NBER WORKING PAPER SERIES

THE MACROECONOMIC IMPACT OF CLIMATE CHANGE: GLOBAL VS. LOCAL TEMPERATURE

Adrien Bilal
Diego R. Känzig

Working Paper 32450
http://www.nber.org/papers/w32450

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2024

Adrien Bilal gratefully acknowledges support from the Chae Family Economics Research Fund at Harvard University. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Macroeconomic Impact of Climate Change: Global vs. Local Temperature
Adrien Bilal and Diego R. Känzig
NBER Working Paper No. 32450
May 2024
JEL No. E01,E23,F18,O44,Q54,Q56

ABSTRACT

This paper estimates that the macroeconomic damages from climate change are six times larger than previously thought. We exploit natural variability in global temperature and rely on time-series variation. A 1°C increase in global temperature leads to a 12% decline in world GDP. Global temperature shocks correlate much more strongly with extreme climatic events than the country-level temperature shocks commonly used in the panel literature, explaining why our estimate is substantially larger. We use our reduced-form evidence to estimate structural damage functions in a standard neoclassical growth model. Our results imply a Social Cost of Carbon of $1,056 per ton of carbon dioxide. A business-as-usual warming scenario leads to a present value welfare loss of 31%. Both are multiple orders of magnitude above previous estimates and imply that unilateral decarbonization policy is cost-effective for large countries such as the United States.

Adrien Bilal
Department of Economics
Harvard University
1805 Cambridge Street
Cambridge, MA 02138
and NBER
adrienbilal@fas.harvard.edu

Diego R. Känzig
Department of Economics
Northwestern University
Kellogg Global Hub
2211 Campus Drive
Evanston, IL 60208
and NBER
dkaenzig@northwestern.edu

1 Introduction

Climate change is frequently described as an existential threat. This view, however, stands in stark contrast to empirical estimates of the impact of climate change on economic activity, which imply that a 1°C rise in the world's temperature reduces world output at most by 1-3%. Under any conventional discounting, such effects seem hardly catastrophic. Why are perceptions of climate change misaligned with empirical estimates? Do existing estimates account for the full impact of climate change? Are the costs of climate change truly so small?

In this paper, we reconcile both views and demonstrate that the macroeconomic impacts of climate change are six times larger than previously documented. We reach this conclusion in two steps. First, we rely on a time-series local projection approach to estimate the impact of global temperature shocks on Gross Domestic Product (GDP). This approach exploits natural variability in global mean temperature—the source of variation closest to climate change—which we show to predict extreme climatic events much more strongly than country-level temperature. We find that a 1°C rise in global temperature lowers world GDP by 12% at peak. Second, we use our reduced-form results to estimate structural damage functions in a simple neoclassical growth model. We find that climate change leads to a present value welfare loss of 31% and a Social Cost of Carbon (SCC) of $1,056 per ton of carbon dioxide (tCO2).

In the first part of the paper, we develop our time-series approach. Our first contribution is to assemble a new climate-economy dataset spannng the last 120 years from sources that are regularly updated and thus allows us to estimate impacts up to recent years. We construct global and country-level measures from high-resolution gridded land and ocean surface air temperature data from Berkeley Earth and combine them with granular reanalysis measures of extreme temperature, wind speed and precipitation from the Inter-Sectoral Impact Model Intercomparison Project (ISIMIP). We obtain economic data on GDP, population, consumption, investment and productivity from the Penn World Tables spanning 173 countries from 1960 onwards and combine them with data from the Jordà-Schularick-Taylor Macrohistory database for select countries since 1900.

To estimate the causal effects of temperature on GDP, we construct global and local

(country-level) temperature shocks. Identification is complicated by the trending behavior of GDP and temperature. Our approach isolates innovations to the temperature process that are orthogonal to their long-run trends and persist for up to two years using the approach in Hamilton (2018).

Our choice of period is motivated by the geoscience literature. Natural climate variability is driven by multiple phenomena. External causes such as solar cycles and volcanic eruptions lead to medium- and short-run fluctuations in the Earth's mean temperature. Internal climate variability—interactions within the climatic system itself—lead to irregular fluctuations in temperature and weather extremes. For instance, the El Niño-La Niña cycle varies unpredictably between 2 to 7 years.

We map out the dynamic causal effects of our global temperature shocks on world GDP using local projections from 1960 onwards. A 1°C innovation to global mean temperatures leads to a gradual decline in world GDP that peaks at 12% after 6 years and does not fully mean-revert even 10 years after the shock.

There are four possible threats to the causal interpretation of our headline results. We address each of them in a series of robustness exercises. First, global temperature shocks may coincide with the global economic and financial cycle. We account for this possibility by controlling for rich measures of world economic performance: indicators for global economic downturns (such as the major oil shocks in the 1970s or the Great Recession) and global macro-financial variables (past world real GDP, commodity prices and interest rates). Our results remain unaffected.

Second, reverse causality may affect our main output estimate. As output declines following an increase in global mean temperature, energy consumption drops, Greenhouse Gas (GHG) emissions fall, lowering temperatures and ultimately increasing output going forward. Thus, if anything, reverse causality leads us to *qualitatively* underestimate the true impact of a global temperature shock. In addition, *quantitatively*, reverse causality is likely to be negligible: for any plausible climate sensitivity, the temperature impact of short-run fluctuations in emissions is small relative to typical global temperature shocks. We confirm these arguments by explicitly adjusting for the lagged impact of past emissions and find virtually identical results.

Third, our estimated output response may be specific to a particular period of time.

