carbon price. By the end of 2020 those countries had used policies with on average about 23.9 instrument type-sector combinations (Figure 4.5b). In 2015, which is the average year of adoption of carbon pricing, they had used on average about 19.7 combinations.

To compare countries that adopted carbon pricing in a given year with those that did not more systematically, we match countries based on a random assignment. This has the advantage that we also consider countries as possible control countries that had not adopted carbon pricing in year t but adopted it in year t' with $t < t' < 2020$. For inference, we calculate bootstrapped confidence intervals (Section 5.2). We find that countries that adopted a carbon price in a given year had policy portfolios that were on average about 11.12 instrument type-sector combinations larger than the portfolios of those that did not. This difference is significant at a confidence level of $\alpha = 0.05$. Overall, countries that adopted carbon pricing in a given year hence tended to have significantly larger climate policy portfolios than those that did not.

As a robustness test, we also estimate a logit model with the adoption of carbon pricing as binary dependent variable and different sets of explanatory variables. As for the matching, we find a statistically significant positive association between the size of the policy portfolio and the adoption of carbon pricing (Table D.31 in the Appendix). Results of a Lasso estimation suggest that the size of portfolio and the prevalence of gas reserves are the strongest determinants of the adoption of carbon pricing. As we include additional

control variables, the estimated coefficient of the size of the portfolio becomes smaller and less significant. This suggests that at least part of the positive association is due to country characteristics that influence both the size of the policy portfolio and the adoption of carbon pricing.

4.3.4 Heterogeneity among adopters of carbon pricing

The results above suggest that the size of climate policy portfolios is positively associated with the probability of adopting carbon pricing in a given country in a given year. One possible explanation is that policies other than carbon pricing allow countries to remove barriers to a subsequent adoption of carbon pricing. To further illuminate this explanation we attempt to explain differences in the size of policy portfolios at the time of adoption of carbon pricing with a linear regression model with variables possibly influencing some of these barriers or acting as confounders in the analysis. That is, we estimate a model with the size of policy portfolio at the time of adoption as dependent variable and GDP per capita, education, control of corruption, and the prevalence of fossil fuels in a country as explanatory variables. To examine trends over time, we also include the year of adoption in the model. Because of the high number of variables relative to the number of observations, we also estimate a reduced model and use a Lasso model for variable selection.

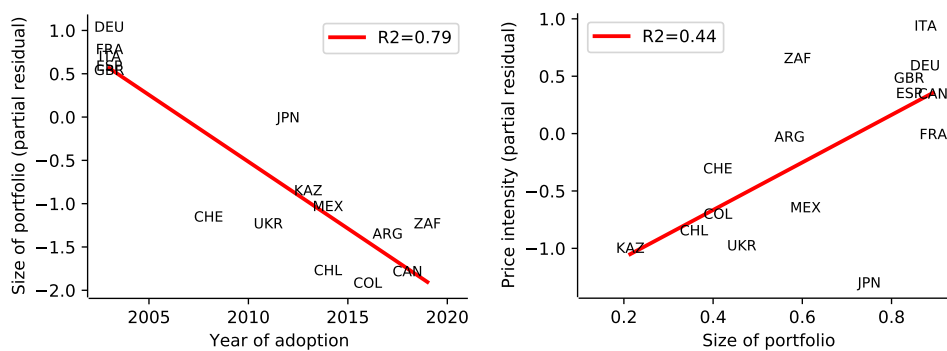


Figure 4.6: **Scatterplots of statistical associations between the year of the adoption of carbon pricing, the size of climate policy portfolios in that year, and the average carbon price in the first year of implementation;** to control for several variables including GDP per capita and reserves of fossil fuels, the figure shows partial residuals of the model in Column 2 (left) and Column 6 (right) in Table D.32 in the Appendix.

We use this model to examine certain patterns in the data. Specifically, we find that after controlling for country characteristics the size of policy portfo-

lios at the time of adopting a carbon price has decreased over the last 20 years (Figure 4.6 left). For example, when Canada adopted a national carbon price in 2019, its policy portfolio was substantially smaller than France's portfolio in 2003, after controlling for several country characteristics. This pattern is robust to different model specifications (Table D.32 in the Appendix). Visual inspection also suggests that this pattern is relatively robust to dropping possible outliers. For example, it can also be identified if the group of EU ETS countries is considered as one observation or dropped from the sample (Figure 4.6 left).

We next regress the average carbon price of the first year of implementation on the size of countries' policy portfolios, controlling for the same country characteristics as in the previous regression. The carbon price is calculated as an economy-wide average price using information on the price level and the coverage from the World Bank. We find a positive association, meaning that countries with a larger policy portfolio at the time of adoption tended to implement carbon prices with higher price levels (Figure 4.6 right). A notable outlier is Japan which implemented a relative low price given its relatively large policy portfolio. The pattern is again robust to the different model specifications (Table D.32 in the Appendix) but appears overall less significant and less robust to dropping individual countries from the sample (Figure 4.6 right).

4.4 Discussion and Conclusions

While carbon pricing is only one of many policy instruments to achieve internationally agreed climate targets, economic theory and empirical evidence on its effectiveness (Andersson, 2019; Mideksa, 2021; Green, 2021) suggest an important role for it in future national climate policy portfolios. The relationship between carbon pricing and other climate policies is generally multifaceted. Specifically, alternative instrument types can be considered as second-best substitutes of first-best policies (Bennear and Stavins, 2007; Fischer et al., 2021), complementary instruments that target different market failures (Bertram et al., 2015; Bataille et al., 2018; Stiglitz, 2019), and as instruments that remove barriers to a first-best policy (Meckling et al., 2017; Pahle et al., 2018).

Here we contribute empirical evidence on this latter idea of climate policy sequences preceding carbon pricing by examining policy adoption of G20

economies and other large emitters focusing on countries that eventually adopted a carbon price. Our analysis also builds on previous more normative work on the benefits of alternative instrument types (Peñasco et al., 2021) and policy mixes (van den Bergh et al., 2021) including complementarities between carbon pricing and other instruments (Bertram et al., 2015) and possible negative effects from their combination (Fankhauser et al., 2010), and more generally improves our understanding of climate policy adoption by examining its temporal dimension. Furthermore, we contribute to debates around the political economy of carbon pricing, and the feasibility of climate policy more broadly (for example, Carattini et al. (2018); Klenert et al. (2018); Dolphin et al. (2019); Levi et al. (2020); Ostry et al. (2021)).

Our results for the first time provide quantitative and international empirical evidence on how countries have built up their portfolios over time before eventually adopting a national carbon tax or permit system. The results suggest that carbon pricing was indeed adopted relatively late in countries individual policy sequence. Furthermore, we find qualitatively different trajectories of how countries built up their climate policy portfolios over time. While some countries did so gradually, other countries implemented national carbon pricing at the end of a quick expansion of their portfolios. A third group of countries expanded their portfolios quickly but then waited several years before eventually adopting a national price on carbon. We suspect that these more gradual or sudden expansions of portfolios reflect a country's exposure to and relative timing of domestic and international events, domestic barriers to a carbon price, and whether carbon pricing was part of a long-term climate strategy before its adoption.

Furthermore, we find that countries that adopted a carbon price in a specific year tended to have significantly larger climate policy portfolios than those that did not adopt carbon pricing. Because this methodology is not able to control for all possible confounders, we do not consider the relationship between policy portfolios and the adoption of carbon pricing as necessarily causal. Nevertheless, we illustrate in Figure 4.1 how the results are generally consistent with the idea that certain barriers to carbon pricing can be removed with other climate policies.

Examining heterogeneity among adopters of carbon pricing, we find large variation in the size of countries' policy portfolios at the time of adoption. Furthermore, over the last 20 years the size of these portfolios ap-

appears to have declined. We note that this pattern is consistent with generally declining abatement costs and international influences including the international diffusion of technological innovation (Dechezleprêtre et al., 2011; Barrett, 2021) and growing economic opportunities for green technologies (Yamazaki, 2017). Furthermore, previous research suggests that policies themselves diffuse internationally as countries learn and mirror each other (Fankhauser et al., 2016; Thisted and Thisted, 2020).

Furthermore, we find that countries that had larger climate policy portfolios at the time they adopted a carbon pricing policy tended to implement an overall higher initial economy-wide average carbon price. Our results are therefore also consistent with the idea that climate policies prior to the adoption of a pricing policy can pave the way for higher stringency. This is especially important because prior evidence suggests that carbon prices tend to be relatively sticky (Dolphin et al., 2019).

Motivated by these insights, we find that several countries including some of the world's largest emitters of GHG have reached the stage at which other countries went on to adopt a price on carbon, in terms of the instrument types of implemented policies (Figure D.41 in the Appendix). This could mean that those countries have exhausted the extensive margin of their climate policy portfolios in terms of sectors and instrument types other than carbon pricing, leaving essentially three avenues for future climate policy: higher stringency of existing policies, additional policies of existing sector and instrument type combinations, and carbon pricing as the last step in the sequence.

Our results suggest important avenues for future empirical research. Because of data limitations, we are not able to examine the experience of some of the earliest adopters of carbon pricing. As data collection efforts are ongoing, future research might focus on these countries. Furthermore, we expect that early climate policies not only lower barriers for the adoption of a pricing policy but also likely influence its effectiveness (Kriegler et al., 2018; Roelfsema et al., 2018). This influence on effectiveness might be in addition to the effect of prior policies on the initial price level that we report here. For example, broad sectoral coverage of mitigation policies can address emission leakage of pricing policies (Rajagopal, 2017). Future research might explore how the size of climate policy portfolios influences the reductions in GHG emissions obtained with a certain carbon price.

Chapter 5

The international diffusion of policies for climate change mitigation

In this paper, we study the international diffusion of carbon pricing policies. In the first part, we empirically examine to what extent the adoption of carbon pricing in a given country can explain the subsequent adoption of the same policy in other countries. In the second part, we quantify the global benefits of policy diffusion in terms of greenhouse gas emission reductions elsewhere. To do so, we combine a large international dataset on carbon pricing with several other datasets. For causal identification, we estimate semi-parametric Cox proportional hazard models. We find robust and statistically significant evidence for policy diffusion. The magnitude of the estimated effects is substantial. For two neighbouring countries, policy adoption in one country increases the probability of subsequent adoption in the other country on average by several percentage points. Motivated by this result, we use Monte Carlo simulations based on our empirical estimates to quantify both direct domestic and indirect foreign emission reductions of policy adoption and subsequent diffusion. The results based on our central empirical estimates suggest that for most countries indirect emission reductions of carbon pricing can exceed direct emission reductions. Overall, our results provide additional support for the adoption of stringent climate policies, especially in countries where climate change mitigation policies might so far have been considered as being of relatively little importance because of a relatively small domestic economy.

5.1 Introduction

Despite the need for more stringent climate policies to achieve the Paris climate target (IPCC 2021), many countries appear reluctant to ratchet up their mitigation efforts. Possible reasons include concerns about political backlashes, about international competitiveness, and about the limited effectiveness of domestic policies in reducing global greenhouse gas (GHG) emissions. Indeed, in 2021 the top 10% largest emitters contributed about 80% percent of global greenhouse gas emissions, suggesting that policies in relatively small countries will have small effects on future climate change. However, this perspective neglects that countries' domestic climate policies can also influence GHG emissions elsewhere. For example, domestic policy adoption can demonstrate political feasibility and lower concerns about international competitiveness, thereby increasing the likelihood that the same or a similar policy is adopted in other countries. Existing empirical evidence on climate policy diffusion is however mixed (Baldwin et al., 2019; Dolphin and Pollitt, 2021; Fankhauser et al., 2016; Sauquet, 2014; Thisted and Thisted, 2020) and its effectiveness in terms of GHG emission reductions has not yet been quantified.

In this paper we empirically examine the international diffusion of climate policies from 1988 to 2020 and quantify indirect emission reductions that can plausibly be attributed to policy diffusion. We focus on carbon pricing policies, which can be considered the most salient and possibly most stringent policies for climate change mitigation. We first construct a global dataset on carbon pricing, countries' characteristics, and geographic and trade linkages between countries. We then estimate Cox proportional hazard models that include spatial lags of policy adoption. The spatial lags are constructed using alternative metrics of the proximity of countries. Possible concerns about causality are addressed with a series of robustness tests and a placebo test. In the last part, we use our empirical estimates to calculate the expected emission reductions due to policy diffusion using a back-of-the-envelope methodology and Monte Carlo simulations. We consider these indirect emission reductions as a proxy for the international leverage of a country's domestic climate policy and examine its variation across countries.

We find robust statistical evidence for an international diffusion of carbon pricing policies. Countries are more likely to adopt carbon pricing if other countries that are relatively close to them in terms of geography or trade

links adopted the policy previously. We find the best model fit for a proximity metric in the spirit of a gravity model that combines the GDP of countries with the geographic distances between them. The magnitude of the diffusion effect is substantial. For example, according to our main estimates adoption of carbon pricing in Canada increases the probability of subsequent adoption in the USA by about 16 percent.

We use several robustness checks to corroborate our main findings. Possible violations of the proportional hazard assumption are addressed with covariates and stratification of the Cox proportional hazard model and systematically assessed with statistical tests. In the main specification, we use carbon pricing policies at the national and subnational level, but the results are robust to using only national pricing schemes. Furthermore, in the main specification we use both carbon taxes and ETS. If we restrict the sample to only either of the two types of carbon pricing, we find coefficients with similar magnitude but no significance. We also conduct placebo tests and do not find any evidence that would suggest spurious diffusion (Braun and Gilardi, 2006).

Our main contribution to the literature is the quantification of indirect emission reductions that can be attributed to policy diffusion, which we derive with a back-of-the-envelope calculation and Monte Carlo simulations. The indirect emission reductions quantify the emission reductions elsewhere that can be attributed to the adoption of carbon pricing in a given country. To isolate the effect of diffusion, we simulate and compare scenarios with and without policy adoption in the given country. Overall, our results suggest that the global benefits of policy diffusion are substantial. In a first set of simulations, we assume for every country that it was the first to adopt carbon pricing in 1988. We find that the indirect emission reductions due to diffusion are larger than domestic emission reductions in about 85 % of countries (1988-2019). We next examine scenarios in which a country is the next to adopt carbon pricing in 2020, given the actual distribution of pricing policies by the end of 2020, and find that indirect emission reductions exceed direct emission reductions in 76 % of the remaining countries (2020-2050).

In the last part of the analysis, we use Monte Carlo simulations to quantify to what extent policy diffusion as observed in the past can help to increase the geographical coverage of carbon pricing policies in the future. Based on the distribution of policies by the end of 2020 and the dynamic of policy

adoption and diffusion over the period 1988-2020, we simulate policy adoption for future scenarios with and without diffusion (2020-2050). We find that by 2050, about 11 percentage points more countries will adopt carbon pricing in the scenario with diffusion than in the scenario without diffusion. While for individual countries the global benefits from policy diffusion are therefore substantial, the possible contribution of policy diffusion to the achievement of a high geographical coverage of carbon pricing policies over the next decades appears limited.

We also contribute new empirical evidence on international policy diffusion, specifically diffusion of climate policies (Sauquet, 2014; Fankhauser et al., 2016; Kammerer and Namhata, 2018; Skovgaard et al., 2019; Baldwin et al., 2019; Abel, 2021; Steinebach et al., 2021; Torney, 2015; Thisted and Thisted, 2020). In agreement with the quantitative analysis of Steinebach et al. (2021) and the qualitative analysis of Thisted and Thisted (2020) we find evidence for an international diffusion of carbon pricing policies. In this respect, our results differ from the results obtained by Dolphin and Pollitt (2021) who report no evidence for diffusion of either carbon taxes or ETS, which they consider as two distinct policies. Our results therefore reconcile this seemingly contradictory prior evidence by accommodating for different implementations of the same policy. This choice is supported by the fact that in the EU (Harrison, 2010) and possibly in other cases, the decision to adopt carbon pricing was made before the instrument design was chosen. Furthermore, Skovgaard et al. (2019) find no systematic differences between countries that adopted either a tax or an ETS and observe that both designs were used in all waves of carbon pricing adoption.

Our findings also contribute new evidence using quantitative methods to prior more qualitative work that has often focused on few selected countries. This literature suggests that international coordination has been part of climate policy from its beginning, most prominently represented by the Kyoto protocol and the Paris climate agreement. This coordination in turn provides a supportive context for policy diffusion. For example, Harrison (2010) points out strong mutual influences among the world's first adopters of carbon pricing policies in Scandinavia after climate change attracted global attention for the first time in the 1980s. According to Thisted and Thisted (2020), the subsequent adoption of carbon pricing by other countries can at least partially be explained with emulation of existing policies and learning from prior experiences. International diffusion has also been actively promoted by early

adopters themselves and through multilateral initiatives such as the World Bank’s Partnership for Market Readiness (PMR) (Biedenkopf et al., 2017). Strong evidence for international diffusion has been reported for example for California (Bang et al., 2017), Kazakhstan (Gulbrandsen et al., 2017), and China (Heggelund et al., 2019), and the influence of multilateral initiatives has been acknowledged for carbon pricing policies in Latin America (Ryan and Micozzi, 2021). We consider these mechanisms and channels of international climate policy diffusion reported in prior literature as possible explanations of our results.

The remainder of the paper is structured as follows. In Section 5.2, we introduce the econometric model and estimation techniques before describing and illustrating our data. In Section 5.3, we present first our empirical results on past international diffusion of carbon pricing including several robustness tests and then the results from our back-of-the-envelope calculations and Monte Carlo simulations. We discuss and conclude in Section 5.4.

5.2 Methods

5.2.1 Empirical analysis of policy diffusion

Theories of policy diffusion propose several mechanisms through which the adoption of a policy in one jurisdiction can influence the adoption of the same or a similar policy elsewhere. These mechanisms are often grouped and referred to as learning, competition, emulation, and coercion (Braun and Giarli, 2006; Simmons et al., 2006; Shipan and Volden, 2008; Volden et al., 2008; Shipan and Volden, 2012; Jordan and Huitema, 2014). Prior literature on climate policies has especially focused on emulation and learning (Biedenkopf et al., 2017; Thisted and Thisted, 2020), which has also been identified as important mechanisms for similar diffusion processes, for example for the diffusion of cash transfer programs in Latin America (Sugiyama, 2011). Depending on the mechanism, adoption in one jurisdiction is more relevant for some jurisdictions than for others. For example, diffusion through competition suggests that policy adoption has a larger influence on jurisdictions with similar specialisation, while diffusion through coercion suggests that this influence is restricted to those jurisdiction over which a jurisdiction has a power advantage.

To identify diffusion we estimate an econometric model that relates adop-

tion of a policy in a country i at time t to the adoption of the same policy in other countries $j = 1, \dots, N_c, j \neq i$ prior to time t (with N_c being the number of countries in the sample). This is a common empirical strategy to identify policy diffusion and has been used in the literature on climate policy (Sauquet, 2014; Kammerer and Namhata, 2018; Abel, 2021; Dolphin and Pollitt, 2021). Technically, the model accounts for the mutual influences between countries with spatial lags, which are calculated as a weighted average of prior policy adoption in all other countries. We use alternative weighting schemes based on geographic proximity and trade which we consider as potentially representing some of the alternative diffusion mechanisms mentioned above.

The choice of our model is informed by some characteristics of our data. The first characteristic is that policy adoption is only observed up until 2021, the most recent year in our sample. This means that our dependent variable is generally right-censored. The second characteristic is that our dependent variable is binary taking on only values 0 or 1. Both these characteristics are common in survival analysis, which is also referred to as event history analysis, and can be addressed with proportional hazard models.

We thus follow previous work on policy diffusion and model policy diffusion with semi-parametric Cox proportional hazard models (Sugiyama, 2011; Sauquet, 2014; Abel, 2021; Dolphin and Pollitt, 2021). As compared to parametric proportional hazard models, the Cox model does not require an assumption about a specific functional form of the survival function and the results can therefore be considered more robust to model misspecification (Lee and Wang, 2003). Formally, we estimate models of the general form

$$h(t, X_{i,t}, W_{i,t}) = h_0(t) \exp(X_{i,t-1}\beta_X) \exp(W_{i,t-1}\beta_W) \quad (5.1)$$

The hazard function $h(\cdot)$ of a unit i in year t represents the probability that the policy is adopted by that unit in that year conditional on it not yet being implemented at time $t - 1$. This hazard rate is composed of a baseline hazard rate $h_0(t)$ and a second partial hazard term that includes the time-dependent matrixes $X_{i,t-1}$ and $W_{i,t-1}$.

In the Cox model, the functional form of the baseline hazard is not prescribed a-priori and not necessarily smooth, but estimated based on the patterns of policy adoption in the data. The matrix $X_{i,t-1}$ accounts for possible domestic influences in country i in year $t - 1$. Informed by prior literature on domestic

influence on the adoption of carbon pricing (Dolphin et al., 2019; Best et al., 2020), we include GDP per capita, the growth rate of GDP per capita, emissions of CO₂ per GDP, the service share of GDP and the export share of GDP. All explanatory variables are lagged by one year to address concerns about reverse causality. As a robustness test, we obtain similar results with models with longer lag times (Appendix Table E.21).

The matrix $W_{i,t-1}$ is a weighted average of policies adopted in other countries $j = 1, \dots, N_c, i \neq j$ at time $t - 1$, sometimes also referred to as a spatial lag. We explain the construction of this matrix further below.

For both the left-hand side and the right-hand side of Equation 5.1 we model adoption as a binary variable that takes on the value 1 for all years $t, t + 1, \dots, T$ if a policy has been adopted prior to or in year t . In this panel setting with time-varying covariates, observations of the same unit in subsequent years are implemented as independent of each other. To account for their dependency, we cluster the standard errors of our estimates at the level of individual units.

The model is estimated from panel data on countries' adoption of climate policies by maximising a likelihood function. Unbiasedness of the estimated coefficients relies on the proportional hazard assumption. This assumption is satisfied if conditional on all explanatory variables the hazard ratio of two units is constant over time. We address possible violations of this assumption with our set of control variables and with stratification. The control variables include GDP per capita, the growth rate of GDP per capita, emissions of CO₂ per GDP, the service share of GDP and the export share of GDP. The stratified version of our model

$$h(t, X_{i,t-1}, W_{i,t-1}) = h_{0,k}(t) \exp(X_{i,t-1}\beta_X) \exp(W_{i,t-1}\beta_W) \quad (5.2)$$

allows for different baseline hazards $h_{0,k}(t)$ for different strata with index k in our sample. For the stratified version of the model, the hazards are assumed to be proportional within strata but not necessarily across them. We use a division of the world into six continents North-America, Latin-America, Europe, Africa, Asia, and Oceania for stratification. We consider countries on the same continent as likely exposed to the same shocks that are unrelated to climate policy making. For every model, we use a statistical test based on Schoenfeld residuals to identify possible violations of the proportional

hazard assumption (Grambsch and Therneau, 1994).

The matrix W is constructed from several data sources, depending on which channel is investigated. For trade, we use data on annual bilateral trade flows from the IMF and calculate the export share $x_{i,j,t}$ and import share $m_{i,j,t}$ (percentage of exports from country i into destination j in year t out of all exports from country i in year t , analogously for imports) for every pair of countries in the data (i, j) and every year t . We then calculate a weighted average:

$$W_{i,t} = \frac{\sum_{j=1, j \neq i}^{N_c} w_{i,j,t} Y_{j,t}}{\sum_{j=1, j \neq i}^{N_c} w_{i,j,t}} \quad (5.3)$$

with $w_{i,j,t} = x_{i,j,t}$ and $w_{i,j,t} = m_{i,j,t}$ for exports and imports respectively. Note that unlike the weights described below, the weights based on trade are generally not symmetric for a pair of countries, i.e. $w_{i,j,t} \neq w_{j,i,t}$.

For geographical proximity we construct similar measures using two alternative definitions of proximity. For the first measure we use a binary variable indicating whether two countries (i, j) share a land border. The second measure is calculated from the distance between centroids of countries $d_{i,j}$ as:

$$w_{i,j} = \frac{1}{d_{i,j}}. \quad (5.4)$$

Furthermore, we construct an additional metric that is based on geographic proximity but also take the size of countries into account. This is motivated by the hypothesis that policies in larger economies have a stronger effect on policy adoption elsewhere. The size of countries is expressed by the GDP of a country. In mathematical terms, we define another set of weights

$$w_{i,j,t} = \frac{\text{GDP}_{j,t}}{d_{i,j}} \quad (5.5)$$

where $d_{i,j}$ is again the distance between countries. A country is therefore considered more influential for domestic policy adoption the closer it is in space and the larger its economy is. This metric is generally related to gravity models of international trade that make similar assumptions (Baier and Standaert, 2020).

The number of carbon pricing policies has continuously increased over the last thirty years. To address concerns about spurious diffusion (Braun and

Gilardi, 2006), we conduct a placebo test. For this purpose we construct an additional matrix $W_{i,t}$ for which we assign a random value for proximity to every country pair $w_{i,j}$ by drawing from a Weibull distribution that we fit to the empirical distribution of the distances between countries.

5.2.2 Modelling the effect of policy diffusion on GHG emissions

Back-of-the-envelope calculations

In the second step of the analysis, we use our empirical estimates to calculate the expected CO₂ emission reductions that can be causally attributed to policy diffusion. We do so in two ways, first with a back-of-the-envelope calculation and then with Monte Carlo simulations. Both methods are briefly described here and in more detail in Appendix E.1.1.

For the back-of-the-envelope calculation, we compare a scenario in which country i adopts carbon pricing in year t with a scenario in which country i does not do so. For each of the two scenarios, we calculate the hazard rate of policy adoption at time $t + 1$ for all other countries $j \neq i$ based on Equation 5.1. The difference between the hazard rates of the two scenarios can then be considered the additional hazard of policy adoption in country j that can be attributed to policy diffusion from country i .

To map the hazard rates onto greenhouse gas emissions, we assume that carbon pricing reduces emissions in all countries by the same percentage $r = 1\%$. This assumption has been made in the literature prior to our study (Eskander and Fankhauser, 2020; Best et al., 2020). Its major limitation is that it does not take into account that countries that adopt more stringent carbon pricing policies in terms of the price and sectoral coverage of the policy are likely to achieve proportionally larger emission reductions. In our idealised simulations we cannot directly use information on the stringency of policies, as for many countries no carbon pricing policy has been adopted yet. Nevertheless, we can use past carbon pricing policies to examine whether in the past earlier adopters tended to implement more or less stringent pricing policies than later adopters. If this was the case and if it was more generally representative for the international diffusion of this policy, our simulated indirect emission reductions would be biased.

We hence examine trends in the economy-wide average price in the year of the first implementation of carbon pricing policies, which we consider the

best proxy for the stringency of the policy. The results are shown in the Appendix in Figure E.22. Reassuringly for our assumption, we do not find any clear trend in the data. While some of the first adopters implemented relatively stringent policies, the trend in more recent years appears slightly positive, especially if members of the EU ETS are considered as only one observation.

Monte-Carlo simulations

The back-of-the-envelope calculations neglect differences between countries in terms of their socioeconomic characteristics and associated baseline hazard and also neglects that policies can diffuse iteratively from one country to the next. To address these limitations, we do a more comprehensive quantification of indirect emission reductions. For this purpose, we use the estimated coefficients of all control variables and the spatial lag and feed them into Monte Carlo simulations of policy adoption and policy diffusion using the model in Equation 5.1. As for the back-of-the-envelope calculations we construct counterfactual scenarios that allow us to quantify the emission reductions that can be attributed to diffusion. More details can be found in Appendix E.1.2.

For every scenario, we simulate policy adoption and diffusion over the time period 1988 and 2021, which is the time period for which we obtain our empirical estimates of diffusion. We again assume that adoption of the policy reduces greenhouse gas emissions by one percent per year and compute the cumulative emission reductions up to the year 2021.

5.2.3 Data

We use data on carbon pricing including carbon taxes and ETS from the Carbon Pricing Dashboard of the World Bank. The dataset includes pricing policies at the national and subnational level. We use the year in which a policy was adopted, i.e. in which the corresponding law was passed, because we consider this to be the point in time at which a policy can start to diffuse to other countries. We assign subnational pricing schemes to the corresponding countries and then drop for every country all but the first national or subnational pricing policy from the sample. For EU member countries, we set the year of adoption to 2003 regardless their year of ascension to avoid that the staggered EU ascension might be interpreted as diffusion in our data. The

adoption of carbon pricing over time in our sample is illustrated in Figure 5.1.

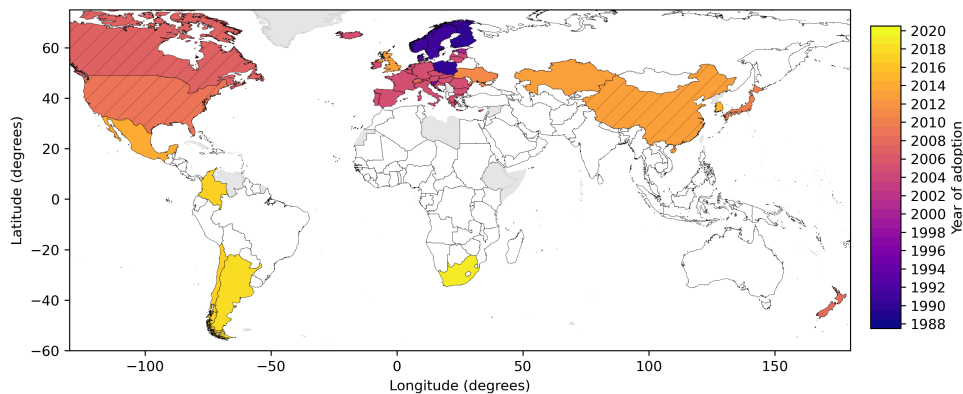


Figure 5.1: **Time of adoption of the first carbon pricing policy by country.** Hashes indicate countries in which the first policy was adopted at the subnational level.

For a robustness test, we ignore subnational pricing policies. Furthermore, for another two robustness tests we keep only either carbon tax or ETS policies in the sample.

For the explanatory variables we use additional data from the World Development Indicators of the World Bank, which we complement with replication data from a comprehensive study on carbon pricing effectiveness across countries (Best et al., 2020). Descriptive statistics of all covariates are shown in Table D.21.

Table 5.1: Descriptive statistics. The sample contains 179 countries and covers the years 1988 to 2021. A map of countries is shown in Appendix Figure E.25.

Variable	Unit	Mean	Std.	Min.	Max.	No. obs.
log GDP per capita PPP	2010 USD	8.42	1.50	5.23	11.63	6086
GDP per capita PPP growth rate	-	0.02	0.05	-1.05	0.88	6086
Exports share of GDP	percent	39.82	27.92	0.01	228.99	6086
Imports share of GDP	percent	46.85	28.96	0.00	424.82	6086
Services share of GDP	percent	21.31	13.63	0.15	55.47	6086
Emissions CO ₂ eq per GDP	t per k 2010 USD	0.62	0.95	0.00	18.39	6086

5.3 Results

5.3.1 Descriptive evidence

The diffusion of policies can be thought of as a web of leader-follower relationships, whereby policy adoption in the leading country increases the

likelihood of adoption in the following country. To better understand patterns of policy diffusion in our data, we illustrate some of those bilateral leader-follower relationships. In the empirical model that we estimate below, adoption by a follower is influenced by all leaders, but for ease of visualisation here we only plot diffusion from the leader that is closest to each follower. Proximity is based on the gravity model, which emerges as our preferred metric from the econometric analysis in the next Section.

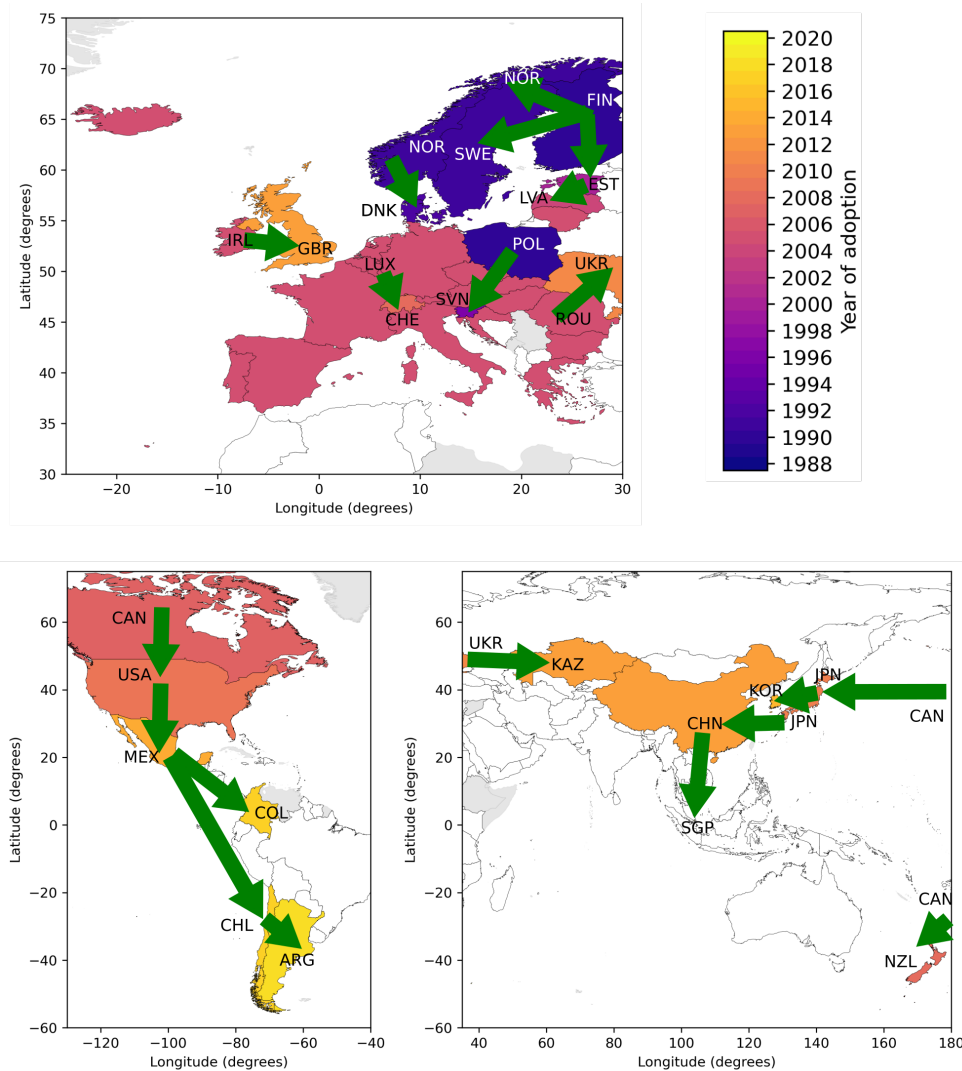


Figure 5.2: Descriptive evidence on possible leader-follower relationships among adopters of carbon pricing. Arrows point from earlier adopters to later adopters, but arrows are only shown from the leader that is closest to each of the followers according to the gravity metric. To make the figure readable, policy diffusion to members of the EU-ETS is not shown.

We focus on Europe, the American continent, and Asia and Oceania, which encompasses all carbon pricing schemes in our data except the one in South

Africa (Figure 5.2). We ignore diffusion to member countries of the EU-ETS to make the figure more readable. In Europe, policies appear to have diffused initially from Finland to other Scandinavian countries and in the Baltics. This is supported by Harrison (2010), who highlights the importance of the pioneering adoption in Finland, which was soon "emulated by its Nordic neighbors" (p. 515). In addition, carbon pricing in Poland appears to have had a relatively large influence on carbon pricing in Slovenia. Furthermore, the EU-ETS appears to have influenced the adoption of carbon pricing in the UK, Switzerland, and Ukraine, most strongly through the respective neighbouring countries Ireland, Luxemburg, and Romania.

On the American continent, pricing policies appear to have diffused from North to South, starting with subnational policies in Canada and the USA. Furthermore, Mexico appears to have played a central role in the subsequent adoption of pricing policies in South-America, specifically Colombia, Chile, and Argentina. In Asia, countries appear to have initially emulated pricing policies in Europe and North-America. Moreover, carbon pricing in Japan appears to have had a relatively large influence on its subsequent adoption in Korea, China, and Singapore.

This analysis of policy diffusion based on Figure 5.2 is of course simplistic. In the next Section, we better account for the possible complexity of the drivers of policy adoption by estimating Cox proportional hazard models, which simultaneously model the influence of a year-specific baseline hazard, several country characteristics, and prior policy adoption in all other countries.

5.3.2 Model estimates

We first examine whether there is evidence for international policy diffusion and if so, which metric of the connectedness of countries describes the diffusion of carbon pricing best. To do so, we estimate the Cox proportional hazard model as in Equation 5.1 with our six explanatory variables and the spatial lag of carbon pricing constructed from six alternative metrics of the proximity between countries: the inverse geographic distance, the presence of a shared land border, import shares, export shares, the average proximity based on these four metrics, and as gravity metric the product of the inverse distance and the GDP of a country. Similar to a gravity model, this latter metric reflects the idea that a country is more influential for domestic policy adoption the closer it is in space and the larger its economy is. The results

are presented in Columns 1-6 in Table 5.2.

Table 5.2: Results of estimation of Cox proportional hazard models with different metrics used for the construction of the spatial lag.

Policy:	Carbon price					
	Proximity	Border	Exports	Imports	Average	Gravity
Column:	1	2	3	4	5	6
Spatial lag of carbon pricing	6.8155*** (1.3603)	1.9726*** (0.4822)	2.4389** (1.0013)	2.8316** (1.2728)	4.8575*** (1.1399)	6.7053*** (1.3469)
GDP per capita PPP	11.5349*** (3.6650)	10.4090*** (3.5146)	9.9847*** (3.6180)	10.1008*** (3.6419)	10.6275*** (3.6460)	11.5662*** (3.6787)
GDP per capita PPP sq.	-0.5512*** (0.1887)	-0.4926*** (0.1820)	-0.4666** (0.1871)	-0.4751** (0.1891)	-0.5032*** (0.1880)	-0.5529*** (0.1895)
GDP per capita PPP growth	2.0603 (3.1070)	2.2431 (3.1997)	2.0444 (3.1228)	2.1894 (3.0227)	2.0292 (3.1429)	2.0597 (3.1010)
Export share	-0.0056 (0.0045)	-0.0085* (0.0049)	-0.0078* (0.0047)	-0.0068 (0.0047)	-0.0071 (0.0047)	-0.0054 (0.0045)
Services share of GDP	0.0273** (0.0136)	0.0309** (0.0135)	0.0276** (0.0127)	0.0288** (0.0128)	0.0297** (0.0139)	0.0271** (0.0136)
Emissions CO2 per GDP	-0.0127 (0.1140)	-0.0064 (0.1078)	-0.0443 (0.1005)	-0.0076 (0.1063)	0.0021 (0.1070)	-0.0121 (0.1140)
Time at risk	5277	5277	5277	5277	5277	5277
log-likelihood	-177.6	-181.3	-183.5	-183.7	-179.5	-177.5
AIC	369.2	376.5	381.0	381.5	372.9	369.1
N	5239	5207	5075	5078	5082	5252

Notes: Standard errors clustered by country in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

For all metrics we find a statistically significant and positive coefficient of the spatial lag of policy adoption. We interpret this as evidence in favour of an international diffusion of carbon pricing policies. To identify which metric describes this diffusion best, we examine the model fits using the AIC statistic. We find the best model fit for the gravity metric, followed by the inverse geographical distance between countries and the average metric. In the remainder of the paper, we therefore use the gravity metric as our preferred metric and consider the corresponding estimates in Column 6 in Table 5.2 as our baseline estimates.

We next quantify the magnitude of the estimated coefficients of the spatial lag of carbon pricing. To this aim, we select a few pairs of countries and calculate how much the adoption of carbon pricing in one country changes the hazard of policy adoption in the other country, given that no other country has previously adopted the policy. To do so, we multiply the estimated coefficient of the spatial lag of carbon pricing in Column 6 in Table 5.2 by the corresponding weight of the other country and exponentiate the result. We find that in the USA prior adoption of carbon pricing by Canada increases the hazard by about 16%, or by a factor of 1.16 (95% CI of 1.10 to 1.22). In Germany, prior adoption by France increases the hazard by 17% (10% to 24%), while in China prior adoption by Japan increases it by 10% (6% to 14%). For comparison, in the USA prior adoption by China increases the hazard by 3% and in Germany prior adoption by Japan by slightly more than 1%.

Furthermore, the estimated coefficients suggest that GDP per capita has a negative quadratic association with the hazard of carbon pricing adoption (Column 6 in Table 5.2). To illustrate the magnitude of the estimated coefficients and the declining marginal effect of higher income, the results suggest that an increase of average income from 20,000 USD to 30,000 USD is associated with an increase of the hazard by about 17 % and an increase from 30,000 USD to 40,000 USD by about 0.2 %. Furthermore, we find statistically significant coefficients for the service share of GDP, which tends to increase the hazard of carbon pricing adoption.

Variation in the hazard over time that cannot be explained by these covariates is in the model represented by the baseline hazard. We find that the baseline hazard is relatively flat except a peak in the year 2003 (Figure E.21 in the Appendix). This year coincides with the adoption of the EU ETS, which cannot sufficiently well be explained by the covariates in the model.

To test for violation of the proportional hazard assumption, we conduct a statistical test based on Schoenfeld residuals (Grambsch and Therneau, 1994). We first estimate a model that only includes the spatial lag of carbon pricing, for which we can reject proportional hazards with high confidence ($p = 0.01$). This results therefore supports our decision to include covariates in our model. For the models with six covariates whose results are shown in Table 5.2, we cannot reject the null hypothesis of proportional hazards for any of the metrics.

As a first robustness test of our main estimates, we exclude all subnational carbon pricing schemes (Column 1 in Table A.11). We find that the estimated coefficients are very similar to the model including subnational pricing policies (Column 2 in Table 5.2). Carbon pricing has first been implemented as a tax in 22 countries and as an ETS in 38 countries in our sample. We next estimate one model based on the adoption of carbon taxes alone (Column 2 in Table A.11) and one based on the adoption of ETS (Column 3). We find positive but insignificant coefficients of the spatial lag for both models, suggesting that it is important to allow for alternative implementations of carbon pricing when examining its international diffusion.

Table 5.3: Results of estimation of Cox proportional hazard models with different policies, with stratification and an additional control variable, and with a placebo spatial lag.

Policy:	Carbon price	Tax	ETS	Carbon price		
Proximity metric:	Gravity			Gravity		Placebo
Administrative level:	National	All		All		
Stratification:	None			Continents		None
Column:	1	2	3	4	5	6
Spatial lag of carbon pricing	5.3678*** (1.3377)	3.9343 (3.9888)	1.6888 (2.1695)	6.4526*** (2.4602)	5.5863** (2.1772)	4.4097 (7.9662)
GDP per capita PPP	13.5128*** (4.0646)	6.1164* (3.5528)	13.0358*** (3.9409)	10.7602*** (2.4706)	10.2111*** (2.4188)	10.6679*** (3.3850)
GDP per capita PPP sq.	-0.6514*** (0.2087)	-0.2705 (0.1886)	-0.6155*** (0.2013)	-0.5171*** (0.1246)	-0.5065*** (0.1194)	-0.4994*** (0.1755)
GDP per capita PPP growth	1.4117 (2.4841)	-3.3774 (2.1766)	7.6394** (3.4996)	2.0850 (4.4817)	0.8861 (4.6117)	2.1416 (2.3829)
Export share	0.0009 (0.0032)	-0.0090 (0.0084)	-0.0042 (0.0048)	0.0003 (0.0044)	0.0061 (0.0042)	-0.0044 (0.0032)
Services share of GDP	0.0545*** (0.0173)	0.0509 (0.0313)	0.0511*** (0.0179)	0.0419* (0.0221)	0.0446** (0.0216)	0.0453*** (0.0166)
Emissions CO2 per GDP	0.3475 (0.5183)	0.3735 (0.5549)	0.4084 (0.4692)	0.6553*** (0.1935)	0.5766** (0.2679)	0.6332* (0.3245)
Kyoto Annex I					40.5713* (20.6120)	
Time at risk	5277	5600	5395	5277	5277	5277
log-likelihood	-177.5	-98.2	-157.6	-129.3	-78.9	-186.1
AIC	369.1	210.4	329.3	272.6	173.9	386.3
N	5252	5575	5332	5252	5252	5074

Notes: Standard errors clustered by country in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

As additional robustness tests, we next estimate a stratified model as in Equation 5.2. We stratify the sample with a division of the world into the six continents North-America, Latin-America, Europe, Africa, Asia, and Oceania. We choose continents because we assume that countries on the same continent are likely to be affected similarly by possibly confounding annual shocks that are not absorbed well by the flexible baseline hazard of an unstratified model. This stratified model allows for possibly different baseline hazards on different continents after adjusting for the covariates included in the model. We find that stratification barely changes the results (Column 4 in Table A.11). Next, we allow for even more heterogeneity in the hazard rate by including an additional dummy variable that indicates whether a country is listed on Annex I of the Kyoto protocol and therefore has specific obligations under this framework. Our hypothesis is that countries with such obligations had a systematically higher baseline hazard of adopting carbon pricing than countries without such obligations. We find that the results are also robust to this additional variable (Column 5). Furthermore, we increase the lag time of the spatial lag and find similar results for periods between 1 and 5 years, with possibly the best model fit according to the AIC statistics for a lag time of 3 years (Table E.21 in the Appendix).

As a last robustness check, we conduct a placebo test. For this test, we construct the spatial lag of policy adoption by assigning random numbers to the proximities between countries. If our previous results are due to spurious diffusion, for example because of certain trends in the data, we would expect that we also find a statistically significant coefficient of prior policy adoption in this exercise. Reassuringly, we find no significance for this placebo spatial lag (Column 6 in Table A.11).

5.3.3 Emission reductions

The results from the empirical analysis above suggest that between 1988 and 2020, carbon pricing policies diffused internationally, possibly due to the learning and emulation mechanisms that we discuss. We next examine how this diffusion can contribute to reductions of greenhouse gas emissions globally. To this aim, we quantify the emission reductions that can be attributed to the adoption of carbon pricing in a given country distinguishing between direct (domestic) emissions reduction and indirect (foreign) emission reductions (due to diffusion). All results are based on the empirical estimates from the econometric analysis. We first do some back-of-the-envelope

calculations and then use Monte-Carlo simulations.

For the back-of-the-envelope calculations, we use the estimated coefficient of the diffusion of carbon pricing from the model with proximity calculated from an average metric i.e. $\beta_W = 6.7053$ (Column 6 in Table 5.2). Moreover we assume a baseline hazard of $h_0^* = 0.01$. In an additional robustness check, we set the baseline hazard to 0.05. Furthermore, we assume that adopting carbon pricing reduces total annual emissions of GHG by $r = 1$ percent per year, irrespective the total emissions of a country. This assumption is in more detail discussed in Section 5.2.2. We emphasise that this value does not influence the comparison of direct and indirect emission reductions, as both values scale with this number. We assume that the policy was implemented at the end of the year $t = 2018$ and base our calculations on actual domestic emissions E in the year $t + 1 = 2019$.

With these assumptions, we calculate direct and indirect emission reductions (Equations E.6 and E.5, respectively, in Appendix E.1.1). Because we calculate indirect emissions from diffusion for every country separately, indirect emission reductions for different countries are not additive. We find that indirect emission reductions can be substantial and similar in size to direct emission reductions. For a baseline hazard of $h_0^* = 0.01$, indirect emission reductions exceed direct emission reductions for about 38 percent of countries (Figure 5.3 left). For a baseline hazard of 0.05, the share of countries increases to 73 percent (Figure 5.3 right).

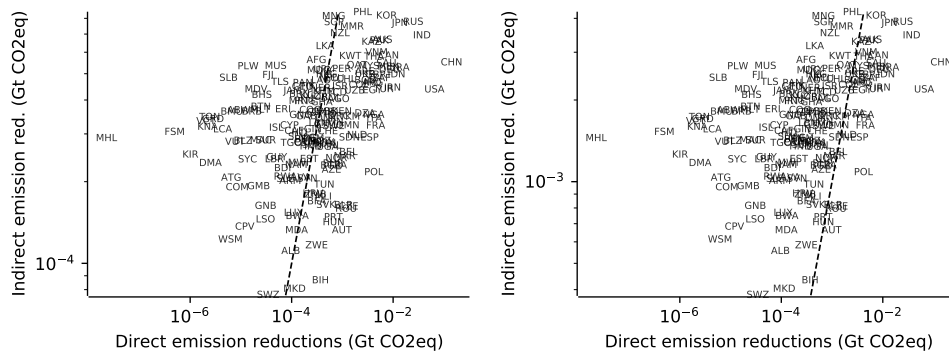


Figure 5.3: **Direct and indirect emission reductions from a back-of-the-envelope calculation.** Emission reductions calculated over one year for a policy with effectiveness of $r = 0.01$ and a baseline hazard of $h_0^* = 0.01$ (left) and 0.05 (right). For countries to the left of the straight lines indirect emission reductions exceed direct emission reductions.

For this quantification we assumed an equal and constant baseline hazard.

Furthermore, we examined only the emission reductions over the year immediately following the introduction of the policy. For this reason, the results do not account for different probabilities of adoption due to different socioeconomic contexts of countries (covariates in the empirical analysis) and ignore the possibly cascading effects of policy diffusion over several years. To address these limitations, we next conduct Monte Carlo simulations of policy diffusion.

We first assume that carbon pricing is for the first time introduced in a given country in 1988 and then diffuses from there. For the coefficient of the spatial lag and the baseline hazard, we estimate the model shown in Column 6 in Table 5.2. For simplicity, we assume a constant baseline hazard (exponential survival function), which means that in these forward simulations differences in the hazard of policy adoption stem from the spatial lag and the covariates only. The indirect emission reductions for different countries are again not additive.

The Monte Carlo simulations result in probabilities of policy adoption which we translate into expected direct and indirect emission reductions (Equations E.7 and E.8, respectively, in Appendix E.1.2). The results are shown in Figure 5.4. We find that indirect emission reductions are as large as or even larger than direct emission reductions in the majority of countries. Overall, 89 % of countries have larger indirect than direct emission reductions (Figure 5.4 left). For most of these countries, indirect emission reductions exceed direct emission reductions by a factor of 1-100, but we also find few small economies with even larger factors (Figure 5.5 left).

Countries with large indirect emission reductions tend to be relatively centrally located and close to countries with relatively large emissions. For example, the two countries with the largest indirect emission reductions are Belgium and Czech Republic. Most of the world's largest emitters are members of the G20. Those countries show a wide range of indirect emission reductions (Figure 5.4 left). Owing to their large economies, most of these countries have larger direct than indirect emission reductions, but for many of them the two tend to be of a similar order of magnitude.

This first exercise simulates policy diffusion for fictitious scenarios in which a given country is the first and only country to adopt carbon pricing in 1988. We next conduct a similar exercise which starts in 2020 from the actually observed adoption of carbon pricing by the end of 2020. We again examine two




Figure 5.4: **Direct and indirect emission reductions from Monte Carlo simulations.** Left: Emission reductions calculated over period 1988-2019 assuming no policies prior to 1988. Right: Emission reductions calculated over period 2020-2050 starting from implemented policies by the end of 2020. Parameter $r = 0.01$. G20 economies are shown in blue.

counterfactual scenarios for every country without a carbon price in 2020. In the first scenario, the country adopts carbon pricing in 2020, whereas in the second scenario it does not. The main differences to the previous simulations are therefore that policies diffuse from countries that already adopted carbon pricing by 2020 but policies cannot diffuse to them, which reduces the indirect emission reductions from diffusion for all countries but more so for some than for others. We again assume a constant baseline hazard and keep the values of all covariates at their value in 2019 (see also Appendix Figure E.26). The results are qualitatively similar to the previous results (Figure 5.4 right). Indirect emission reductions are larger than direct emission reductions in 76 % of countries in the sample (Figure 5.5 left).

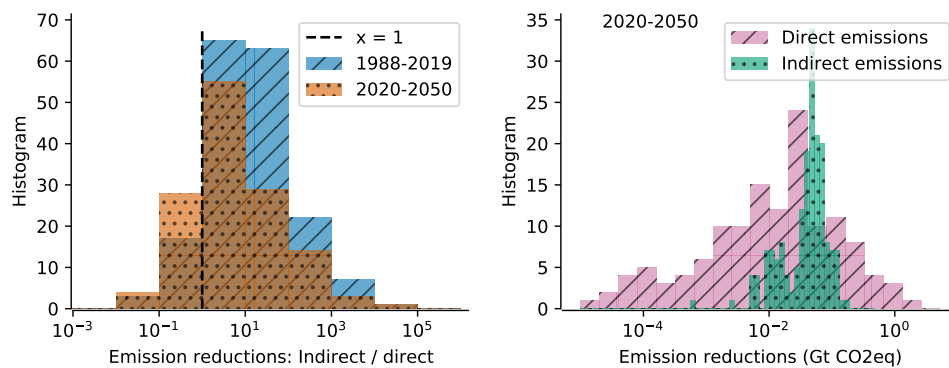


Figure 5.5: **Frequency distribution of emission reductions from Monte Carlo simulations.** Left: Histogram of ratio of indirect to direct emission reductions. Right: Histogram of direct and indirect emission reductions for the sample of 179 countries. Based on emission reductions shown in Figure 5.4.

Furthermore, we find that indirect emission reductions are far more equally distributed across countries than direct emission reductions (Figure 5.5 right). This is the case for the exercise starting in 1988 and the exercise starting in 2020. This distribution of emission reductions suggests that total emission reductions from policy adoption are more equally distributed across countries if one takes into account the emission reductions from international diffusion.

In the last part of the analysis, we examine how diffusion affects the future geographical coverage of carbon pricing policies. To this aim, we again conduct Monte Carlo simulations starting in 2020 and compare two counterfactual scenarios, one in which we use our empirical estimate of the diffusion parameter ($\beta_W = 6.7053$) and one in which we set this parameter to zero ($\beta_W = 0$). Both simulations start from carbon pricing policies that were implemented by the end of 2020. In contrast to the previous exercise, we do not need to run these simulations separately for every country because we are not interested in the effect of diffusion if a specific country adopts carbon pricing next, but instead in the effect of simultaneous diffusion from all countries with existing carbon pricing policies. All other parameter values are chosen as in the previous exercise, including the baseline hazard of policy adoption.

It appears plausible that the probability of carbon pricing adoption is generally larger in 2020-2050 than it was in 1988-2020. In a sensitivity analysis, we therefore double the baseline hazard. Importantly, in the sensitivity analysis we double the baseline hazard in the scenario with diffusion and in the scenario without diffusion to be able to again isolate the effect of diffusion.

We find that policy diffusion substantially increases the geographical coverage of carbon pricing over the time period 2020-2050 (Figure 5.6). By 2030, carbon pricing policies cover about 3.5 percentage points more countries and a 3 percentage points larger share of global greenhouse gas emissions in the scenario with diffusion than in the scenario without diffusion. By 2050, the effect of diffusion increases to 11 and 9 percentage points, respectively. Furthermore, with diffusion a similar share of countries has adopted carbon pricing by 2030 as without diffusion by 2050.

These estimates are obtained with the baseline hazard over the period 1988-2020 and the values of covariates in 2019. In the sensitivity analysis with twice the baseline hazard, the benefits of diffusion become several times




Figure 5.6: Geographical coverage of carbon pricing policies from Monte Carlo simulations for 2020-2050 with and without diffusion. The diagram shows the share of countries (left) and the share of global emissions (right) covered by carbon pricing policies for scenarios with diffusion and without diffusion. All scenarios start from carbon pricing policies implemented by the end of 2020. Based on sample of 179 countries and baseline hazard as estimated for period 1988-2020. Sensitivity analysis uses twice that baseline hazard.

larger, especially in 2030. For example, the share of countries with carbon pricing in 2030 is about 23 percentage points larger in the scenario with diffusion than in the scenario without diffusion.

These results add another nuance to the importance of international policy diffusion. While our results suggest that policy diffusion can substantially increase the geographical coverage of carbon pricing policies, this coverage increases only by about 11 percentage points of countries by 2050 relative to a scenario without diffusion (29 percentage points in the sensitivity analysis).

5.4 Discussion and Conclusions

A possible reason for the slow progress in mitigating global climate change are concerns about limited effectiveness of emission abatements in relatively small economies. Countering that concern, researchers have identified additional global benefits of a country's leadership in climate change mitigation beyond domestic emission reductions (Schwerhoff, 2016; Höhne et al., 2018). For example, stringent climate policies can support international diffusion of technological innovations that reduce mitigation costs in other countries (Dechezleprêtre et al., 2011; Barrett, 2021), demonstrate political feasibility, and create incentives related to trade (Steinebach et al., 2021) and diplomacy (Kammerer and Namhata, 2018) that nudge other countries to adopt the same

or similar policies. Overall, adoption of a climate policy at home is likely to also reduce some emissions abroad, possibly also because of the international diffusion of that policy.

In this paper, we empirically examine the diffusion of carbon pricing policies over the last 30 years and quantify the indirect emission reductions that can be attributed to policy diffusion. As compared to previous work on domestic influences on climate policy adoption (Dolphin et al., 2019; Best and Zhang, 2020; Eskander and Fankhauser, 2020), we focus on international influences. Our results are however in line with this earlier work and provide support for the importance of domestic factors, suggesting for example a positive influence of the level of GDP per capita on the adoption of carbon pricing with a declining marginal effect at higher values.

The empirical part of our paper builds on prior work on the diffusion of climate policies. Some of this prior work has also used proportional hazard models (Sauquet, 2014; Dolphin and Pollitt, 2021). Three studies have examined the diffusion of carbon pricing using qualitative (Thisted and Thisted, 2020) and similar quantitative methods (Dolphin and Pollitt, 2021; Steinebach et al., 2021). With the exception of Dolphin and Pollitt (2021), who find mixed evidence, all prior work reports evidence in support of an international diffusion of climate policies. We find robust statistical evidence for an international diffusion of carbon pricing policies. The magnitude of this diffusion is substantial: according to our estimates prior adoption of the policy by a neighbouring country increases the probability of adoption in a given year by on average about 10 %.

In contrast to the most similar prior work on the international diffusion of carbon pricing (Dolphin and Pollitt, 2021), we consider carbon taxes and ETS as two alternative designs of the same policy. This is informed by earlier findings that there are no systematic differences between countries that chose either of the two designs (Skovgaard et al., 2019). Furthermore, we consider it likely that in many cases the decision to adopt carbon pricing is likely made before the choice of instrument design, as in the case of the EU ETS (Harrison, 2010).

To some extent, the plausibility of this assumption also depends on the mechanism of diffusion. Our work does not propose a specific mechanism, but previous work suggests that learning and emulation are important for the diffusion of carbon pricing (Biedenkopf et al., 2017; Thisted and Thisted,

2020). On the one hand, if the observed diffusion is mostly due to learning from earlier experiences, as it might have been the case for the ETS in Kazakhstan that was modelled after the EU ETS (Gulbrandsen et al., 2017) and the ETS in California that intentionally differed from the EU ETS in some design parameters (Bang et al., 2017), instrument design might play a relatively more important role in diffusion. On the other hand, to the extent that diffusion is explained by emulation, for example due to an emerging international norm of carbon pricing (Thisted and Thisted, 2020), specific design parameters might be relatively less important for diffusion. Given the relatively short time periods between the adoption of carbon pricing policies in neighbouring countries in our sample, which leave little time for learning, we consider emulation as the more important process.

International coordination of climate policy is likely to be an important factor underlying this observed diffusion. Especially the Kyoto protocol and Paris climate agreement created incentives for countries to ratchet up their mitigation efforts. Ratcheting up alone can however not explain our main results, because trends over time are absorbed by the (stratified) Cox baseline term in our empirical model and our results also pass a related Placebo test. Instead, we consider it likely that the diffusion of carbon pricing can partially be explained by the efforts of early adopters to promote carbon pricing in other countries (Biedenkopf et al., 2017) and by multilateral initiatives that supported exchange of knowledge such as the International Carbon Action Partnership.

In additional analysis, we use our empirically estimated coefficients to quantify emission reductions that can plausibly be attributed to diffusion, which we refer to as indirect emission reductions. By comparing the results of a treatment and a counterfactual scenario we are able to isolate the effect of international diffusion. However, the resulting values should not be considered at face value as estimates of actual emission reductions. Above all, our results for the time period 1988-2019 are based on hypothetical scenarios in which a country adopted carbon pricing as the first and only country in 1988. To address this limitation, we also conduct simulations for the time period 2020-2050 that start from the actual adoption of carbon pricing policies in 2020. This analysis is limited in turn by the use of empirical estimates obtained from the earlier period which are extrapolated into the future. For simplicity, we also assume that carbon pricing in all countries reduces GHG emissions proportionally with a uniform annual rate. We address this limita-

tion by comparing direct and indirect emission reductions which both scale with this parameter. Lastly, due to the construction of the scenarios indirect emission reductions attributed to policy diffusion for a specific pioneering country are not additive with those indirect emission reductions attributed to other pioneering countries.

Given these limitations, our main objective here is to derive an order of magnitude of indirect emission reductions that is based on our empirical estimates and on assumptions that we consider plausible, and that takes the heterogeneous socioeconomic environments, linkages between countries, and the cascading nature of policy diffusion into account. Our results suggest that these indirect emission reductions can be substantial: using Monte Carlo simulations we find that for the majority of countries (89 % for 1988-2019 and 76 % for 2020-2050) indirect emission reductions are larger than the direct domestic reductions.

Moreover, our results suggest that indirect emission reductions due to policy diffusion are much more equally distributed across countries than domestic GHG emissions. This means that policy diffusion tends to matter relatively more in relatively small economies. Furthermore, it means that if one accounts for policy diffusion, the overall effectiveness of domestic policy adoption becomes more equal across countries. These results take into account our empirical findings that suggest that larger economies tend to have a larger effect on international diffusion. The benefits of a larger economy appear however small in our empirical results, somewhat consistent with the observation of Skovgaard et al. (2019) and the insights obtained from our descriptive analysis that many of the early adopters of carbon pricing were relatively small countries.

This insight that the emission reductions from international diffusion are relatively more important for small countries does not suggest that emission reductions in large economies are not important. Indeed, these results for small countries embody an intentional adoption of carbon pricing by large economies that is influenced by prior adoption in smaller countries. Furthermore, this insight does not conflict with possible barriers to the adoption of stringent climate policies in small countries which might be particularly exposed to competition on international markets, but highlights the possible benefits of overcoming those barriers.

In the last part of the analysis, we examine to what extent international pol-

icy diffusion as observed in the past can increase the geographical coverage of carbon pricing policies in the future. To isolate the effect of diffusion, we simulate scenarios with diffusion and without diffusion over the period 2020-2050. Our results suggest that diffusion can increase the share of countries with carbon pricing by about 11 percentage points by 2050 relative to the scenario without diffusion. As a sensitivity test, we repeat the same exercise for scenarios in which the future baseline hazard is twice as large as the historical baseline hazard, in which case the effect of diffusion increases to 29 percentage points. The results similarly show that with diffusion a similar number of countries adopts carbon pricing by 2030 as without diffusion by 2050. We emphasise again that these estimates should not be considered at face value, but indicate an order of magnitude of the effects. Overall, our results suggest that while for individual countries the global benefits from policy diffusion are therefore substantial, the possible contribution of policy diffusion to the achievement of a high geographical coverage of carbon pricing policies over the next decades appears limited.

Our study is subject to certain limitations and our results point to some avenues for future research. The empirical analysis necessarily focuses on the time period 1988 to 2020, over which the diffusion of carbon pricing might have benefitted from a generally cooperative international political environment. To what extent a possibly more fragmented international political environment will affect similar diffusion processes in the future remains an open question. More generally, any extrapolation from past policy diffusion to future diffusion should of course be made and interpreted with caution.

We focus on the adoption decision of carbon pricing policies and do not account for differences in the stringency of carbon pricing policies. For example, for the calculation of direct and indirect emission reductions, we assume the same effectiveness of carbon pricing policies for domestic emission reductions as for emission reductions in other countries. This assumption is motivated by the fact that there is no information about stringency for many countries for which we simulate policy adoption as those countries, by the end of 2021, have not yet adopted such policy. Reassuringly, we examine data on all carbon pricing policies implemented by the end of 2020 and do not find any clear trend in the initial carbon price over time, which justifies our assumption that followers implement policies with similar stringency as their leaders. Future research might examine how the stringency of pricing policies affects their diffusion and possibly the stringency of later policies.

Our analysis focuses on carbon pricing policies and subsequent work might extend this work to other climate policies. Earlier work has focused, for example, on the ratification decisions of the Kyoto protocol (Sauquet, 2014), feed-in-tariffs and renewable energy quotas (Baldwin et al., 2019; Dolphin and Pollitt, 2021), and local funding schemes for solar photovoltaic (Abel, 2021). We consider it plausible that the international political environment of climate policy will be similarly supportive to the diffusion of other types of policies and that emulation and learning will play an important role in the adoption of those policies too. More generally, we consider it plausible that those processes also matter for ratcheting up the stringency of existing climate policy, for example increases in carbon prices.

Similarly, future research might study the diffusion of carbon pricing policies at the sectoral level. International competition and the possibility of linking national ETS suggests some coordination on the inclusion of specific sectors. For example, Bullock (2012) point out how New Zealand synchronised the inclusion of agriculture in their ETS with its inclusion in the ETS considered by Australia at that time. At the same time, sectoral coverage is an important part of the stringency of a pricing scheme and might therefore influence the perceived ambitiousness of a pricing policy relative to prior policies adopted elsewhere. For example, according to Crowley (2013), the sectoral coverage relative to the EU ETS was an important consideration of the proposed ETS in Australia. Future research might therefore want to study sectoral coverage in the context of international diffusion also as a determinant of the stringency of policies.

This study focuses on international influences on climate policy adoption. Several previous studies have examined domestic influences (Fankhauser et al. (2015); Klenert et al. (2018); Dolphin et al. (2019); Levi et al. (2020), among others). Furthermore, over the last 20 years countries tended to adopt carbon pricing at the end of climate policy sequences, in most cases after the adoption of a variety of other instrument types including regulatory instruments, subsidies, research and development, and procurement and investment (Linsenmeier et al., 2022). Future research might attempt to study those international and domestic influences in one empirical framework. Furthermore, we focus on geographic proximity and trade relationships and future work might consider additional channels through which countries learn and imitate each other. For example, Kammerer and Namhata (2018) examine to what extent international climate diplomacy plays a role in diffusion.

Our results provide evidence for large positive spillovers of domestic climate policy adoption. They can be interpreted as additional support for the adoption of stringent climate policies, especially in countries where climate policies might so far have been considered as being of relatively little importance because of a relatively small domestic economy.

Concluding remarks

The development of new methods to infer causality has been of great benefit to economists studying climate change empirically. This research has been and continues to be of high importance for society. Despite recent progress on addressing climate change, such as the Paris agreement, much more evidence is needed to improve our understanding of climate change through the lenses of economics and to inform policies on mitigation and adaptation.

The five papers collected in this thesis make important contributions to the field of climate econometrics. The first two chapters push the frontier by proposing and applying novel strategies to identify the causal effect of climate variability on economies at different time scales. Their results improve our understanding of how climate influences economic activity and represent important evidence for policy makers. They also motivate paying more attention to the variability of temperature in empirical and theoretical research. A particular remaining challenge is the quantification of the costs of uncertainty and its projected changes under future global warming.

The results of the third chapter point to important heterogeneities in the response of economies to unusually cold or warm weather. Specifically, the results suggest that cold regions in Europe do not necessarily benefit from warmer-than-average years and might therefore bear relatively large costs from global warming. Furthermore, the results point to some adaptation to very cold and very hot days. These insights suggest several avenues for future research, including replication in similar contexts, more detailed studies of mechanisms, and a closer examination of adaptation.

Without additional policies, the world will likely fail to limit global warming to less than 2 degrees Celsius above pre-industrial temperature levels. The fourth and fifth chapter address specific barriers to stringent climate policies, including several simultaneous market failures, concerns about limited

effectiveness, and concerns about international competitiveness. They focus on both domestic and international factors that influence climate policy adoption. The results suggest that both perspectives can help to explain past policies to a certain extent. Future work might combine the two perspectives on mitigation policies, for example in more detailed case studies of individual countries or groups of countries, to contribute additional qualitative evidence on the importance of sequencing and diffusion of climate policies.

Bibliography

- Abel, D. (2021). The diffusion of climate policies among German municipalities. *Journal of Public Policy*, 41(1):111–136.
- Acevedo, S., Mrkaic, M., Novta, N., Pugacheva, E., and Topalova, P. (2020). The Effects of Weather Shocks on Economic Activity: What are the Channels of Impact? *Journal of Macroeconomics*, 65:103207.
- Acevedo Mejia, S., Mrkaic, M., Novta, N., Pugacheva, E., and Topalova, P. (2018). The Effects of Weather Shocks on Economic Activity: What are the Channels of Impact? *IMF Working Papers*, 18(144):1.
- Aghion, P., Bacchetta, P., Rancière, R., and Rogoff, K. (2009). Exchange rate volatility and productivity growth: The role of financial development. *Journal of Monetary Economics*, 56(4):494–513.
- Almås, I., Auffhammer, M., Bold, T., Bolliger, I., Dembo, A., Hsiang, S., Kitamura, S., Miguel, E., and Pickmans, R. (2019). Destructive Behavior, Judgment, and Economic Decision-making under Thermal Stress. *NBER Working Paper*, 25785 (revised version 2020).
- Alvarez, J. L. C. and Rossi-Hansberg, E. (2021). The Economic Geography of Global Warming. Technical Report w28466, National Bureau of Economic Research, Cambridge, MA.
- Anderson, S., Marinescu, I., and Shor, B. (2021). Can Pigou at the Polls Stop Us Melting the Poles? *Working Paper*.
- Andersson, J. J. (2019). Carbon Taxes and CO₂ Emissions: Sweden as a Case Study. *American Economic Journal: Economic Policy*, 11(4):1–30.
- Ashraf, Q. and Michalopoulos, S. (2015). Climatic Fluctuations and the Diffusion of Agriculture. *Review of Economics and Statistics*, 97(3):589–609.

- Auffhammer, M. (2018). Quantifying Economic Damages from Climate Change. *Journal of Economic Perspectives*, 32(4):33–52.
- Auffhammer, M., Hsiang, S. M., Schlenker, W., and Sobel, A. (2013). Using Weather Data and Climate Model Output in Economic Analyses of Climate Change. *Review of Environmental Economics and Policy*, 7(2):181–198.
- Baier, S. and Standaert, S. (2020). *Gravity Models and Empirical Trade*. Oxford University Press.
- Baldwin, E., Carley, S., and Nicholson-Crotty, S. (2019). Why do countries emulate each others' policies? A global study of renewable energy policy diffusion. *World Development*, 120:29–45.
- Bang, G., Victor, D. G., and Andresen, S. (2017). California's Cap-and-Trade System: Diffusion and Lessons. *Global Environmental Politics*, 17(3):12–30.
- Baranzini, A. and Carattini, S. (2017). Effectiveness, earmarking and labeling: Testing the acceptability of carbon taxes with survey data. *Environmental Economics and Policy Studies*, 19(1):197–227.
- Barrett, P. (2021). Can International Technological Diffusion Substitute for Coordinated Global Policies to Mitigate Climate Change? *IMF Working Papers*, 2021(173).
- Barsky, R. B. and Miron, J. A. (1989). The Seasonal Cycle and the Business Cycle. *Journal of Political Economy*, 97(3):503–534.
- Bataille, C., Guivarch, C., Hallegatte, S., Rogelj, J., and Waisman, H. (2018). Carbon prices across countries. *Nature Climate Change*, 8(8):648–650.
- Bathiany, S., Dakos, V., Scheffer, M., and Lenton, T. M. (2018). Climate models predict increasing temperature variability in poor countries. *Science Advances*, 4(5):1–10.
- Beaulieu, J. J., Mackie-Mason, J. K., and Miron, J. A. (1992). Why Do Countries and Industries with Large Seasonal Cycles Also Have Large Business Cycles? *The Quarterly Journal of Economics*, 107(2):621–656.
- Beaulieu, J. J. and Miron, J. A. (1992). A Cross Country Comparison of Seasonal Cycles and Business Cycles. *The Economic Journal*, 102(413):772–788.
- Bechtel, M. M., Scheve, K. F., and van Lieshout, E. (2020). Constant carbon

- pricing increases support for climate action compared to ramping up costs over time. *Nature Climate Change*, 10(11):1004–1009.
- Behrer, A. P. and Park, J. (2019). Will We Adapt - Temperature, Labor and Adaptation to Climate Change. *Working Paper*, page 39.
- Bellemare, M. F., Barrett, C. B., and Just, D. R. (2013). The Welfare Impacts of Commodity Price Volatility: Evidence from Rural Ethiopia. *American Journal of Agricultural Economics*, 95(4):877–899.
- Bennear, L. S. and Stavins, R. N. (2007). Second-best theory and the use of multiple policy instruments. *Environmental and Resource Economics*, 37(1):111–129.
- Bertram, C., Luderer, G., Pietzcker, R. C., Schmid, E., Kriegler, E., and Edenhofer, O. (2015). Complementing carbon prices with technology policies to keep climate targets within reach. *Nature Climate Change*, 5(3):235–239.
- Best, R., Burke, P. J., and Jotzo, F. (2020). Carbon Pricing Efficacy: Cross-Country Evidence. *Environmental and Resource Economics*, 77(1):69–94.
- Best, R. and Zhang, Q. Y. (2020). What explains carbon-pricing variation between countries? *Energy Policy*, 143:111541.
- Bezabih, M. and Di Falco, S. (2012). Rainfall variability and food crop portfolio choice: Evidence from Ethiopia. *Food Security*, 4(4):557–567.
- Biedenkopf, K., Müller, P., Slominski, P., and Wettestad, J. (2017). A Global Turn to Greenhouse Gas Emissions Trading? Experiments, Actors, and Diffusion. *Global Environmental Politics*, 17(3):1–11.
- Bollmeyer, C., Keller, J. D., Ohlwein, C., Wahl, S., Crewell, S., Friederichs, P., Hense, A., Keune, J., Kneifel, S., Pscheidt, I., Redl, S., and Steinke, S. (2015). Towards a high-resolution regional reanalysis for the European CORDEX domain: High-Resolution Regional Reanalysis for the European CORDEX Domain. *Quarterly Journal of the Royal Meteorological Society*, 141(686):1–15.
- Bowen, A. and Hepburn, C. (2014). Green growth: An assessment. *Oxford review of economic policy*, 30(3):407–422.
- Braun, D. and Gilardi, F. (2006). Taking ‘Galton’s Problem’ Seriously: Towards a Theory of Policy Diffusion. *Journal of Theoretical Politics*, 18(3):298–322.

- Braun, R. (1995). Seasonality and equilibrium business cycle theories. *Journal of Economic Dynamics and Control*, 19:503–531.
- Buggle, J. C. and Durante, R. (2021). Climate Risk, Cooperation and the Co-Evolution of Culture and Institutions. *The Economic Journal*, 131(637):1947–1987.
- Bullock, D. (2012). Emissions trading in New Zealand: Development, challenges and design. *Environmental Politics*, 21(4):657–675.
- Bundesbank (2012). Kalendarische Einflüsse auf das Wirtschaftsgeschehen. *Monatsbericht*, Dezember 2012:11.
- Bundesbank (2014). Wettereffekte auf das Bruttoinlandsprodukt im Winterhalbjahr 2013/2014. *Monatsbericht*, Mai 2014:168.
- Burke, M., Dykema, J., Lobell, D. B., Miguel, E., and Satyanath, S. (2015a). Incorporating Climate Uncertainty into Estimates of Climate Change Impacts. *Review of Economics and Statistics*, 97(2):461–471.
- Burke, M. and Emerick, K. (2016). Adaptation to Climate Change: Evidence from US Agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.
- Burke, M., Hsiang, S. M., and Miguel, E. (2015b). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.
- Byrne, M. P. (2021). Amplified warming of extreme temperatures over tropical land. *Nature Geoscience*, <https://doi.org/10.1038/s41561-021-00828-8>.
- Calel, R., Chapman, S. C., Stainforth, D. A., and Watkins, N. W. (2020). Temperature variability implies greater economic damages from climate change. *Nature Communications*, 11(1):5028.
- Carattini, S., Carvalho, M., and Fankhauser, S. (2018). Overcoming public resistance to carbon taxes. *WIREs Climate Change*, 9(5).
- Carleton, T. A. and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, 353(6304):1–15.
- Carrea, L., Embury, O., and Merchant, C. (2015). GloboLakes: High-resolution global limnology data. Centre for Environmental Data Analysis. <http://catalogue.ceda.ac.uk/uuid/06cef537c5b14a2e871a333b9bc0b482>.
- Carvalho, D., Cardoso Pereira, S., and Rocha, A. (2021). Future surface tem-

- peratures over Europe according to CMIP6 climate projections: An analysis with original and bias-corrected data. *Climatic Change*, 167(1-2):10.
- Cattaneo, C. and Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122:127–146.
- Center For International Earth Science Information Network-CIESIN-Columbia University (2018). Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11.
- Chatterjee, S. and Ravikumar, B. (1992). A neoclassical model of seasonal fluctuations. *Journal of Monetary Economics*, 29:59–86.
- Chen, X. and Nordhaus, W. D. (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594.
- Chen, X. and Nordhaus, W. D. (2019). VIIRS Nighttime Lights in the Estimation of Cross-Sectional and Time-Series GDP. *Remote Sensing*, 11(9):1057.
- Cinelli, C. and Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1):39–67.
- Clark, N. J., Kerry, J. T., and Fraser, C. I. (2020). Rapid winter warming could disrupt coastal marine fish community structure. *Nature Climate Change*, 10(9):862–867.
- Colacito, R., Hoffmann, B., and Phan, T. (2019). Temperature and Growth: A Panel Analysis of the United States. *Journal of Money, Credit and Banking*, 51(2-3):313–368.
- Colmer, J. (2021). Rainfall Variability, Child Labor, and Human Capital Accumulation in Rural Ethiopia. *American Journal of Agricultural Economics*, pages 858–877.
- Commission, E. (2021). Communication from the commission to the european parliament, the council, the european economics and social committee and the committee of the regions: Forging a climate-resilient europe - the new EU strategy on adaptation to climate change. Communication COM(2021) 82 final, Publications Office of the European Union.
- Connolly, M. (2008). Here Comes the Rain Again: Weather and the Intertemporal Substitution of Leisure. *Journal of Labor Economics*, 26(1):73–100.

- Conte, B., Desmet, K., Nagy, D., and Rossi-Hansberg, E. (2020). Local sectoral specialization in a warming world. *CEPR Discussion Paper*, DP15491.
- Crowley, K. (2013). Irresistible Force? Achieving Carbon Pricing in Australia. *Australian Journal of Politics & History*, 59(3):368–381.
- Cruz, J. and Rossi-Hansberg, E. (2021). The economic geography of global warming. *CEPR Discussion Paper*, DP15803.
- Cubadda, G., Savio, G., and Zelli, R. (2002). Seasonality, Productivity Shocks, and Sectoral Comovements in a Real Business Cycle Model for Italy. *Macroeconomic Dynamics*, 6(3):337–356.
- Dechezleprêtre, A., Glachant, M., Hašič, I., Johnstone, N., and Ménière, Y. (2011). Invention and Transfer of Climate Change–Mitigation Technologies: A Global Analysis. *Review of Environmental Economics and Policy*, 5(1):109–130.
- Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics*, 4(3):66–95.
- den Elzen, M., Kuramochi, T., Höhne, N., Cantzler, J., Esmeijer, K., Fekete, H., Fransen, T., Keramidas, K., Roelfsema, M., Sha, F., van Soest, H., and Vandyck, T. (2019). Are the G20 economies making enough progress to meet their NDC targets? *Energy Policy*, 126:238–250.
- Deryugina, T. and Hsiang, S. (2017). The Marginal Product of Climate. Technical Report w24072, National Bureau of Economic Research, Cambridge, MA.
- Deschênes, O. and Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–185.
- Dingel, J. I., Meng, K. C., and Hsiang, S. M. (2020). Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate. *Working paper*, page 85.
- Dolphin, G., Pollitt, M. G., and Newbery, D. M. (2019). The political economy of carbon pricing: A panel analysis. *Oxford Economic Papers*, page gpz042.
- Dolphin, G. G. and Pollitt, M. G. (2021). The international diffusion of climate policy: Theory and evidence. page 60.

- Douenne, T. and Fabre, A. (2022). Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion. *AEJ:Economic Policy* [forthcoming].
- Druckenmiller, H. and Hsiang, S. (2018). Accounting for Unobservable Heterogeneity in Cross Section Using Spatial First Differences. *NBER Working Paper Series*, 25177:1–42.
- Dwyer, J. G., Biasutti, M., and Sobel, A. H. (2012). Projected Changes in the Seasonal Cycle of Surface Temperature. *Journal of Climate*, 25(18):6359–6374.
- Elvidge, C. D., Baugh, K., Zhizhin, M., Hsu, F. C., and Ghosh, T. (2017). VIIRS night-time lights. *International Journal of Remote Sensing*, 38(21):5860–5879.
- Emerick, K. (2018). Agricultural productivity and the sectoral reallocation of labor in rural India. *Journal of Development Economics*, 135:488–503.
- Eskander, S. M. S. U. and Fankhauser, S. (2020). Reduction in greenhouse gas emissions from national climate legislation. *Nature Climate Change*, 10(8):750–756.
- Eskander, S. M. S. U., Fankhauser, S., and Setzer, J. (2020). Global Lessons from Climate Change Legislation and Litigation. *Environmental and Energy Policy and the Economy*, 2(1):44–82.
- Fankhauser, S., Gennaioli, C., and Collins, M. (2015). The political economy of passing climate change legislation: Evidence from a survey. *Global Environmental Change*, 35:52–61.
- Fankhauser, S., Gennaioli, C., and Collins, M. (2016). Do international factors influence the passage of climate change legislation? *Climate Policy*, 16(3):318–331.
- Fankhauser, S., Hepburn, C., and Park, J. (2010). Combining Multiple Climate Policy Instruments: How Not to Do It. *Climate Change Economics*, 01(03):209–225.
- Fekete, H., Kuramochi, T., Roelfsema, M., den Elzen, M., Forsell, N., Höhne, N., Luna, L., Hans, F., Sterl, S., Olivier, J., van Soest, H., Frank, S., and Gusti, M. (2021). A review of successful climate change mitigation policies in major emitting economies and the potential of global replication. *Renewable and Sustainable Energy Reviews*, 137:110602.

- Fischer, C., Hübler, M., and Schenker, O. (2021). More birds than stones – A framework for second-best energy and climate policy adjustments. *Journal of Public Economics*, 203:104515.
- Franses, P. H. (1996). Recent Advances in Modelling Seasonality. *Journal of Economic Surveys*, 10(3):299–345.
- Gädeke, A., Langer, M., Boike, J., Burke, E. J., Chang, J., Head, M., Reyer, C. P. O., Schaphoff, S., Thiery, W., and Thonicke, K. (2021). Climate change reduces winter overland travel across the Pan-Arctic even under low-end global warming scenarios. *Environmental Research Letters*, 16(2):024049.
- Gibson, J. (2020). Better Night Lights Data, For Longer. *Oxford Bulletin of Economics and Statistics*, pages 1–22.
- Gibson, J., Olivia, S., Boe-Gibson, G., and Li, C. (2021). Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149:102602.
- Graff Zivin, J. and Neidell, M. (2014). Temperature and the Allocation of Time: Implications for Climate Change. *Journal of Labor Economics*, 32(1):1–26.
- Grambsch, P. M. and Therneau, T. M. (1994). Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika*, 81(3):515–526.
- Green, J. F. (2021). Does carbon pricing reduce emissions? A review of ex-post analyses. *Environmental Research Letters*, 16(4):043004.
- Guellec, D. (2009). *Policy Responses to the Economic Crisis Investing in Innovation for Long-Term Growth*. OECD Digital Economy Papers, No.159. OECD Publishing, Paris.
- Gulbrandsen, L. H., Sammut, F., and Wettestad, J. (2017). Emissions Trading and Policy Diffusion: Complex EU ETS Emulation in Kazakhstan. *Global Environmental Politics*, 17(3):115–133.
- Hamilton, J. D. (2018). Why You Should Never Use the Hodrick-Prescott Filter. *The Review of Economics and Statistics*, 100(5):831–843.
- Hanna, E. G. and Tait, P. W. (2015). Limitations to Thermoregulation and Acclimatization Challenge Human Adaptation to Global Warming. *Int. J. Environ. Res. Public Health*, 12:8034–8074.

- Hänsel, M. C., Drupp, M. A., Johansson, D. J., Nesje, F., Azar, C., Freeman, M. C., Groom, B., and Sterner, T. (2020). Climate economics support for the UN climate targets. *Nature Climate Change*, 10(8):781–789.
- Harrison, K. (2010). The Comparative Politics of Carbon Taxation. *Annual Review of Law and Social Science*, 6(1):507–529.
- Hastings, D. A., Dunbar, P. K., Elphinstone, G. M., Bootz, M., Murakami, H., Maruyama, H., Masaharu, H., Holland, P., Payne, J., Bryant, N. A., et al. (1999). The global land one-kilometer base elevation (GLOBE) digital elevation model, version 1.0. *National Oceanic and Atmospheric Administration, National Geophysical Data Center*, 325.
- Heggelund, G., Stensdal, I., Duan, M., and Wettestad, J. (2019). China’s Development of ETS as a GHG Mitigating Policy Tool: A Case of Policy Diffusion or Domestic Drivers? *Review of Policy Research*, 36(2):168–194.
- Henderson, J. V., Squires, T., Storeygard, A., and Weil, D. (2018). The Global Distribution of Economic Activity: Nature, History, and the Role of Trade¹. *The Quarterly Journal of Economics*, 133(1):357–406.
- Henderson, J. V., Storeygard, A., and Deichmann, U. (2017). Has climate change driven urbanization in Africa? *Journal of Development Economics*, 124:60–82.
- Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring Economic Growth from Outer Space. *American Economic Review*, 102(2):994–1028.
- Hepburn, C., O’Callaghan, B., Stern, N., Stiglitz, J., and Zenghelis, D. (2020). Will COVID-19 fiscal recovery packages accelerate or retard progress on climate change? *Oxford review of economic policy*.
- Hjort, J., Karjalainen, O., Aalto, J., Westermann, S., Romanovsky, V. E., Nelson, F. E., Eitzelmüller, B., and Luoto, M. (2018). Degrading permafrost puts Arctic infrastructure at risk by mid-century. *Nature communications*, 9(1):5147–5147.
- Hjort, J., Streletskiy, D., Doré, G., and al., e. (2022). Impacts of permafrost degradation on infrastructure. *Nature reviews. Earth & environment*, 3(4):24–38.
- Höhne, N., Fekete, H., den Elzen, M. G., Hof, A. F., and Kuramochi, T. (2018).

- Assessing the ambition of post-2020 climate targets: A comprehensive framework. *Climate Policy*, 18(4):425–441.
- Hovdahl, I. (2020). Deadly Variation: The Effect of Temperature Variability on Mortality. *CAMP Working Paper Series*, 1/2020:1–56.
- Hsiang, S. (2016). Climate Econometrics. *Annual Review of Resource Economics*, 8(1):43–75.
- Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.
- Hylleberg, S., Jorgensen, C., and Sorensen, N. K. (1993). Seasonality in macroeconomic time series. *Empirical Economics*, 18(2):321–335.
- IPCC (2013). Chapter 2: Observations: Atmosphere and Surface. In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (Eds.)]*. Cambridge University Press, Cambridge, UK and New York, NY, USA.
- Jordan, A. and Huitema, D. (2014). Innovations in climate policy: The politics of invention, diffusion, and evaluation. *Environmental Politics*, 23(5):715–734.
- Kalkuhl, M. and Wenz, L. (2020). The Impact of Climate Conditions on Economic Production. Evidence from a Global Panel of Regions. *Journal of Environmental Economics and Management*, 103:1–20.
- Kammerer, M. and Namhata, C. (2018). What drives the adoption of climate change mitigation policy? A dynamic network approach to policy diffusion. *Policy Sciences*, 51(4):477–513.
- Kelly, S., Moore, T. N., de Eyto, E., Dillane, M., Goulon, C., Guillard, J., Lasne, E., McGinnity, P., Poole, R., Winfield, I. J., Woolway, R. I., and Jennings, E. (2020). Warming winters threaten peripheral Arctic charr populations of Europe. *Climatic Change*, 163(1):599–618.
- Kikstra, J. S., Waidelich, P., Rising, J., Yumashev, D., Hope, C., and Brierley, C. M. (2021). The social cost of carbon dioxide under climate-economy

- feedbacks and temperature variability. *Environmental Research Letters*, 16(9):1–33.
- Klenert, D., Mattauch, L., Combet, E., Edenhofer, O., Hepburn, C., Rafaty, R., and Stern, N. (2018). Making carbon pricing work for citizens. *Nature Climate Change*, 8(8):669–677.
- Kotchen, M. J., Turk, Z. M., and Leiserowitz, A. A. (2017). Public willingness to pay for a US carbon tax and preferences for spending the revenue. *Environmental Research Letters*, 12(9):094012.
- Kotz, M., Wenz, L., and Levermann, A. (2021a). Footprint of greenhouse forcing in daily temperature variability. *Proceedings of the National Academy of Sciences*, 118(32):1–8.
- Kotz, M., Wenz, L., Stechemesser, A., Kalkuhl, M., and Levermann, A. (2021b). Day-to-day temperature variability reduces economic growth. *Nature Climate Change*, 11(4):319–325.
- Kriegler, E., Bertram, C., Kuramochi, T., Jakob, M., Pehl, M., Stevanović, M., Höhne, N., Luderer, G., Minx, J. C., Fekete, H., Hilaire, J., Luna, L., Popp, A., Steckel, J. C., Sterl, S., Yalaw, A. W., Dietrich, J. P., and Edenhofer, O. (2018). Short term policies to keep the door open for Paris climate goals. *Environmental Research Letters*, 13(7):074022.
- Lee, E. T. and Wang, J. W. (2003). *Statistical Methods for Survival Data Analysis*. Wiley Series in Probability and Statistics. Wiley, 3rd edition.
- Legates, D. R. and Willmott, C. J. (1990). Mean seasonal and spatial variability in global surface air temperature. *Theoretical and Applied Climatology*, 41(1-2):11–21.
- Levi, S., Flachsland, C., and Jakob, M. (2020). Political Economy Determinants of Carbon Pricing. *Global Environmental Politics*, 20(2):128–156.
- Linsenmeier, M., Mohommad, A., and Schwerhoff, G. (2022). Policy sequencing towards carbon pricing - empirical evidence from G20 economies and other major emitters. *IMF Working Paper No. 2022/066*.
- Lumsdaine, R. L. and Prasad, E. S. (2003). Identifying the Common Component of International Economic Fluctuations: A new Approach. *The Economic Journal*, 113(484):101–127.

- Malik, A. and Temple, J. R. (2009). The geography of output volatility. *Journal of Development Economics*, 90(2):163–178.
- Mann, M. E. and Park, J. (1994). Global modes of surface temperature variability on interannual to century time scales. *Journal of Geophysical Research*, 99(D12):25819–25833.
- McKinnon, K. A., Stine, A. R., and Huybers, P. (2013). The Spatial Structure of the Annual Cycle in Surface Temperature: Amplitude, Phase, and Lagrangian History. *Journal of Climate*, 26(20):7852–7862.
- Meckling, J., Kelsey, N., Biber, E., and Zysman, J. (2015). Winning coalitions for climate policy. *Science*, 349(6253):1170–1171.
- Meckling, J., Sterner, T., and Wagner, G. (2017). Policy sequencing toward decarbonization. *Nature Energy*, 2(12):918–922.
- Mendelsohn, R., Basist, A., Dinar, A., Kurukulasuriya, P., and Williams, C. (2007a). What explains agricultural performance: Climate normals or climate variance? *Climatic Change*, 81(1):85–99.
- Mendelsohn, R., Kurukulasuriya, P., Basist, A., Kogan, F., and Williams, C. (2007b). Climate analysis with satellite versus weather station data. *Climatic Change*, 81(1):71–83.
- Mendelsohn, R. O. and Massetti, E. (2017). The Use of Cross-Sectional Analysis to Measure Climate Impacts on Agriculture: Theory and Evidence. *Review of Environmental Economics and Policy*, 11(2):280–298.
- Mideksa, T. K. (2021). Pricing for a Cooler Planet: An Empirical Analysis of the Effect of Taxing Carbon. *CESifo Working Paper no. 9172*, page 25.
- Mildenberger, M., Lachapelle, E., Harrison, K., and Stadelmann-Steffen, I. (2022). Limited impacts of carbon tax rebate programmes on public support for carbon pricing. *Nature Climate Change*.
- Miron, J. A. and Beaulieu, J. J. (1996). What Have Macroeconomists Learned about Business Cycles from the Study of Seasonal Cycles? *The Review of Economics and Statistics*, 78(1):54–66.
- Mitton, T. (2016). The wealth of subnations: Geography, institutions, and within-country development. *Journal of Development Economics*, 118:88–111.

- Moberg, A., Jones, P. D., Barriendos, M., Bergström, H., Camuffo, D., Cocheo, C., Davies, T. D., Demarée, G., Martin-Vide, J., Maugeri, M., Rodriguez, R., and Verhoeve, T. (2000). Day-to-day temperature variability trends in 160- to 275-year-long European instrumental records. *Journal of Geophysical Research: Atmospheres*, 105(D18):22849–22868.
- Nascimento, L., Kuramochi, T., Iacobuta, G., den Elzen, M., Fekete, H., Weishaupt, M., van Soest, H. L., Roelfsema, M., Vivero-Serrano, G. D., Lui, S., Hans, F., Jose de Villafranca Casas, M., and Höhne, N. (2021). Twenty years of climate policy: G20 coverage and gaps. *Climate Policy*, pages 1–17.
- Nordhaus, W. and Chen, X. (2015). A sharper image? Estimates of the precision of nighttime lights as a proxy for economic statistics. *Journal of Economic Geography*, 15(1):217–246.
- Nordhaus, W. D. (2006). Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, 103(10):3510–3517.
- Nunn, N. and Puga, D. (2012). Ruggedness: The Blessing of Bad Geography in Africa. *The Review of Economics and Statistics*, 94(1):20–36.
- Ostry, J., Furceri, D., and Ganslmeier, M. (2021). Are Climate Change Policies Politically Costly? *IMF Working Papers*, 2021(156).
- Pahle, M., Burtraw, D., Flachsland, C., Kelsey, N., Biber, E., Meckling, J., Edenhofer, O., and Zysman, J. (2018). Sequencing to ratchet up climate policy stringency. *Nature Climate Change*, 8(10):861–867.
- Pankratz, N. M. C. and Schiller, C. (2021). Climate Change and Adaptation in Global Supply-Chain Networks. *SSRN Electronic Journal*.
- Parker, W. S. (2016). Reanalyses and Observations: What’s the Difference? *Bulletin of the American Meteorological Society*, 97(9):1565–1572.
- Peñasco, C., Anadón, L. D., and Verdolini, E. (2021). Systematic review of the outcomes and trade-offs of ten types of decarbonization policy instruments. *Nature Climate Change*, 11(3):257–265.
- Pretis, F., Schwarz, M., Tang, K., Haustein, K., and Allen, M. R. (2018). Uncertain impacts on economic growth when stabilizing global temperatures at 1.5°C or 2°C warming. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2119):20160460.

- Rajagopal, D. (2017). A synthesis of unilateral approaches to mitigating emissions leakage under incomplete policies. *Climate Policy*, 17(5):573–590.
- Ramankutty, N., Evan, A., Monfreda, C., and Foley, J. (2010). Global Agricultural Lands: Pastures, 2000.
- Ramankutty, N., Evan, A. T., Monfreda, C., and Foley, J. A. (2008). Farming the Planet: 1. Geographic Distribution of Global Agricultural Lands in the Year 2000: Global Agricultural Lands in 2000. *Global Biogeochemical Cycles*, 22(1).
- Ramey, G. and Ramey, V. (1994). Cross-Country Evidence on the Link Between Volatility and Growth. *The American Economic Review*, 85(5):1138–1151.
- Rode, A., Carleton, T., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Jina, A., Kopp, R. E., McCusker, K. E., Nath, I., Rising, J., and Yuan, J. (2021). Estimating a social cost of carbon for global energy consumption. *Nature*, 598(7880):308–314.
- Roelfsema, M., Fekete, H., Höhne, N., den Elzen, M., Forsell, N., Kuramochi, T., de Coninck, H., and van Vuuren, D. P. (2018). Reducing global GHG emissions by replicating successful sector examples: The ‘good practice policies’ scenario. *Climate Policy*, 18(9):1103–1113.
- Roelfsema, M., van Soest, H. L., Harmsen, M., van Vuuren, D. P., Bertram, C., den Elzen, M., Höhne, N., Iacobuta, G., Krey, V., Kriegler, E., Luderer, G., Riahi, K., Ueckerdt, F., Després, J., Drouet, L., Emmerling, J., Frank, S., Fricko, O., Gidden, M., Humpenöder, F., Huppmann, D., Fujimori, S., Fragkiadakis, K., Gi, K., Keramidas, K., Köberle, A. C., Aleluia Reis, L., Rochedo, P., Schaeffer, R., Oshiro, K., Vrontisi, Z., Chen, W., Iyer, G. C., Edmonds, J., Kannavou, M., Jiang, K., Mathur, R., Safonov, G., and Vishwanathan, S. S. (2020). Taking stock of national climate policies to evaluate implementation of the Paris Agreement. *Nature Communications*, 11(1):2096.
- Rudik, I., Lyn, G., Tan, W., and Ortiz-Bobea, A. (2021). Heterogeneity and Market Adaptation to Climate Change in Dynamic-Spatial Equilibrium. *Working paper June 2021 [https://osf.io/usghb]*.
- Ryan, D. and Micozzi, M. (2021). The politics of climate policy innovation:

- The case of the Argentine carbon tax. *Environmental Politics*, 30(7):1155–1173.
- Sauquet, A. (2014). Exploring the nature of inter-country interactions in the process of ratifying international environmental agreements: The case of the Kyoto Protocol. *Public Choice*, 159(1-2):141–158.
- Schlenker, W., Hanemann, W. M., and Fisher, A. C. (2006). The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions. *The Review of Economics and Statistics*, 88(1):113–125.
- Schwerhoff, G. (2016). The economics of leadership in climate change mitigation. *Climate Policy*, 16(2):196–214.
- Seo, S. N. and Mendelsohn, R. (2008). An analysis of crop choice: Adapting to climate change in South American farms. *Ecological Economics*, 67(1):109–116.
- Shaver, A., Carter, D. B., and Shawa, T. W. (2018). Terrain ruggedness and land cover: Improved data for most research designs.
- Shipan, C. R. and Volden, C. (2008). The Mechanisms of Policy Diffusion. *American Journal of Political Science*, 52(4):840–857.
- Shipan, C. R. and Volden, C. (2012). Policy Diffusion: Seven Lessons for Scholars and Practitioners. *Public Administration Review*, 72(6):788–796.
- Simmons, B. A., Dobbin, F., and Garrett, G. (2006). Introduction: The International Diffusion of Liberalism. *International Organization*, 60(04).
- Skovgaard, J., Ferrari, S. S., and Knaggård, Å. (2019). Mapping and clustering the adoption of carbon pricing policies: What polities price carbon and why? *Climate Policy*, 19(9):1173–1185.
- Steinebach, Y., Fernández-i-Marín, X., and Aschenbrenner, C. (2021). Who puts a price on carbon, why and how? A global empirical analysis of carbon pricing policies. *Climate Policy*, 21(3):277–289.
- Stiglitz, J. E. (2019). Addressing climate change through price and non-price interventions. *European Economic Review*, 119:594–612.
- Stine, A. R. and Huybers, P. (2012). Changes in the Seasonal Cycle of Temperature and Atmospheric Circulation. *Journal of Climate*, 25(21):7362–7380.

- Sugiyama, N. B. (2011). The diffusion of Conditional Cash Transfer programs in the Americas. *Global Social Policy*, 11(2-3):250–278.
- Thisted, E. V. and Thisted, R. V. (2020). The diffusion of carbon taxes and emission trading schemes: The emerging norm of carbon pricing. *Environmental Politics*, 29(5):804–824.
- Torney, D. (2015). Bilateral Climate Cooperation: The EU’s Relations with China and India. *Global Environmental Politics*, 15(1):105–122.
- van den Bergh, J., Castro, J., Drews, S., Exadaktylos, F., Foramitti, J., Klein, F., Konc, T., and Savin, I. (2021). Designing an effective climate-policy mix: Accounting for instrument synergy. *Climate Policy*, pages 1–20.
- Van Passel, S., Massetti, E., and Mendelsohn, R. (2017). A Ricardian Analysis of the Impact of Climate Change on European Agriculture. *Environmental and Resource Economics*, 67(4):725–760.
- Volden, C., Ting, M. M., and Carpenter, D. P. (2008). A Formal Model of Learning and Policy Diffusion. *American Political Science Review*, 102(3):319–332.
- Waldinger, M. (2022). The Economic Effects of Long-Term Climate Change: Evidence from the Little Ice Age. *Journal of Political Economy*, forthcoming.
- Wan, H., Kirchmeier-Young, M. C., and Zhang, X. (2021). Human influence on daily temperature variability over land. *Environmental Research Letters*, 16(9):094026.
- Wheeler, T. R., Craufurd, P. Q., Ellis, R. H., Porter, J. R., and Vara Prasad, P. (2000). Temperature variability and the yield of annual crops. *Agriculture, Ecosystems & Environment*, 82(1-3):159–167.
- [World Bank] (2022). Carbon Pricing Dashboard. <https://carbonpricingdashboard.worldbank.org/>.
- Xu, Z., Han, Y., Tam, C.-Y., Yang, Z.-L., and Fu, C. (2021). Bias-corrected CMIP6 global dataset for dynamical downscaling of the historical and future climate (1979–2100). *Scientific Data*, 8(1):293.
- Yamazaki, A. (2017). Jobs and climate policy: Evidence from British Columbia’s revenue-neutral carbon tax. *Journal of Environmental Economics and Management*, 83:197–216.

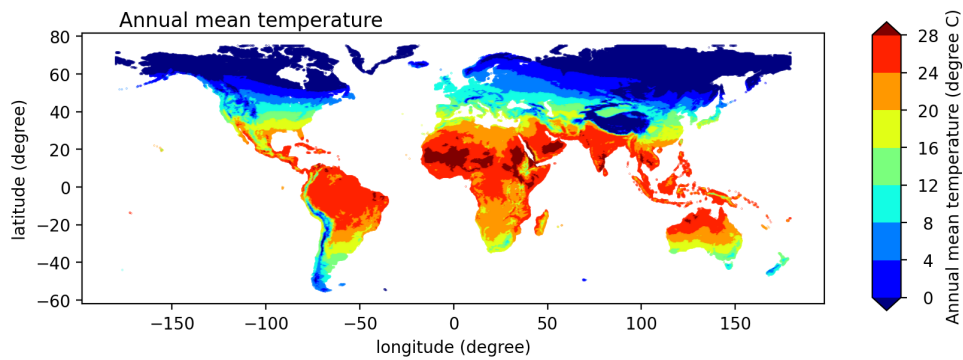
Yao, J. and Zhao, Y. (2022). Structural Breaks in Carbon Emissions: A Machine Learning Analysis. *IMF Working Papers*, 2022(9).

Appendix A

Appendix of Chapter 1

A.1 Map of annual mean temperature

Figure A.1: Global map of annual mean temperature.



Notes: The figures shows the global distribution of annual mean temperature using the same bins as in Section 1.5.2. Source is ERA-5 reanalysis for 1985-2014 (see Section 1.3.2).

A.2 Omitted variable biases and multicollinearity

To illustrate the strengths of the SFD framework, I estimate a simple model in which I explain variation of nightlights by day-to-day temperature variability. The exercise focuses on day-to-day variability as I can use recent estimates of its effect on regional GDP per capita using variation across time for identification as benchmark (Kotz et al., 2021b). I use two models, one without any other explanatory variables and one that also includes annual mean temperature. I first estimate the two models using levels of all variables and then with the SFD estimator.

I first focus on the model with only day-to-day variability (Columns 1 and 2 in Table A.1). I find that using levels yields a significantly positive coefficient, contrary to the result by Kotz et al. (2021b). Visual inspection of Figure 1.2 shows that levels of day-to-day variability are relatively low in the tropics and tend to increase with latitude. Estimates using levels could thus be confounded by any other variable correlated with latitude, such as institutions/colonial legacies. By contrast, the SFD estimator yields a significantly negative coefficient, consistent with the previously reported result.

Table A.1: Results of a simple model estimated with levels and SFD.

Dependent variable:	<i>log Nightlight density</i>			
	levels		SFD	
Estimator:				
Column	1	2	3	4
Day-to-day variab. of T	0.11857*** (0.03601)	0.04641** (0.02297)	-0.76479*** (0.19528)	-0.68200*** (0.18304)
Annual mean temperature		0.12265*** (0.02853)		0.89835*** (0.11474)
R2	0.0673	0.1055	0.0041	0.0082
df	224454	224453	448909	448908

*Notes: The table shows the results of a linear model similar to Equation 1.4 but without interaction terms estimated with spatial first-differences. Differences in West-East and North-South are pooled. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.*

Because I use different data and focus on long-term rather than short-term effects of day-to-day variability, the comparability of these estimates with previous results is limited. Given that the assumptions that need to be satisfied for an unbiased estimate from levels are stronger than those of the SFD estimator, I expect that the estimates obtained from levels are more prone to

omitted variable biases. The results presented in Table A.1 include some evidence for it: including annual mean temperature in the model changes the estimated coefficient of day-to-day variability obtained from levels by about 60 percent, while its inclusion changes the SFD estimate by much smaller 10 percent.

This latter result points to another advantage of the SFD estimator as compared to the levels-estimator. Because temperature variability at all frequencies (day-to-day, seasonal, interannual) and annual mean temperature are all influenced by latitude (Section 1.2.2), levels of these variables tend to be highly correlated. This raises concerns about multicollinearity, which has been recognised as a major challenge of empirically disentangling the effect of multiple climate variables (Auffhammer et al., 2013).

Table A.2: Variance inflation factors for a model including all geographic and climatic controls in addition to day-to-day, seasonal, and interannual temperature variability.

Variable	Levels	Spatial first-differences		
		Pooled	NS	WE
Day-to-day var. of temperature	59.11	1.87	1.94	1.81
Seasonal var. of temperature	58.94	1.91	2.00	1.81
Interannual var. of temperature	35.02	1.47	1.50	1.42
Annual mean temperature	141.12	3.83	4.13	3.48

Notes: The table shows the VIF of linear models including annual mean temperature, its day-to-day, seasonal, and interannual variability, as well as all climatic and geographic control variables shown in Table 1.1. Estimates obtained from spatial first-differences are shown for differences in West-East (WE) and North-South (NS) direction and for differences in the two directions pooled.

A common indicator of multicollinearity in a model is the Variance Inflation Factor (VIF) which is a measure of how much variation of one explanatory variable in a model is explained by all the other explanatory variables. I calculate the VIF for a model in which I include the annual mean of temperature and its day-to-day, seasonal, and interannual variability as well as linear terms of all climatic and geographic control variables (Table 1.1). Typical critical thresholds for multicollinearity are 5 and 10, corresponding to 80 and 90 percent of all variation being explained by other explanatory variables. I find that multicollinearity is indeed a major concern for the levels-estimator, but is mitigated by using spatial first-differences (Table A.2). This analysis of

the VIF simultaneously accounts for the correlation of temperature variability at different time scales and its correlation with any of the climatic and geographic control variables. Focusing only on the correlation of temperature variability across time scales, I find that spatial first-differencing also substantially reduces their cross-correlations (Table A.3).

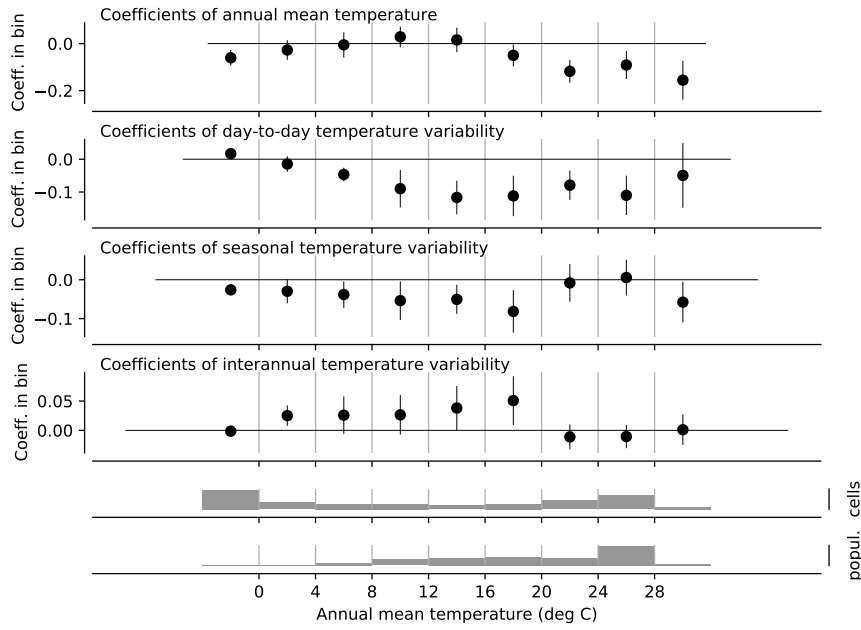
Table A.3: Pearson correlation coefficients between different temperature variables.

Variable	(0)	(1)	(2)	(3)
<i>Pearson correlation of levels</i>				
(0) Annual mean temperature		-0.83	-0.86	-0.81
(1) Day-to-day var. of temperature	-0.83		0.90	0.89
(2) Seasonal var. of temperature	-0.86	0.90		0.83
(3) Interannual var. of temperature	-0.81	0.89	0.83	
<i>Pearson correlation of spatial first-differences (WE)</i>				
(0) Annual mean temperature		-0.07	0.08	0.05
(1) Day-to-day var. of temperature	-0.07		0.60	0.44
(2) Seasonal var. of temperature	0.08	0.60		0.49
(3) Interannual var. of temperature	0.05	0.44	0.49	
<i>Pearson correlation of spatial first-differences (NS)</i>				
(0) Annual mean temperature		-0.13	0.11	0.01
(1) Day-to-day var. of temperature	-0.13		0.60	0.50
(2) Seasonal var. of temperature	0.11	0.60		0.51
(3) Interannual var. of temperature	0.01	0.50	0.51	

Notes: n/a.

A.3 Estimated coefficients across bins of annual mean temperature

Figure A.2: Estimated marginal effects of annual mean temperature and day-to-day, seasonal, and interannual temperature variability at different levels of annual mean temperature.



Notes: The figure shows the estimated coefficients of a model with linear terms for annual mean temperature, day-to-day, seasonal, and interannual temperature variability in bins of annual mean temperature (Equation 1.4). The coefficients can be interpreted as marginal effects. Error bars show 95 percent confidence intervals. The bottom two rows show histograms of grid cells and population within the same temperature bins. The geographic distribution of these bins is shown in Figure A.1 in Appendix A.1.

A.4 Alternative measures for seasonal variability

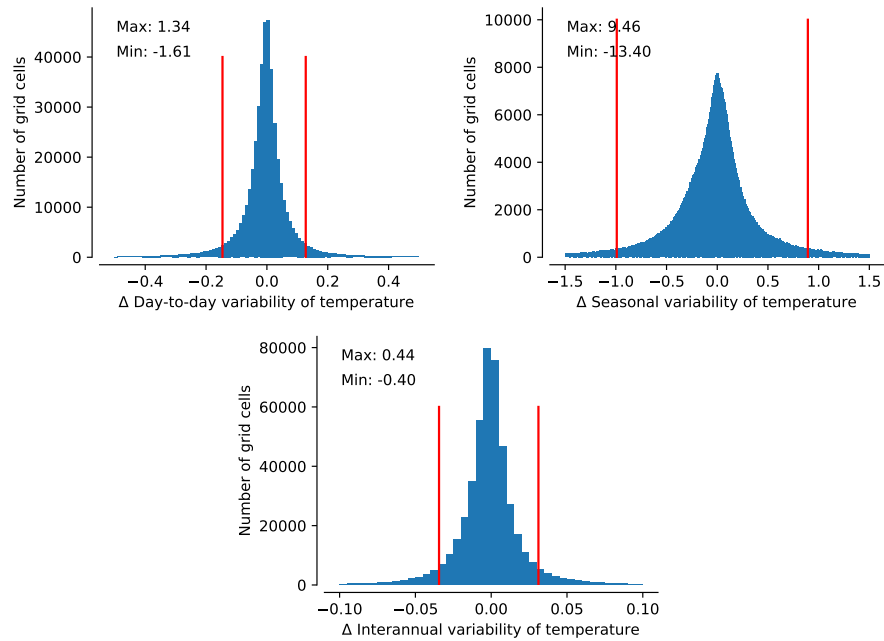
Table A.4: Results of a model estimated with SFD with two alternative measures for seasonal variability of temperature.

Dependent variable: Seasonal variability: Column	<i>log Nightlight density</i>	
	range	std
	1	2
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.50031*** (0.13259)
Seasonal variab. of T (std)		-0.30724* (0.16225)
Seasonal variab. of T (range)	-0.28016 (0.17127)	
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.16939*** (0.04337)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.25686** (0.10038)
Day-to-day variab. of T	-0.11631	-0.11535
Seasonal variab. of T (std)		-0.01906
Seasonal variab. of T (range)	-0.00632	
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.19185	0.18710
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.27857	-0.28372
Climate controls (linear)	x	x
Climate controls (quadratic)	x	x
Geographic controls (linear)	x	x
Geographic controls (quadratic)	x	x
R2	0.0249	0.0250
df	448877	448877

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Discontinuities of treatment

Figure A.3: Histograms of first-differences in temperature variability.



Notes: The figure shows histograms of observations in terms of spatial first-differences in temperature variability pooling differences in WE and NS. Red lines indicate 5% and 95% percentile.

Table A.5: Results of a model estimated with SFD with subsamples, excluding the bottom 5% and top 5% of observations in terms of the first-difference in temperature variability.

Dependent variable:	<i>log Nightlight density</i>		
Spatial first differences:	Pooled	Pooled	Pooled
Sampling based on:	Day-to-day	Season.	Inter-ann.
Day-to-day variab. of T	-0.79884*** (0.14882)	-0.76855*** (0.15450)	-0.61238*** (0.16967)
Seasonal variab. of T	-0.41895*** (0.13591)	-0.08313 (0.16946)	-0.31921** (0.14780)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.16195*** (0.03496)	0.11883*** (0.04223)	0.16700* (0.08586)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.13970 (0.08815)	-0.20035** (0.09252)	-0.07319 (0.09005)
Climate controls (linear)	x	x	x
Climate controls (quadratic)	x	x	x
Geographic controls (linear)	x	x	x
Geographic controls (quadratic)	x	x	x
R2	0.0183	0.0217	0.0212
df	403988	403988	403985

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. Standard errors in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Results of a model estimated with SFD with and without the inclusion of spatial lags.

Dependent variable:	<i>log Nightlight density</i>	
	Pooled	Pooled
Columns:	1	2
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.50290*** (0.11963)
Seasonal variab. of T	-0.28016 (0.17127)	-0.37287** (0.18787)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.18756*** (0.04990)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.25231** (0.09813)
Climate controls (linear)	x	x
Climate controls (quadratic)	x	x
Geographic controls (linear)	x	x
Geographic controls (quadratic)	x	x
Spatial lags		x
R2	0.0249	0.0251
df	448877	447216

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. The model with lags includes the first spatial lag of annual mean temperature, its squared term, and all variables that measure temperature variability. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Results of a model estimated with SFD where spatial differences are calculated by skipping Δi observations.

Dependent variable:	<i>log Nightlight density</i>	
	Pooled	Pooled
Spatial first differences:		
Δi :	0	1
Column:	1	2
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.55108*** (0.14430)
Seasonal variab. of T	-0.28016 (0.17127)	-0.29782* (0.16841)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.12445*** (0.03979)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.19516* (0.10010)
Climate controls (linear)	x	x
Climate controls (quadratic)	x	x
Geographic controls (linear)	x	x
Geographic controls (quadratic)	x	x
R2	0.0249	0.0302
df	448877	433257

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. $\Delta i = 0$ corresponds to the main specification in this paper. $\Delta i = 1$ means that for taking spatial differences, grid cells are not matched with their immediate neighbour, but with their neighbour's neighbour. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.6 Climate data for 1955-1984

Table A.8: Results of a model estimated with SFD for different climate periods.

Dependent variable:	<i>log Nightlight density</i>	
	1985-2014	1955-1984
Time period (climate):		
Column	1	2
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.47408*** (0.13475)
Seasonal variab. of T	-0.28016 (0.17127)	-0.27510 (0.19847)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.08998 (0.07277)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.19300*** (0.07264)
<i>Effect of increase by 1 deg. C on log nightlights</i>		
Day-to-day variab. of T	-0.11209	-0.10534
Seasonal variab. of T	-0.00623	-0.00612
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.19115	0.09902
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.27755	-0.21240
Climate controls (linear)	x	x
Climate controls (quadratic)	x	x
Geographic controls (linear)	x	x
Geographic controls (quadratic)	x	x
R2	0.0249	0.0247
df	448877	448877

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. Standard errors in parentheses. Nightlights are for the year 2015. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.7 Nightlights for 2015-2019

Table A.9: Results of a model estimated with SFD using different versions and time periods of the nightlights data.

Dependent variable:	<i>Nightlight density</i>			
	Pooled	Pooled	Pooled	Pooled
Spatial first differences:	log	arcsinh	log	log
Transformation of dep. variable:	v1	v1	v2	v2
VIIRS version	2015	2015	2015	2015-2019
Nightlights time period				
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.48670*** (0.13155)	-0.49211*** (0.13332)	-0.48630*** (0.13045)
Seasonal variab. of T	-0.28016 (0.17127)	-0.26858 (0.16780)	-0.26720 (0.16633)	-0.26709 (0.16823)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.15849*** (0.04734)	0.16691*** (0.04755)	0.16386*** (0.04906)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.25988** (0.10321)	-0.25554** (0.10221)	-0.27256*** (0.10035)
Climate controls (linear)	x	x	x	x
Climate controls (quadratic)	x	x	x	x
Geographic controls (linear)	x	x	x	x
Geographic controls (quadratic)	x	x	x	x
R2	0.0249	0.0252	0.0255	0.0274
df	448877	448877	448877	448877

Notes: The table shows the results of a model as shown in Equation 1.5 estimated with spatial first-differences, pooling differences in WE and NS. Standard errors in parentheses. Main estimates are obtained with nightlights in version 1 for the year 2015. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.8 Sensitivity analysis

The sensitivity of my estimated coefficients of temperature variability to the inclusion of omitted variables in the model is quantified using the robustness value (Cinelli and Hazlett, 2020). The robustness value is the partial R^2 that an omitted variable would need to have with both temperature variability and nightlights to reduce the estimated coefficient of temperature variability to zero. I also quantify the robustness value that would make the estimated coefficients insignificant (at $\alpha = 0.05$). The robustness values of day-to-day, seasonal, and interannual variability are shown in Table A.10. Furthermore, the table shows the partial R^2 of the control variables of my model as benchmarks and the critical values for those variables for which the robustness value is exceeded by one of the two partial R^2 .

For example, to make my estimated coefficient of day-to-day variability insignificant, a variable that was added to the model would need to have a partial R^2 of at least 2.71 (robustness value) with both out dependent variable (nightlights) and my treatment variable (day-to-day temperature variability). To reduce the estimated coefficient to zero, the robustness value is 2.99. To interpret these values, I can use my climatic and geographic variables that are already included in the model as benchmarks. The first section of the table shows that there are only three variables (annual mean temperature, annual mean relative humidity, and annual mean solar radiation) that have a partial $R^2 \geq 2.71$ (robustness value for significance) for day-to-day variability (Column 2). For these three variables the partial R^2 with day-to-day variability exceeds the robustness value, which means that even a partial R^2 with nightlights smaller than the robustness value could make my estimated coefficients zero/insignificant (if these variables were an omitted confounder). I therefore quantify the critical value for the partial R^2 of these variables with nightlights, which are shown in the bottom section of the table. The critical values are 0.77, 0.80, and 1.76 respectively. I can see in the top part of the table (Column 1) that these critical values are not exceeded by the corresponding partial R^2 .

Also the results for seasonal and interannual variability are reassuring. I find that four and one of my control variables respectively explain enough of the variation of temperature variability to be able to render the estimated coefficient insignificant (Columns 4, 6, and 8) but for none of these variables my model has a strong enough association with both nightlights and my temper-

ature variability variables to make my estimated coefficients of temperature variability insignificant (Columns 3, 5, and 7).

The benchmarking of this sensitivity analysis can also be interpreted as a balancing test. Specifically, Columns 2 and 4 in Table A.10 show the partial R^2 of all covariates for regressions on my treatment variables, day-to-day (Column 2), seasonal variability (Column 4) and interannual variability (Columns 6 and 8). The results suggest that no single covariate can explain more than 10 percent of the residual variation of day-to-day variability, 4 percent for seasonal variability, or 6 percent of the residual variation of interannual variability.

Table A.10: Results of a sensitivity analysis as proposed by Cinelli and Hazlett (2020).

Variable	Day-to-day variability		Seasonal variability		Interann. var. ($\bar{T} < 20$)		Interann. var. ($\bar{T} \geq 20$)	
	$R^2_{Y-X_j \sigma^d, X_{-j}}$	$R^2_{\sigma^d-X_j X_{-j}}$	$R^2_{Y-X_j \sigma^m, X_{-j}}$	$R^2_{\sigma^m-X_j X_{-j}}$	$R^2_{Y-X_j \sigma_A^y, X_{-j}}$	$R^2_{\sigma_A^y-X_j X_{-j}}$	$R^2_{Y-X_j \sigma_B^y, X_{-j}}$	$R^2_{\sigma_B^y-X_j X_{-j}}$
Annual mean temperature	0.04115	9.52375	0.04115	1.19277	0.04915	0.97199	0.03630	0.11316
Annual total precipitation	0.00305	0.23419	0.00305	0.09516	0.00411	0.37066	0.00175	0.14781
Annual mean rel. hum.	0.00614	9.25113	0.00614	3.33312	0.01823	1.73334	0.01449	5.28454
Solar rad. annual mean	0.06589	4.59948	0.06589	0.67956	0.06872	0.21369	0.10980	0.68666
Elevation	0.41257	2.64177	0.41257	3.18982	0.55558	0.06992	0.21424	0.17217
Ruggedness	0.27404	0.09342	0.27404	0.19660	0.27126	0.16736	0.29237	0.03908
Distance from nearest coast	0.02488	0.06632	0.02488	2.02610	0.02687	0.43462	0.03704	0.35219
Distance from nearest lake/river	0.24401	0.29917	0.24401	0.20186	0.43069	0.02648	0.17840	0.01711
<i>Robustness values</i>	2.99583		1.03060		1.27824		0.94189	
<i>Critical values</i>								
Annual mean temperature	0.87896		0.88902					
Annual mean rel. hum.	0.90759		0.31125		0.93829		0.16052	
Solar rad. annual mean	1.91905							
Elevation			0.32571					
Distance from nearest coast			0.51895					
<i>Robustness values (significance)</i>	2.71166		0.74064		0.91828		0.45353	
<i>Critical values (significance)</i>								
Annual mean temperature	0.76793		0.77734					
Annual mean rel. hum.	0.79475		0.24619		0.82353		0.11478	
Solar rad. annual mean	1.75567							
Elevation			0.25908					
Distance from nearest coast			0.43408					

Notes: The table shows the results of a sensitivity analysis using a model with spatial first differences pooled in North-South and West-East directions. Standard errors in parentheses. See text for explanation and an example of how to read the table. All values are shown in percent.

A.9 Control variables

Table A.11: Results of models estimated with different control variables.

Dependent variable:	<i>log Nightlight density</i>					
Spatial first differences:	Pooled	WE	NS	Pooled	Pooled	Pooled
Column	1	2	3	4	5	6
Day-to-day variab. of T	-0.50448*** (0.12930)	-0.69073*** (0.15170)	-0.41483*** (0.11675)	-0.60062*** (0.15917)	-0.55172*** (0.13816)	-1.03221*** (0.20010)
Seasonal variab. of T	-0.28016 (0.17127)	-0.14383 (0.17498)	-0.32325* (0.17247)	-0.22194 (0.15297)	-0.23241* (0.13572)	0.44984*** (0.08412)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17369*** (0.04441)	0.17134*** (0.05661)	0.16404*** (0.04454)	0.19419*** (0.04814)	0.24093*** (0.03951)	0.24093*** (0.05338)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.25220** (0.10077)	-0.23005** (0.09922)	-0.25865** (0.11796)	-0.19005** (0.08443)	-0.18553** (0.09012)	-0.10075 (0.09327)
Climate controls (linear)	x	x	x	x		
Climate controls (quadratic)	x	x	x			
Climate controls (linear in bins)					x	
Geographic controls (linear)	x	x	x	x		
Geographic controls (quadratic)	x	x	x			
Geographic controls (linear in bins)					x	
R2	0.0249	0.0250	0.0254	0.0208	0.0312	0.0051
df	448877	224426	224425	448887	448800	448904

Notes: The table shows the results of a model with linear terms for day-to-day and seasonal variability and an interaction term for interannual variability (Equation 1.5) estimated with spatial first-differences. Standard errors in parentheses. WE = West-East, NS = North-South. Pooled = pooling differences in WE and NS.

A.10 Urban versus rural areas

Table A.12: Results of regressions with subsampling observations based on population density.

Dependent variable:	<i>log Nightlight density</i>		
	> 95	80-95	< 50
Population density (percentiles):			
Spatial first differences:	Pooled	Pooled	Pooled
Column:	1	2	3
Day-to-day variab. of T	-0.91061 (0.76166)	-0.20431 (0.21245)	-0.24728*** (0.08610)
Seasonal variab. of T	-2.77649* (1.46725)	-0.95826*** (0.33788)	-0.16781* (0.10051)
Interann. variab. of $T * \delta(\bar{T} < 20)$	1.25644 (0.91921)	0.23916* (0.13205)	0.05253** (0.02431)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.24114 (0.86465)	-0.02971 (0.18254)	-0.17792** (0.08717)
Climate controls (linear)	x	x	x
Climate controls (quadratic)	x	x	x
Geographic controls (linear)	x	x	x
Geographic controls (quadratic)	x	x	x
R2	0.0926	0.0181	0.0228
df	6220	23428	293981

Notes: The table shows the results of a model with linear terms for day-to-day and seasonal variability and an interaction term for interannual variability (Equation 1.5) estimated with spatial first-differences. Grid cells are sampled based on their ranking (percentile) in terms of population density of the within-country distribution of grid cells. For example, Column 1 shows results for a model which includes only grid cells that are among the 5 percent most densely populated grid cells of the corresponding country. Because I use spatial first-differences, I require that grid cells and their neighbours to the West and North must fulfill this requirements.

A.11 Agricultural land use

Table A.13: Results of regressions to examine the role of agriculture.

Dependent variable: Spatial first differences: Column:	<i>log Nightlight density</i>		
	Pooled 1	Pooled 2	Pooled 3
Day-to-day variab. of T	-0.50895*** (0.12933)	-0.51023*** (0.12783)	-0.50749*** (0.12773)
Seasonal variab. of T	-0.30296* (0.18080)	-0.31293* (0.18653)	-0.31289* (0.18356)
Interann. variab. of $T * \delta(\bar{T} < 20)$	0.17643*** (0.04819)	0.19750*** (0.05019)	0.19352*** (0.05023)
Interann. variab. of $T * \delta(\bar{T} \geq 20)$	-0.24130** (0.10101)	-0.23328** (0.10163)	-0.23376** (0.10149)
Share of cropland	0.06278* (0.03758)		0.05787 (0.03868)
Share of pasture		-0.07376*** (0.02069)	-0.06962*** (0.02169)
Climate controls (linear)	x	x	x
Climate controls (quadratic)	x	x	x
Geographic controls (linear)	x	x	x
Geographic controls (quadratic)	x	x	x
R2	0.0260	0.0263	0.0268
df	445969	445969	445968


Notes: The table shows the results of a model with linear terms for day-to-day and seasonal variability and an interaction term for interannual variability (Equation 1.5) estimated with spatial first-differences. Standard errors in parentheses.

Appendix B

Appendix of Chapter 2

B.1 Detrending of time-series

Figure B.11: Time-series of seasonal production: winter (W) and summer (S) for the USA (left) and Brazil (right).



Notes: Note that there is no clear ordering of summer and winter within a calendar year. The order chosen here for visualisation (winter, summer) is arbitrary and does not affect any of the results.

B.2 Summer peak vs winter peak countries

Table B.21: Results of balancing tests for summer-peak and winter-peak countries.

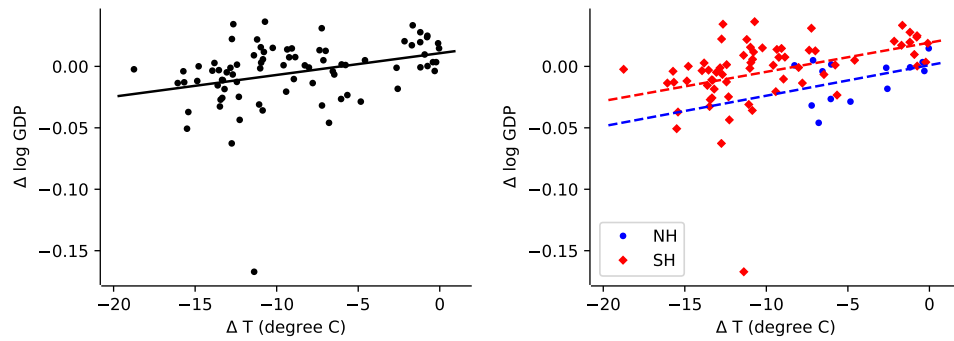
Peak of temperature in:		W		S		p-value
Variable	Unit	Mean	Std.	Mean	Std.	
$\Delta \log \text{GDP}$	USD 2010	0.01	0.01	-0.01	0.03	0.000
ΔT	deg. C	-7.03	4.59	-9.84	4.88	0.011
ΔP	mm day-1	-0.03	0.05	-0.02	0.06	0.315
Annual mean temperature	deg. C	16.38	7.08	14.63	6.23	0.260
Share of agriculture in GDP	percent	4.81	5.33	7.04	6.05	0.090
Share of manufacturing of GDP	percent	14.75	4.90	14.36	6.12	0.762
Share of exports of GDP	percent	44.58	37.11	44.71	29.94	0.987
Share of imports of GDP	percent	42.31	31.25	50.41	27.72	0.239
Share of tourism receipts of GDP	percent	9.18	11.07	14.95	12.62	0.036
Share of tourism expenditures of GDP	percent	6.52	2.41	6.32	3.95	0.778
Real interest rate	percent	7.33	8.02	4.56	5.15	0.096
Share of Christian population	percent	64.04	32.18	63.48	32.95	0.940
Share of Muslim population	percent	11.82	24.36	13.74	27.22	0.745
$\log \text{GDP per capita}$	USD 2010	10.02	0.85	9.61	0.73	0.030
Land area	1E6 km ²	11.96	2.19	11.48	2.28	0.358
Latitude	degrees	22.15	31.34	32.51	24.82	0.120
Latitude, absolute value	degrees	32.36	20.22	38.01	14.79	0.177

Notes: There are 43 summer-peak countries (S) and 34 winter-peak countries in the sample.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.3 Descriptive evidence on statistical associations

Figure B.31: Scatter plot of seasonal differences of temperature and log GDP.



Notes: Colors indicate split of sample into countries in the Northern hemisphere (NH) and Southern hemisphere (SH).

B.4 Estimation results to explore alternative channels

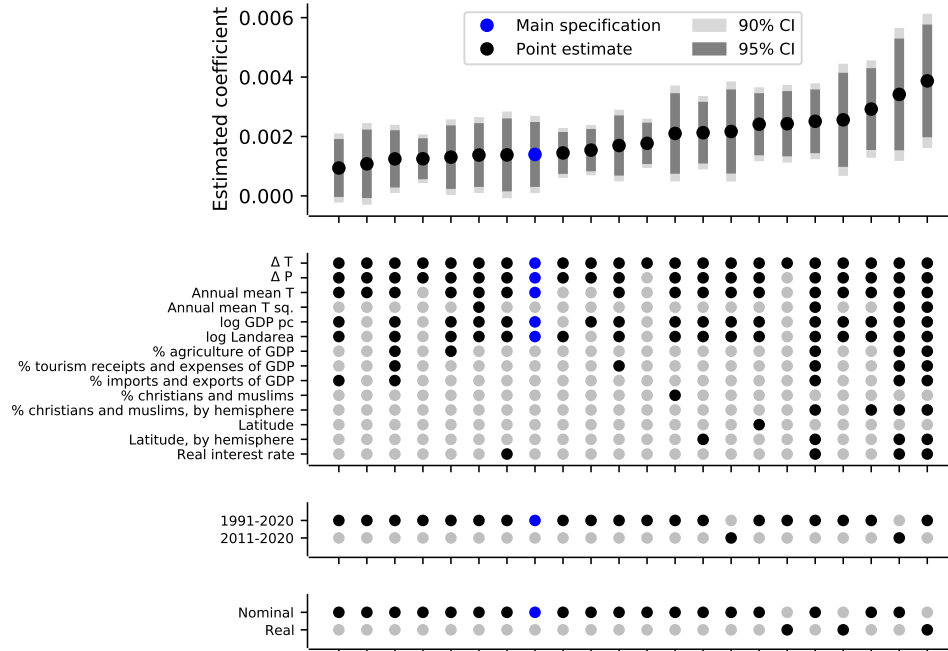
Table B.41: Results of regressions to explore possible channels through which temperature variability might affect seasonal economic cycles.

Dependent variable:	log GDP pc	% agriculture	% imports	% exports	% tourism exp.	% tourism rec.
Column:	1	2	3	4	5	6
ΔT	-0.0152 (0.0383)	0.0945 (0.1254)	-0.7797 (0.6696)	-0.5687 (0.7321)	-0.1271 (0.1459)	0.0798 (0.2524)
Δ Precipitation	2.0623 (2.2562)	21.0396** (9.8009)	-102.0344* (57.8043)	-133.4057** (60.5778)	16.5967* (9.7562)	59.0979** (28.9485)
Annual mean temperature	-0.0187 (0.0275)	0.0189 (0.0821)	-0.1567 (0.6110)	0.0130 (0.6476)	0.1399 (0.1297)	0.2440 (0.1736)
log GDP pc		-5.9706*** (0.6544)	5.1230 (3.5792)	16.1331*** (4.1690)	0.5443 (0.6466)	-5.4378*** (1.4660)
log Landarea	-0.0362 (0.0436)	0.2422 (0.1663)	-9.6960*** (1.3761)	-8.7134*** (1.5763)	0.0582 (0.1698)	-1.8469*** (0.6243)
R2	0.10	0.73	0.59	0.56	0.09	0.33
R2 adj.	0.05	0.71	0.57	0.53	0.03	0.29
N	81	81	81	81	81	81

Notes: Percentages of total GDP. Exp. = expenditures, rec. = receipts. Seasonal differences Δ calculated as winter minus summer. Significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.5 Specification chart

Figure B.51: Specification chart.



Notes: The figure shows the central estimates and 95% and 90% confidence intervals of the estimated effect of ΔT on $\Delta \log GDP$ for models with different explanatory variables (top panel), different underlying data (central panel), and different time periods (bottom panel).

B.6 Results by industry group

Table B.61: Results of regressions using a sample of GVA by industry groups of 35 European economies.

Variable	TOTAL	A	B-E	C	F	G-I	J	K	L	M-N	O-Q	R-U
ΔT	0.0037** (0.0018)	-0.0010 (0.0204)	0.0063 (0.0039)	0.0034 (0.0032)	0.0072 (0.0046)	0.0062 (0.0043)	-0.0038 (0.0023)	0.0008 (0.0033)	0.0009 (0.0015)	-0.0062** (0.0027)	-0.0002 (0.0028)	0.0008 (0.0035)
Δ Precipitation	-0.0807 (0.1565)	0.4764 (1.4129)	0.0677 (0.2710)	-0.0718 (0.2555)	-0.1152 (0.7347)	-0.0157 (0.4187)	0.5887*** (0.1942)	-0.6512 (0.3949)	0.3325 (0.2040)	0.6309** (0.2302)	0.0927 (0.3072)	0.3128 (0.4650)
Annual mean temperature	-0.0022* (0.0012)	0.0068 (0.0127)	-0.0051** (0.0022)	0.0004 (0.0022)	0.0050 (0.0050)	-0.0061** (0.0028)	-0.0028 (0.0017)	0.0033 (0.0034)	-0.0023 (0.0016)	0.0010 (0.0016)	-0.0028 (0.0029)	-0.0019 (0.0030)
log GDP pc	0.0179* (0.0088)	0.1668** (0.0794)	0.0052 (0.0137)	0.0294* (0.0152)	0.0190 (0.0247)	0.0193 (0.0182)	0.0220*** (0.0070)	-0.0141 (0.0128)	-0.0033 (0.0064)	0.0399*** (0.0105)	-0.0169 (0.0166)	-0.0153 (0.0173)
log Landarea	0.0052*** (0.0016)	0.0070 (0.0180)	0.0076*** (0.0026)	0.0094*** (0.0026)	0.0052 (0.0051)	0.0076* (0.0038)	0.0037** (0.0015)	0.0040 (0.0035)	-0.0037 (0.0023)	0.0067*** (0.0019)	0.0025 (0.0034)	0.0012 (0.0044)
R2	0.54	0.19	0.41	0.42	0.22	0.31	0.40	0.12	0.15	0.46	0.12	0.06
R2 adj.	0.46	0.05	0.31	0.32	0.08	0.19	0.30	-0.03	-0.00	0.37	-0.04	-0.11
N	35	35	35	35	35	35	35	35	35	35	35	35

Notes: Seasonal differences Δ calculated as winter minus summer. A: Agriculture, forestry and fishing, B-E: Industry (except construction), C: Manufacturing, F: Construction, G-I: Wholesale and retail trade, transport, accommo., J: Information and communication, K: Financial and insurance activities, L: Real estate activities, M-N: Professional, scientific and technical activit., O-Q: Public administration, defence, education, hum., R-U: Arts, entertainment and recreation; other serv. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.7 Estimation results for different samples

Table B.71: Results of regressions using the global samples of GDP of countries that are used for the cross-sectional and the long differences estimation.

Dependent variable:	$\Delta \log GDP$	
	Cross-section	Long differences
Sample used for:		
Column:	1	2
ΔT	0.0105** (0.0050)	0.0061 (0.0057)
$\Delta T \cdot \log GDP$ pc	-0.0010* (0.0005)	-0.0005 (0.0006)
Δ Precipitation	-0.1787** (0.0835)	-0.1365*** (0.0401)
Annual mean temperature	0.0005 (0.0004)	-0.0000 (0.0005)
$\log GDP$ pc	0.0031 (0.0058)	0.0077 (0.0069)
\log Landarea	0.0020* (0.0011)	0.0018* (0.0009)
R2	0.36	0.40
R2 adj.	0.30	0.34
N	81	60

Notes: Sample of cross-sectional regression is the same as in Table 2.2. Sample of long differences estimation is the same as in Table 2.4. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.8 Robustness check with seasons based on quarters


Table B.81: Results of regressions using a global sample of GDP of countries using differences between quarter with maximum and minimum temperature for identification.

Dependent variable: Column:	$\Delta \log GDP$		
	1	2	3
ΔT	0.0042*** (0.0008)	0.0045*** (0.0009)	0.0237*** (0.0077)
$\Delta T \cdot \log GDP$ pc			-0.0021*** (0.0008)
Δ Precipitation		-0.2406* (0.1364)	-0.2256* (0.1351)
Annual mean temperature		-0.0003 (0.0007)	-0.0005 (0.0007)
$\log GDP$ pc		-0.0218*** (0.0076)	0.0003 (0.0074)
$\log Landarea$		-0.0056** (0.0022)	-0.0048** (0.0021)
R2	0.23	0.43	0.47
R2 adj.	0.23	0.40	0.43
N	81	81	81

Notes: Time period is 1991-2020. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B.9 Past climate trends

Figure B.91: Change in seasonal temperatures between 1981-2000 and 2001-2020 for winter and summer months.



B.10 Estimation results by season

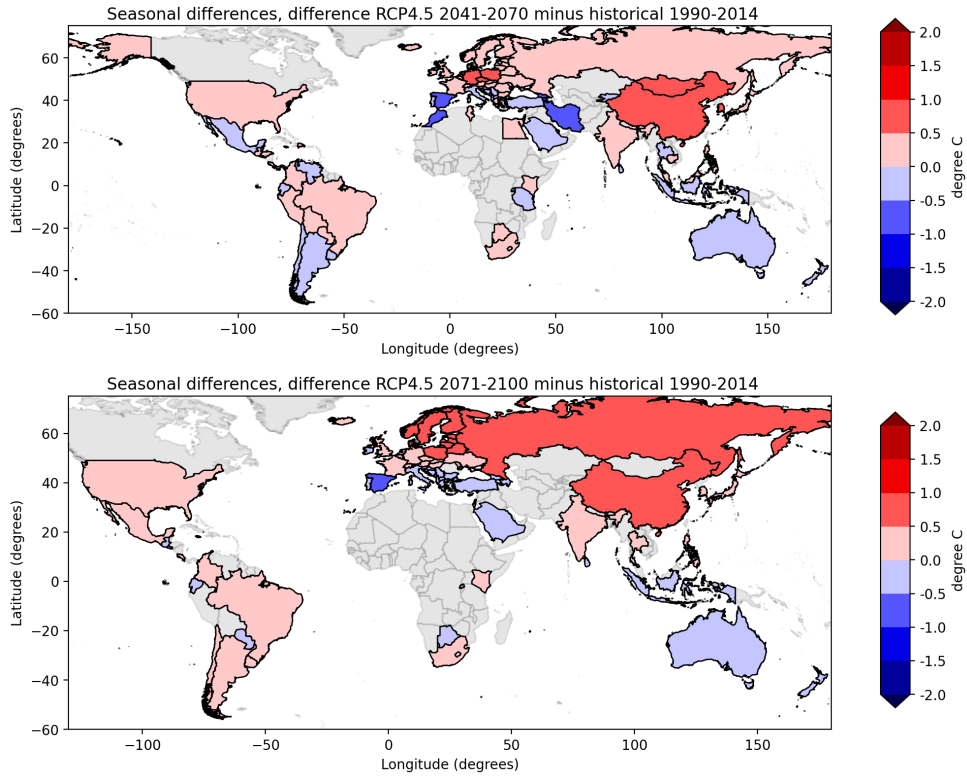
Table B.101: Results of regressions using a global sample of GDP of 60 countries, by season.

Dependent variable: Column:	$\Delta_{LD} \log GDP$	
	1	2
Δ_{LD} Temperature in summer	-0.0005 (0.0007)	0.0161 (0.0110)
Δ_{LD} Temperature in winter	0.0003 (0.0004)	-0.0082 (0.0053)
Δ_{LD} Temperature in summer · log GDP pc		-0.0016 (0.0011)
Δ_{LD} Temperature in winter · log GDP pc		0.0008 (0.0005)
Δ_{LD} Precipitation in summer	-0.0213 (0.0228)	-0.0017 (0.0140)
Δ_{LD} Precipitation in winter	-0.0312 (0.0332)	-0.0384 (0.0259)
log GDP pc		0.2352 (0.1724)
R2	0.05	0.14
R2 adj.	-0.02	0.03
N	60	60

Notes: Long differences Δ_{LD} calculated by subtracting mean over 1981-2000 from mean over 2001-2020. Seasonal differences Δ calculated as winter minus summer. Significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

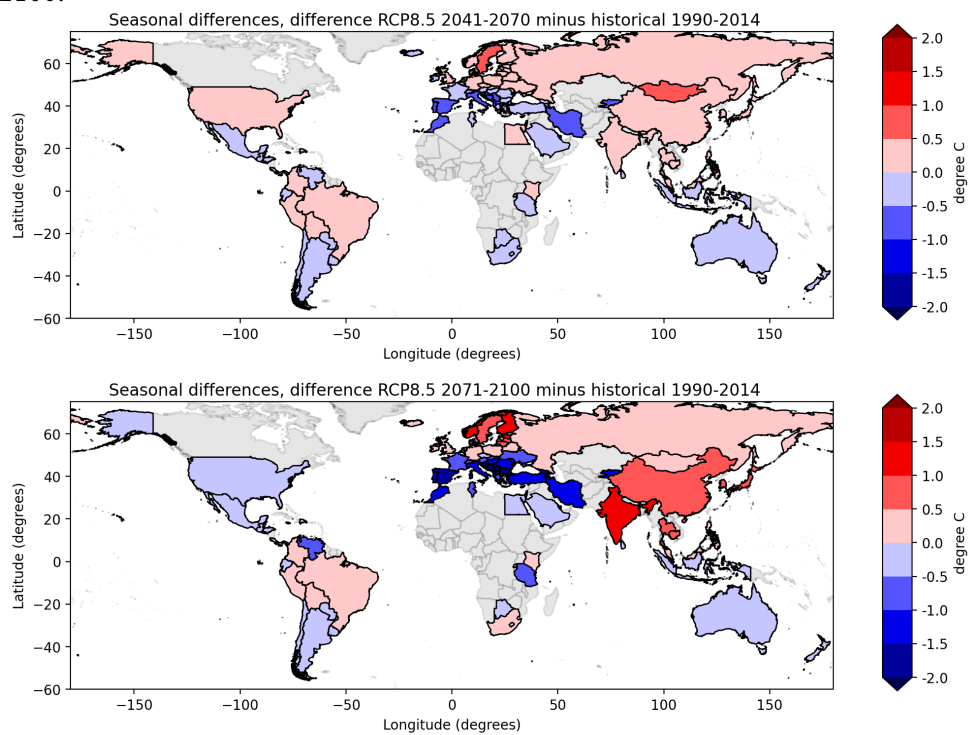
B.11 Climate projections

Figure B.111: Future projections of ΔT for RCP4.5 for 2041-2070 and 2071-2100.



Notes: Seasonal differences calculated as winter minus summer. Positive values mean that temperature in winter is projected to increase more than temperature in summer.


Figure B.112: Future projections of ΔT for RCP8.5 for 2041-2070 and 2071-2100.



Notes: Seasonal differences calculated as winter minus summer. Positive values mean that temperature in winter is projected to increase more than temperature in summer.

B.12 Projected changes to seasonal economic cycles for RCP4.5 and RCP8.5

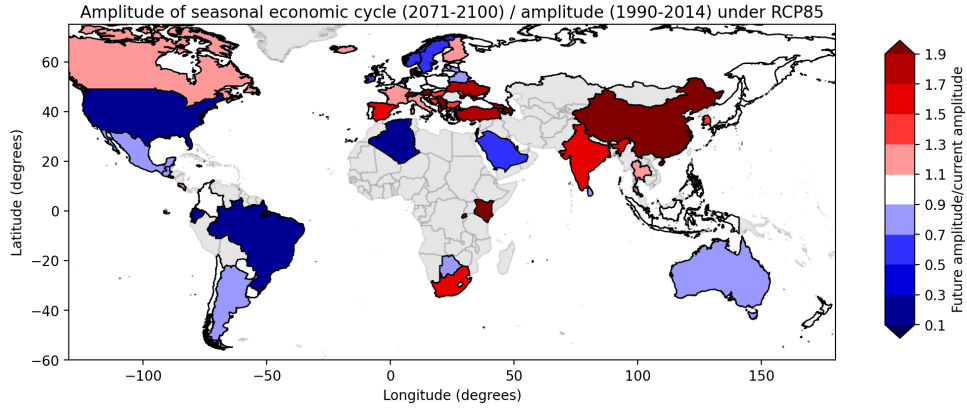
Figure B.121: Projections of Δ GDP for two alternative scenarios of future climate change for the global sample.



Notes: The plot shows the distribution of the projected changes of Δ log GDP for individual countries in the vertical direction (violinplots) based on the results from the long-difference estimation. Horizontal bars indicate the maximum, median, and minimum values. Seasonal differences Δ calculated as winter minus summer. Positive values mean that for the given scenario GDP will be reallocated from summer to winter.

B.13 Map of projected changes to seasonal economic cycles

Figure B.131: Map of projections of Δ GDP for the RCP8.5 scenario of future climate change for the global sample.



Notes: The map shows the distribution of the projected changes of $\Delta \log$ GDP for the RCP 8.5 scenarios based on the results from the long-difference estimation. Seasonal differences Δ calculated as winter minus summer. Positive values mean that for the given scenario GDP will be reallocated from summer to winter.

Appendix C

Appendix of Chapter 3

C.1 Data

Figure C.11: Spatial distribution of annual mean temperature.

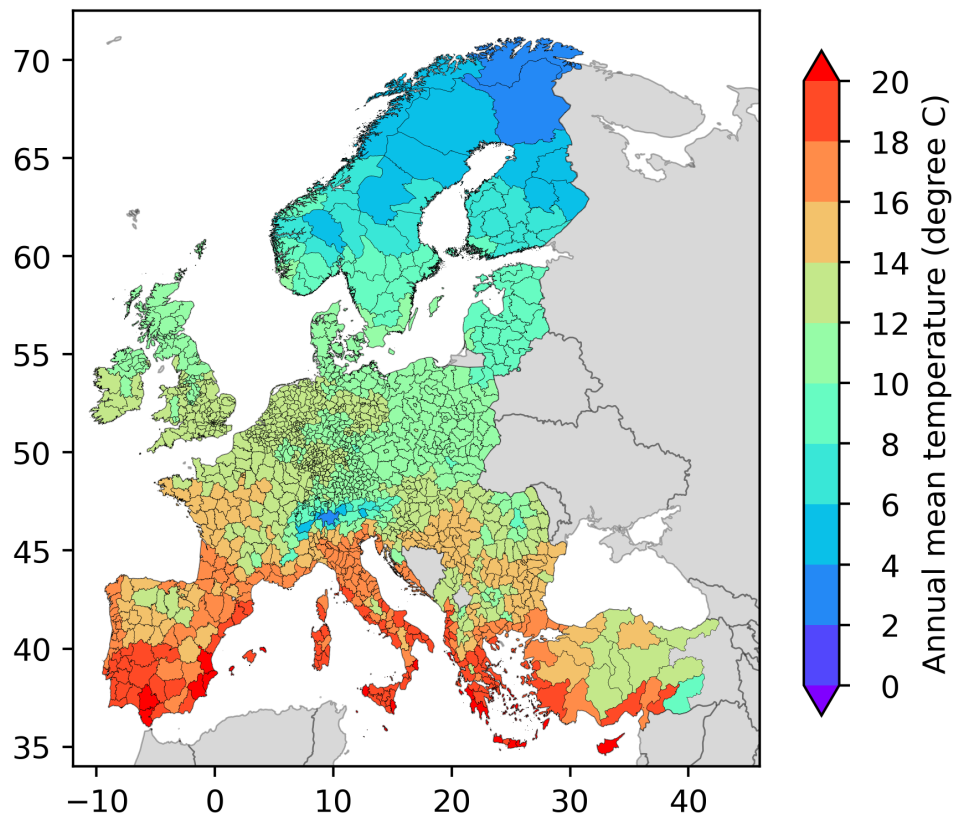
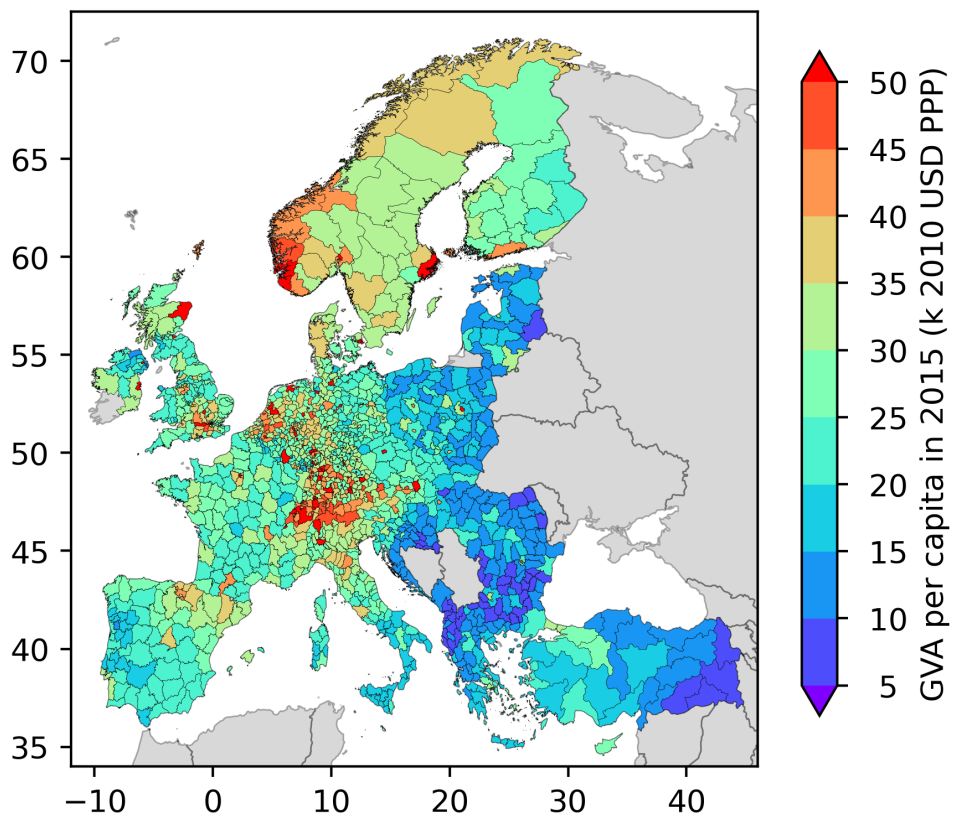


Figure C.12: Spatial distribution of Gross Value Added in Europe in 2015 (2010 USD PPP).



<i>EUROSTAT industry groups and NACE2 sections</i>			
Industry group	Median share	Group	Sections
Agriculture	0.02	A	A: Agriculture, Forestry, and Fishing
Construction	0.06	F	F: Construction
Manufacturing	0.18	C	C: Manufacturing
Mining and Utilities	0.22	B-E	B: Mining and Quarrying D: Electricity, Gas, Steam, and Air Conditioning Supply E: Water Supply; Sewerage, Waste Management and Remediation Activities
Trade	0.19	G-I	G: Wholesale and Retail Trade; Repair of Motor Vehicles and Motorcycles H: Transportation and Storage I: Accommodation and Food Service Activities
Other Services	0.51	J K L M-N O-Q P Q R-U	J: Information and Communication K: Financial and Insurance Activities L: Real Estate Activities M: Professional, Scientific and Technical Activities N: Administrative and Support Service Activities O: Public Administration and Defence; Compulsory Social Security P: Education Q: Human Health and Social Work Activities R: Arts, Entertainment and Recreation S: Other Service Activities T: Activities of Households as Employers; Production activities of Households for own use U: Activities of Extraterritorial Organisations and Bodies

Table C.11: Industry groups used in the analysis.

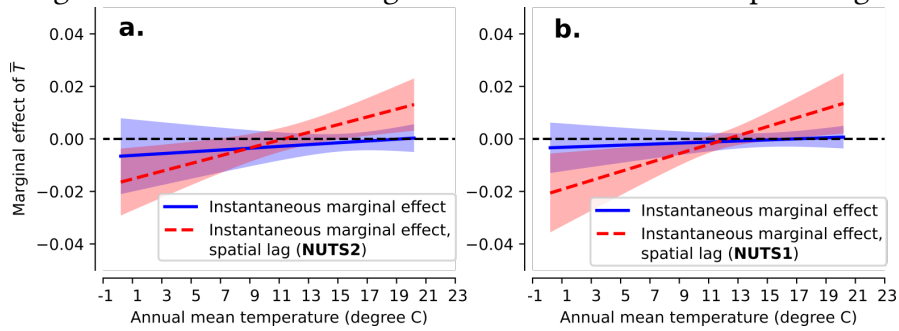
Variable	Unit	Mean	Std.	Min.	Max.	No. obs.
Population density	km-2	561.11	1384.38	1.94	21844.42	25911
Gross value added per capita	2010 USD PPP	27.72	17.08	3.35	440.92	22354
Annual mean temperature	deg C	10.99	2.76	-0.18	20.43	28413
Seasonal mean temperature DJF	deg C	2.83	3.81	-12.59	16.03	28413
Seasonal mean temperature MAM	deg C	10.18	2.70	-2.94	19.10	28413
Seasonal mean temperature JJA	deg C	19.15	3.07	10.62	29.29	28413
Seasonal mean temperature SON	deg C	11.63	3.05	-1.20	23.16	28413
Annual total precipitation	m	2.03	0.73	0.23	9.48	28413

Table C.12: Descriptive statistics.

C.2 Additional results

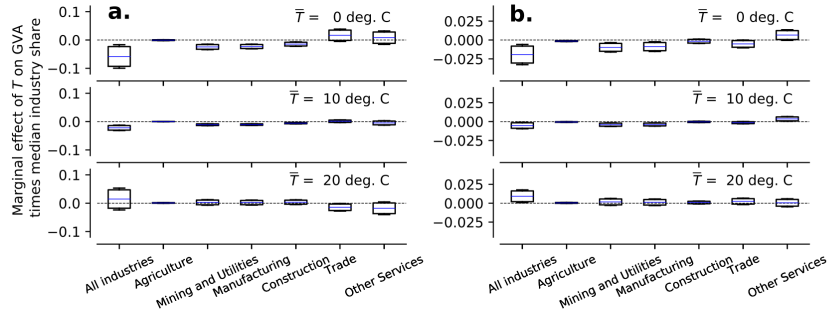
C.2.1 Spatial spill-overs

Figure C.21: Estimated marginal effects of model with spatial lag.



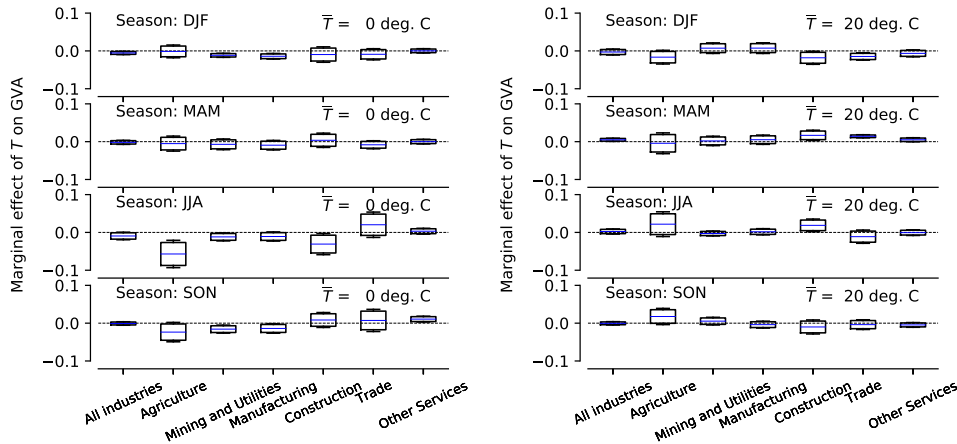
C.2.2 Weighted industry results

Figure C.22: Estimated marginal effects by industry at three different levels of temperature. Estimated coefficients are multiplied with median share of industry GVA of total GVA of whole sample. **a.** Instantaneous effect. **b.** Cumulative effect over six years.



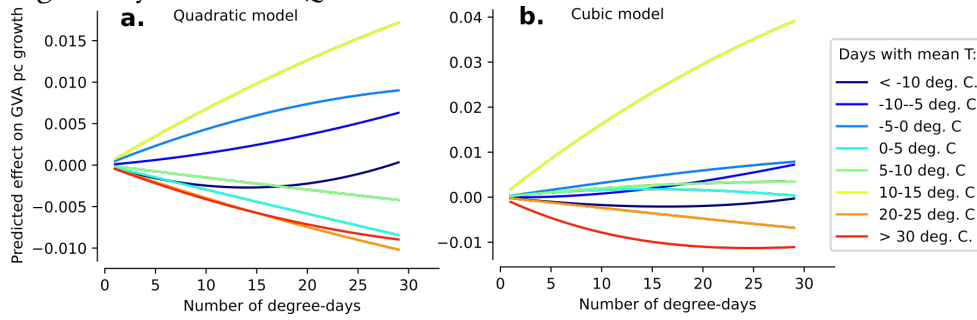
C.2.3 Seasonal effects

Figure C.23: Estimated marginal effect of seasonal temperature at annual mean temperature of 0 degree C (left) and 20 degree C (right).




C.2.4 Adaptation

Figure C.24: Curvature of predicted effects estimated with polynomial degree-day models. **a.** Quadratic models. **b.** Cubic models.



C.2.5 Degree-day model

Figure C.25: Estimated effects from degree-day model.



C.3 Additional robustness checks

C.3.1 Sample of countries and years

Dependent variable: <i>diff log GDP per capita</i>				
Sample:	1950 -	1997 -	1950 -	1997 -
Time period:	World	World	Europe	Europe
Column:	1	2	3	4
T.mean	0.24757*** (0.06328)	0.12241*** (0.04400)	-0.13799* (0.06921)	-0.30506** (0.13067)
T.mean.sq	-0.00043*** (0.00011)	-0.00022*** (0.00008)	0.00024* (0.00012)	0.00053** (0.00023)
P.mean	-0.05238 (0.04250)	-0.04239 (0.05072)	0.03089 (0.09619)	0.06062 (0.07630)
R2	0.2162	0.3655	0.4337	0.5510
df	6338	3503	1110	662

Table C.31: Results of regressions with different samples of countries and years using the growth rate of national GDP per capita as dependent variable. See main text for explanation.


C.3.2 Alternative standard errors

Dependent variable: <i>diff log GVA per capita</i>								
SE:	Clustered by country				Conley, cutoff at ...			
	100 km		200 km		400 km		1000 km	
FE:	r, y	r, y	r, y x c	r, y x c	r, y	r, y	r, y	
Time trends:	no	linear	no	linear	no	no	no	no
Column	1	2	3	4	5	6	7	8
T.mean	-0.41167** (0.15234)	-0.52141** (0.20691)	-0.03990 (0.12926)	-0.06183 (0.13850)	-0.41167*** (0.07239)	-0.41167*** (0.09307)	-0.41167*** (0.12875)	-0.41167** (0.17218)
T.mean.2	0.00072** (0.00027)	0.00091** (0.00036)	0.00007 (0.00023)	0.00011 (0.00024)	0.00072*** (0.00013)	0.00072*** (0.00016)	0.00072*** (0.00023)	0.00072** (0.00030)
P.mean	0.00024 (0.00205)	0.00064 (0.00204)	-0.00090 (0.00132)	-0.00087 (0.00129)	0.00024 (0.00089)	0.00024 (0.00117)	0.00024 (0.00153)	0.00024 (0.00189)
R2	0.2659	0.3130	0.4628	0.4866	0.2659	0.2659	0.2659	0.2659
df	19701	18337	19258	17925	19701	19701	19701	19701

Table C.32: Results of regressions with different fixed effects and time trends (Columns 1-4) and with Conley HAC standard errors with different cutoff values (Columns 5-8). See main text for explanation.

C.3.3 Alternative weather data

Figure C.31: Estimated marginal effect of increasing annual mean temperature by one unit for subnational GVA using two alternative weather data sets: ERA reanalysis (left) and Delaware temperature data (right).



Appendix D

Appendix of Chapter 4

D.1 Causality and the empirical framework

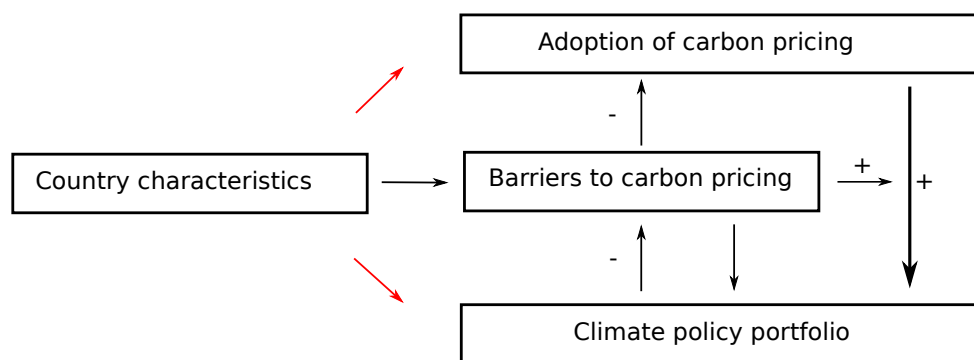


Figure D.11: **Causal diagram with possibly confounding country characteristics.** The red arrows indicate how country characteristics can confound the statistical association between the climate policy portfolio and the adoption of a carbon pricing policy due to an influential variable that is omitted in a linear regression.

D.2 Descriptive statistics

Table D.21: Descriptive statistics of country characteristics included in the regression analyses for all countries in the sample of G20 economies and other major emitters (see Table D.51) covering the years 1988-2020. Sources are listed in Section 5.2.

Variable	Unit	Mean	Std.	Min.	Max.	No. obs.
log GDP pc PPP	2010 USD	9.19	1.29	5.99	11.28	1188
Control of corruption	index	0.17	1.05	-1.60	2.15	1188
Education	index	0.65	0.16	0.20	0.94	1188
Reserves of coal	tons pc	307.01	822.58	0.00	6702.82	1188
Reserves of gas	1000 cubic metres pc	112.65	323.93	0.00	3541.63	1188
Reserves of oil	cubic metres pc	419.63	1371.79	0.00	9586.25	1188

D.3 Results from regression analysis

Table D.31: Results of a logistic regression on the adoption of carbon pricing (binary variable). Standard errors in parentheses (clustered at country level). All models are estimated with Maximum Likelihood. The variable selection of the reduced models in Column 4 is obtained with a Lasso model with shrinkage parameter $\alpha = 0.1$. Size of portfolio is measured as described in Section 5.2. Significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Dependent variable:	Adoption of carbon pricing			
	Portfolio	All	Selected	Lasso
Column:	1	2	3	4
log GDP per capita PPP		0.4495 (0.8221)	0.5640* (0.3402)	
Coal reserves		-1.2343 (1.4830)	-0.5379 (0.5035)	
Oil reserves		-1.5461 (1.1937)		
Gas reserves		-0.0085* (0.0044)	-0.0073 (0.0084)	-0.0108 (0.0157)
Control of corruption		-1.1384 (0.8139)		
Education		19.0651*** (6.0125)		
Policy sequence score	0.1072*** (0.0245)	0.0485 (0.0338)	0.0808*** (0.0221)	0.0994*** (0.0242)
Intercept	-4.4354*** (0.7497)	-20.3759** (8.1861)	-8.8369*** (3.3383)	-3.9841*** (0.7238)
Ps. R2	0.33	0.51	0.39	0.35
AIC	641.66	479.73	595.02	625.47
No. countries	36	36	36	36
N	1188	1188	1188	1188

D.4 Policy portfolios by the end of 2020

Carbon pricing and policy portfolios by end of 2020

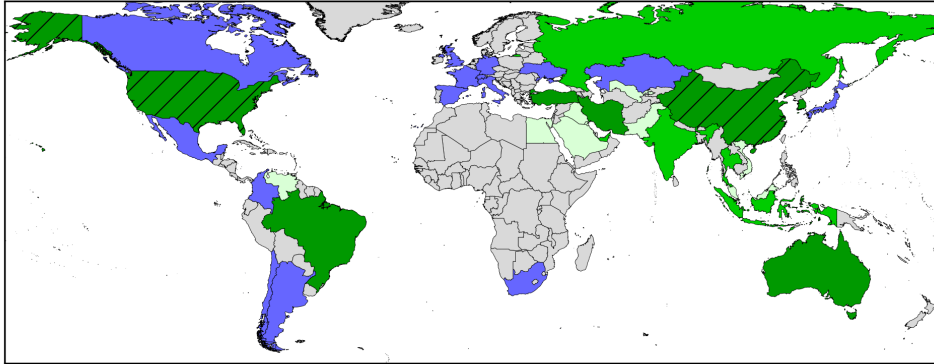



Figure D.41: **Carbon pricing policies and the number of instrument types in countries' policy portfolios.** Map shows the policy portfolios as of end of 2020. Number of instrument types only shown for countries without a carbon price at the national level.

D.5 Sample of countries

Table D.51: List of countries included in the main analysis and their carbon pricing policies by the end of 2020. Sectors: E = Electricity and Heat Production, I = Industry, B = Buildings, T = Transport, A = AFOLU.

ISO code	Sample		Carbon pricing policies						
			Year of first adoption		Sectoral coverage				
	G20	Other	National	Subnational	E	I	B	T	A
ARG	X		2018		X	X	X	X	X
AUS	X								
BRA	X								
CAN	X		2019	2007	X	X	X	X	X
CHN	X			2013					
DEU	X		2005		X	X		X	
FRA	X		2005		X	X		X	
GBR	X		2005		X	X		X	
IDN	X								
IND	X								
ITA	X		2005		X	X		X	
JPN	X		2012	2010	X	X	X	X	X
KOR	X								
MEX	X		2014		X	X	X	X	X
RUS	X								
SAU	X								
TUR	X								
USA	X			2009					
ZAF	X		2019		X	X	X	X	
ARE		X							
CHE		X	2008		X	X	X	X	
CHL		X	2017		X	X			
COL		X	2017		X	X	X	X	X
EGY		X							
ESP		X	2005		X	X		X	
IRN		X							
IRQ		X							
KAZ		X	2013		X	X	X		
KWT		X							
MYS		X							
NIG		X							
PAK		X							
THA		X							
UKR		X	2011		X	X	X		
UZB		X							
VEN		X							
VNM		X							

Table D.52: List of countries not included in the main analysis because of insufficient data on policy adoption, but which had a national carbon price implemented by the end of 2020.

ISO code	EU ETS	Carbon pricing policies							
		Year of first adoption		Sectoral coverage					
		National	Subnational	E	I	B	T	A	
SGP		2019		X	X				
HRV	Yes	2013		X	X		X		
ISL	Yes	2013		X	X	X	X	X	
LIE	Yes	2008		X	X	X	X		
NZL		2008		X	X		X	X	
BGR	Yes	2007		X	X		X		
ROU	Yes	2007		X	X		X		
AUT	Yes	2005		X	X		X		
BEL	Yes	2005		X	X		X		
CYP	Yes	2005		X	X		X		
CZE	Yes	2005		X	X		X		
EST	Yes	2005		X	X		X		
GRC	Yes	2005		X	X		X		
HUN	Yes	2005		X	X		X		
IRL	Yes	2005		X	X		X		
LTU	Yes	2005		X	X		X		
LUX	Yes	2005		X	X		X		
MLT	Yes	2005		X	X		X		
NLD	Yes	2005		X	X		X		
PRT	Yes	2005		X	X		X		
SVK	Yes	2005		X	X		X		
LVA	Yes	2004		X	X				
SVN	Yes	1996					X	X	
DNK	Yes	1992					X	X	
NOR	Yes	1991		X	X	X	X	X	X
SWE	Yes	1991					X	X	
FIN	Yes	1990			X	X	X		
POL	Yes	1990		X	X	X	X	X	X

D.6 Instrument types and sectors




Figure D.61: **Frequency of instrument types in different sectors.** Heatmap is based on all policies adopted by countries in the sample. Notes: AFOLU = Agriculture, Forestry, and other Land Use.

D.7 Instrument types and instrument categories

Table D.71: List of instrument categories, their instrument type, and their frequency in the sample of 37 countries.

Instrument type	Instrument category	Number
Grants, subsidies, other fin. incentives	Grants and subsidies	342
Grants, subsidies, other fin. incentives	Tax relief	191
Grants, subsidies, other fin. incentives	Feed-in tariffs or premiums	128
Grants, subsidies, other fin. incentives	Loans	96
Grants, subsidies, other fin. incentives	Fiscal or financial incentives (other)	91
Grants, subsidies, other fin. incentives	Energy and other taxes	74
Grants, subsidies, other fin. incentives	GHG emission reduction crediting and offsetting mechanism	43
Grants, subsidies, other fin. incentives	GHG emissions allowances	42
Grants, subsidies, other fin. incentives	Market-based instruments (other)	34
Grants, subsidies, other fin. incentives	Tendering schemes	28
Grants, subsidies, other fin. incentives	Net metering	19
Grants, subsidies, other fin. incentives	other CO2 taxes	19
Grants, subsidies, other fin. incentives	Economic instruments (other)	16
Grants, subsidies, other fin. incentives	Retirement premium	11
Grants, subsidies, other fin. incentives	User charges	8
Grants, subsidies, other fin. incentives	Removal of fossil fuel subsidies	1
Information and Education	Information provision	315
Information and Education	Advice or aid in implementation	172
Information and Education	Endorsement label	78
Information and Education	Comparison label	69
Information and Education	Professional training and qualification	39
Information and Education	Information and education (other)	29
Information and Education	Green certificates	27
Information and Education	Performance label (other)	14
Information and Education	White certificates	13
Information and Education	Barrier removal (other)	1

Table D.72: List of instrument categories, their instrument type, and their frequency in the sample of 37 countries (cont.).

Instrument type	Instrument category	Number
Policy Support	Strategic planning	692
Policy Support	Institutional creation	177
Policy Support	Policy support (other)	153
Policy Support	Political & non-binding climate strategy	72
Policy Support	Political & non-binding GHG reduction target	67
Policy Support	Political & non-binding renewable energy target	62
Policy Support	Formal & legally binding renewable energy target	43
Policy Support	Formal & legally binding climate strategy	40
Policy Support	Formal & legally binding GHG reduction target	39
Policy Support	Political & non-binding energy efficiency target	29
Policy Support	GHG reduction target (other)	21
Policy Support	Formal & legally binding energy efficiency target	20
Policy Support	Renewable energy target (other)	20
Policy Support	Coordinating body for climate strategy	18
Policy Support	Target (other)	12
Policy Support	Energy efficiency target (other)	7
Policy Support	Climate strategy (other)	3
Procurement and investment	Infrastructure investments	123
Procurement and investment	Procurement rules	62
Procurement and investment	Funds to sub-national governments	44
Procurement and investment	Direct investment (other)	18

Table D.73: List of instrument categories, their instrument type, and their frequency in the sample of 37 countries (cont.).

Instrument type	Instrument category	Number
Regulatory Instruments	Other mandatory requirements	192
Regulatory Instruments	Monitoring	178
Regulatory Instruments	Product standards	129
Regulatory Instruments	Sectoral standards	115
Regulatory Instruments	Regulatory Instruments (other)	100
Regulatory Instruments	Building codes and standards	99
Regulatory Instruments	Vehicle fuel-economy and emissions standards	93
Regulatory Instruments	Obligation schemes	88
Regulatory Instruments	Auditing	75
Regulatory Instruments	Codes and standards (other)	58
Regulatory Instruments	Grid access and priority for renewables	29
Regulatory Instruments	Industrial air pollution standards	3
Regulatory Instruments	Vehicle air pollution standards	2
Research, Development and Deployment	Technology deployment and diffusion	131
Research, Development and Deployment	Technology development	112
Research, Development and Deployment	Demonstration project	108
Research, Development and Deployment	RD&D funding	93
Research, Development and Deployment	Research & Development and Deployment (RD&D) (other)	72
Research, Development and Deployment	Research programme (other)	11
Voluntary Approaches	Negotiated agreements (public-private sector)	153
Voluntary Approaches	Public voluntary schemes	25
Voluntary Approaches	Voluntary approaches (other)	23
Voluntary Approaches	Unilateral commitments (private sector)	9

Appendix E

Appendix of Chapter 5

E.1 Quantification of the global benefits of diffusion

E.1.1 Back-of-the-envelope calculations

In the second step of the analysis, we use our empirical estimates to calculate the expected CO2 emission reductions that can be causally attributed to policy diffusion. We do so in two ways, first with a back-of-the-envelope calculation and then with Monte Carlo simulations.

For the back-of-the-envelope calculation, we compare two counterfactual scenarios: scenario A in which country i adopts carbon pricing in year t and scenario B in which it does not do so. For both scenarios, we calculate the hazard of policy adoption at time $t + 1$ for all countries $j \neq i$. The additional hazard that is due to policy diffusion from country i to country j can then be calculated as the difference between the hazards of the two scenarios.

Formally, for all countries $j \neq i$ we compare the two hazards (Equation 5.1)

$$h^A(t + 1, X_{j,t}, W_{j,t}^A) = h_0(t + 1) \exp(X_{j,t}\beta_X) \exp(W_{j,t}^A\beta_W) \quad (\text{E.1})$$

and

$$h^B(t + 1, X_{j,t}, W_{j,t}^B) = h_0(t + 1) \exp(X_{j,t}\beta_X) \exp(W_{j,t}^B\beta_W) \quad (\text{E.2})$$

For simplicity, we assume that in scenario B, no country has adopted the policy at time t , i.e. $Y_{j,t} = 0 \forall j$, which implies that the spatial lag is zero for

all countries, i.e. $W_{j,t}^A = 0 \forall j$ (Equation 5.3). Furthermore, we assume that after adjusting for covariates all countries j have the same baseline hazard, i.e. $h_0(t+1) \exp(X_{j,t}\beta_X) = h_0^*(t+1) \forall j$.

With these assumption, we can calculate the additional hazard in country j from policy adoption in country i as

$$\begin{aligned} \Delta h_{j,t+1} &= h^B(t+1, X_{j,t}, W_{j,t}^B) - h^A(t+1, X_{j,t}, W_{j,t}^A) \\ &= h_0^*(t+1) [\exp(W_{j,t}^B \beta_W) - 1] \end{aligned} \quad (\text{E.3})$$

Because in scenario B only country i adopts the policy, i.e. $Y_{j,t} = 0 \forall j \neq i$, we can calculate the spatial lag as (Equation 5.3)

$$W_{j,t}^B = \frac{w_{i,j,t}}{\sum_{i=1, i \neq j}^{N_c} w_{i,j,t}} \forall j \quad (\text{E.4})$$

The total indirect emission reductions due to diffusion can then be calculated as

$$R_{i,t+1}^{\text{indirect}} = r \sum_{j \neq i} \Delta h_{j,t+1} E_{j,t+1} \quad (\text{E.5})$$

where $E_{j,t}$ are the total CO2 emissions of country j in year t and r is the rate at which emissions are reduced per year. We compare these indirect emission reductions with the direct emission reductions obtained with similar assumptions

$$R_{i,t+1}^{\text{direct}} = r E_{i,t+1} \quad (\text{E.6})$$

For these calculations, we use actual CO2 emissions in the year 2019, which is the last year prior to the pandemic with Sars-CoV-2.

For the back-of-the-envelope calculations we only quantify emission reductions in year $t+1$. Subsequent emissions reductions, including those from further diffusion of the policy, are quantified with Monte Carlo simulation as described in the following.

E.1.2 Monte-Carlo simulations

The Monte Carlo simulations are based on Equations E.1 and E.2. We start the simulations in the year $t = 1988$ and assume that no country has adopted the policy prior to that. For every country i , we then conduct simulations for the same two scenarios A and B as above: in scenario A, no country adopts the policy in the year $t = 1988$. In scenario B, only country i adopts the policy at $t = 1988$.

For both scenarios, we then simulate adoption and diffusion of climate policies from the year 1989 onwards. To do so, at every timestep $1989 \leq t \leq 2021$ we update the spatial lag $W_{j,t}$ of every country, calculate its hazard of policy adoption, and use this hazard to draw from a probability distribution to determine whether the country adopts or does not adopt the policy at this timestep.

We conduct 5,000 simulations for every country for scenario B and 10,000 simulations for scenario A, which is the counterfactual of scenario B for all countries. The simulations of scenario B result for every country i in one matrix of probabilities of policy adoption of country j in year t , $P_{i,j,t}^B$ with $\sum_{t=1988}^{2021} P_{i,j,t}^B = 1 \forall i, j$. The simulations of scenario A result in another matrix $P_{j,t}^A$ that again satisfies $\sum_{t=1988}^{2021} P_{j,t}^A = 1 \forall j$. Because there is no difference in the counterfactuals, this matrix $P_{j,t}^A$ is the same for all countries i .

Based on these probabilities, for every country i we subsequently calculate the expected direct emission reductions and the expected indirect emission reductions due to policy diffusion. The indirect emission reductions again refer to emission reductions that can be attributed to the diffusion of the policy from country i to other countries and onwards. For both direct and indirect emission reductions, we use actual emission growth rates and subtract the effect of the carbon pricing policy from them. Formally, for every country i we calculate the direct emission reductions from 1988 - 2019 of implementing the policy in year 1988 as

$$\hat{R}_{i,2019}^{\text{direct}} = \sum_{t=1988}^{2019} \left[E_{i,t} - E_{i,1988} \prod_{l=1988}^t (1 + g_{i,l} - r) \right] \quad (\text{E.7})$$

where $g_{j,t}$ is the actually observed growth rate of CO2 emissions of country j in year t and r is the effectiveness of carbon pricing as in the Section above. For the indirect emission reductions that can be attributed to policy diffusion

from country i to other countries, we use the probabilities of policy adoption $P_{j,t}^A$ and $P_{j,t}^B$ of the scenarios A and B respectively. In mathematical terms, we take the difference between the expected emission reductions between the two scenarios:

$$\hat{R}_{i,2019}^{\text{indirect}} = \sum_{j \neq i} \left[\sum_{\xi=1988}^{2019} \left[(P_{i,j,\xi}^B - P_{j,\xi}^A) \left[\sum_{t=1988}^{\xi} E_{j,t} + E_{j,\xi} \prod_{l=\xi}^{2019} (1 + g_{j,l} - r) \right] \right] \right] \quad (\text{E.8})$$

E.2 Additional results




Figure E.21: Cumulative baseline hazard of the Cox proportional hazard model in Equation 5.1 with six covariates. Estimated coefficients of this model are shown in Column 2 in Table 5.2.

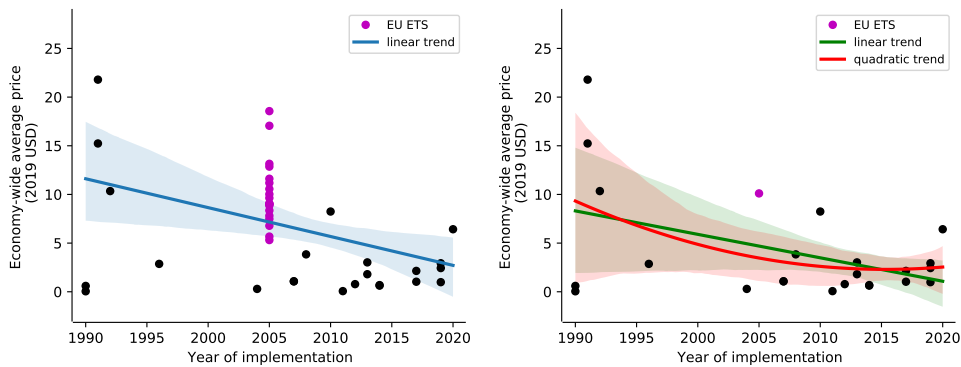


Figure E.22: Scatter plot of economy-wide emission-weighted average carbon prices over time.

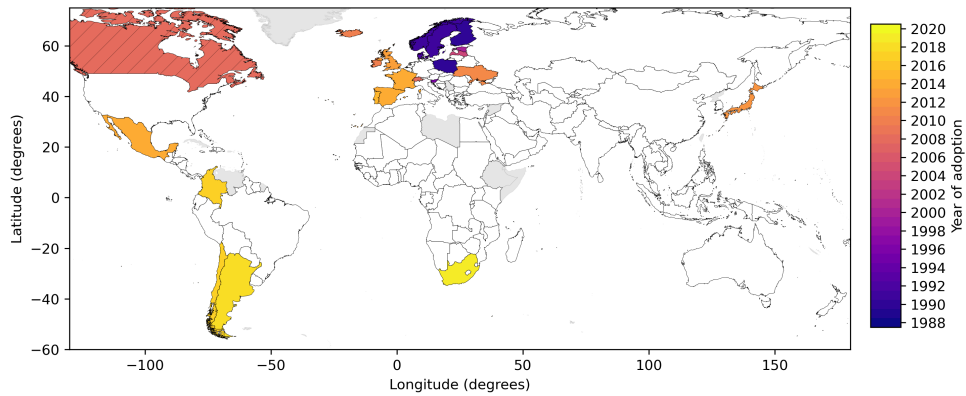


Figure E.23: Time of adoption of the first carbon tax policy by country. Hashes indicate countries in which the first policy was adopted at the subnational level.

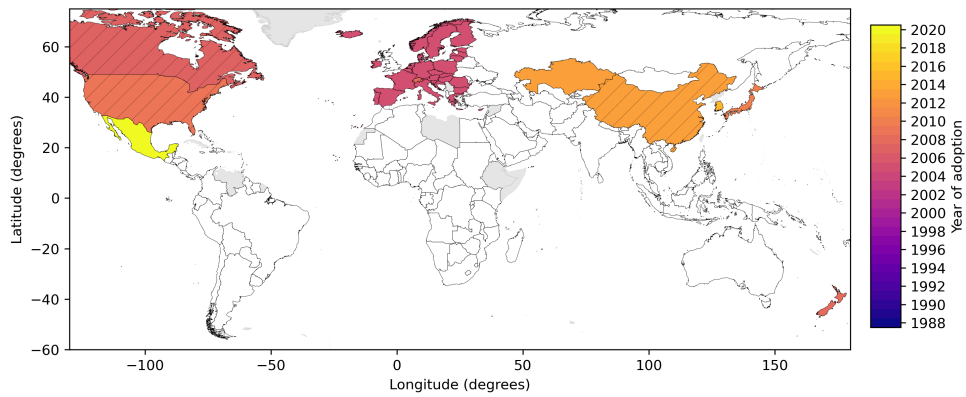


Figure E.24: Time of adoption of the first ETS policy by country. Hashes indicate countries in which the first policy was adopted at the subnational level.

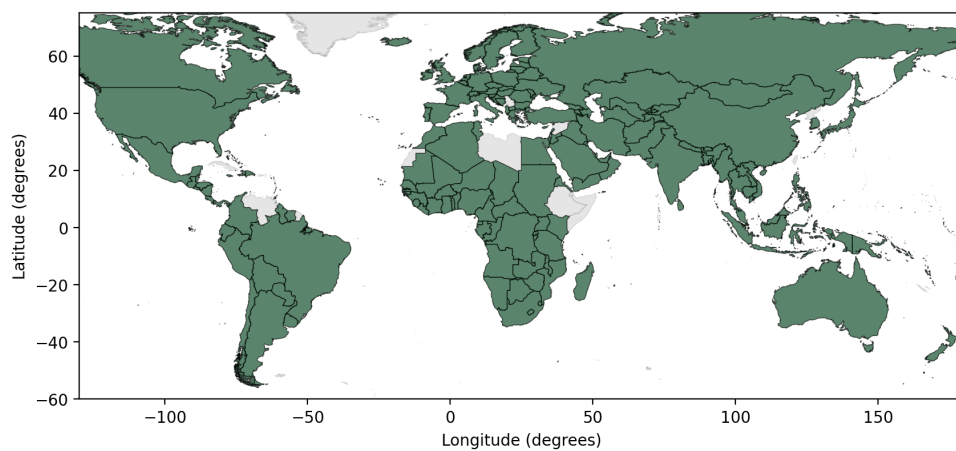


Figure E.25: Map of the sample of 179 countries used in this study.

Table E.21: Results of estimation with different lag times.

Policy:	Carbon price				
Proximity metric:	Gravity				
Lag time:	1	2	3	4	5
Column:	1	2	3	4	5
Spatial lag of carbon pricing	6.7053*** (1.3469)	6.3749*** (1.3561)	6.5309*** (1.4128)	6.5649*** (1.4821)	6.9365*** (1.5248)
GDP per capita PPP	11.5662*** (3.6787)	11.0535*** (3.5089)	11.2919*** (3.3388)	10.1004*** (3.0844)	9.7194*** (2.9444)
GDP per capita PPP sq.	-0.5529*** (0.1895)	-0.5279*** (0.1806)	-0.5368*** (0.1730)	-0.4788*** (0.1610)	-0.4623*** (0.1535)
GDP per capita PPP growth	2.0597 (3.1010)	2.0058 (4.2026)	8.1873*** (1.6269)	5.6635*** (2.1262)	3.9481 (3.8003)
Export share	-0.0054 (0.0045)	-0.0054 (0.0043)	-0.0069 (0.0044)	-0.0057 (0.0045)	-0.0042 (0.0041)
Services share of GDP	0.0271** (0.0136)	0.0251* (0.0135)	0.0230* (0.0124)	0.0219* (0.0123)	0.0206* (0.0115)
Emissions CO2 per GDP	-0.0121 (0.1140)	-0.0083 (0.1107)	-0.0049 (0.1203)	-0.0197 (0.1013)	-0.0008 (0.0976)
Time at risk	5277	5277	5277	5277	5277
log-likelihood	-177.5	-179.5	-177.3	-180.3	-182.1
AIC	369.1	373.0	368.6	374.6	378.1
N	5252	5230	5203	5174	5142

Notes: Standard errors clustered by country in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.




Figure E.26: Histogram of estimated baseline hazard adjusted for covariates in 2020 for the sample of 179 countries. Median and mean values are 0.14 and 0.32 percent, respectively. Probabilities of 0.32, 1, and 5 percent imply a cumulative probability of policy adoption by the end of a period of 30 years of 9, 26, and 79 percent, respectively.