We test whether this is the case by estimating our specification on three separate time frames: 1900-2019, 1960-2019—our main sample—and 1985-2019. We find remarkably similar estimates in all three samples.

Fourth, global temperature shocks may be driven by some countries more than others, and these countries may also have systematically higher or lower GDP growth for unrelated reasons. We account for this possibility by projecting country-level GDP—rather than global GDP—on global temperature. We then control for country fixed effects and region-specific time trends, but, crucially, no time fixed effects to identify the average impact of global temperature. We obtain virtually identical results across all specifications. Taken collectively, our robustness exercises ultimately support the view that our specification captures the causal effect of global temperature shocks on economic activity.

Our estimated effect of temperature shocks on world GDP stands in stark contrast to existing estimates of the cost of climate change. Nordhaus (1992), Dell et al. (2012), Burke et al. (2015) and Nath et al. (2023) find that a 1°C temperature shock reduces GDP by at most 1-3% in the medium run. Why do we find dramatically larger effects?

Our estimate is six times larger than in previous work because we focus on a different source of temperature variation, one that captures the comprehensive impact of climate change: changes in *global* mean temperature. By contrast, previous work exploits variation in *country-level, local* temperatures. It turns out that *global* temperature has much more pronounced impacts on economic activity than *local* temperature. When we estimate the impact of *local* temperature on country-level GDP, based on the same empirical specification and using the same approach to construct temperature shocks, we find similarly small effects to previous studies. Econometrically, previous work that exploits *local* temperature in a panel setting nets out common impacts of *global* temperature shocks through time fixed effects. Instead, we focus on these common impacts.

Why, then, does global temperature depress economic ativity so much more than local temperature? We uncover a novel relationship that rationalizes this difference. Global temperature shocks predict a large and persistent rise in extreme climatic events that cause economic damage: extreme temperature, extreme wind, and extreme precipitation (Deschênes and Greenstone, 2011; Hsiang and Jina, 2014; Bilal and Rossi-Hansberg, 2023). By contrast, local temperature shocks predict a much weaker rise in extreme temperature,

and barely any rise in extreme wind speed and precipitation. This conclusion is consistent with the geoscience literature: extreme wind and precipitation are outcomes of the global climate that depend on ocean temperatures and atmospheric humidity throughout the globe, rather than outcomes of idiosyncratic local temperature realizations.

Consistently with heterogeneous exposure to extreme events, we find suggestive evidence that the impact of global temperature shocks on country-level GDP varies by baseline temperature. Warmer countries are more severely affected than cold countries, while high-income and low-income countries experience similar effects. However, these comparisons are imprecisely estimated and should be interpreted with some caution.

In the second part of the paper, we develop a simple neoclassical growth model to translate our reduced-form estimates into welfare effects. Our model extends the economic block of the Dynamic Integrated Climate Economy (DICE) model of Nordhaus (1992) to include capital depreciation damages in addition to productivity damages from global temperature. Critically however, we use our novel reduced-form effects to obtain new structural damage function estimates.

We recover the underlying productivity and capital depreciation shocks that correspond to a global temperature shock by inverting their mapping to the estimated impulse response function of output and capital. Our framework is simple enough that this mapping has a closed-form expression that guarantees identification. In doing so, we also account for the own internal persistence of global mean temperature. Specifically, we use the approach in Sims (1986) to construct the empirical response to a fully transitory temperature shock as our estimation target. We purposefully remain *conservative* and use a specification with persistent level effects rather than growth effects. We find that a one-time transitory 1°C rise in global mean temperature leads to a 2.5% peak productivity decline and a 0.3 percentage point (p.p.) peak rise in the capital depreciation rate. These effects gradually vanish but, consistently with our persistent impacts on extreme events, persist for nearly 10 years.

With the estimated model at hand, our main counterfactual is a gradual increase in global mean temperature that starts in 2024 and reaches 3°C above pre-industrial levels by 2100—so 2°C above 2024 temperatures—with a 2% discount rate. Climate change implies precipitous declines in output, capital and consumption that exceed 50% by 2100.

These changes imply a 31% welfare loss in permanent consumption equivalent in 2024, that grows to nearly 52% by 2100. These magnitudes are comparable to the economic damage caused by fighting a war domestically and *permanently*. Our results also indicate that world GDP per capita would be 37% higher today had no warming occurred between 1960 and 2019 instead of the 0.75°C observed increase in global mean temperature.

The estimated model lets us characterize the SCC. We follow Folini et al. (2024) and use the temperature response of global mean temperature to a CO2 pulse from Dietz et al. (2021) to map welfare losses into the SCC while remaining consistent with state-of-the-art atmospheric circulation models. We purposefully remain *conservative* and use the lower end of the range of temperature responses from Dietz et al. (2021), which are also consistent with historical emissions and warming data.

We obtain a SCC of \$1,056/tCO2. This value is six times larger than the high end of existing estimates (Rennert et al., 2022). We construct bootstrapped confidence intervals around welfare impacts and the SCC. The 68% confidence interval for the SCC ranges from \$723/tCO2 to \$1,451/tCO2. While the range of plausible values is thus non-trivial, even the lower bound of that confidence interval is multiple times larger than conventional SCC estimates.

Our focus on global temperature shocks accounts for the bulk of this substantial difference. When we re-estimate our model based on the impact of local temperature shocks on productivity only as in previous research, the welfare cost of climate change is 4% and the SCC is \$151/tCO2. By contrast, the impact of gobal temperature on productivity only implies a welfare loss of 24% and a SCC of \$833/tCO2. Including damages to capital depreciation further increases these values to our main results.

How sensitive are these results to changes in the discount rate and the warming scenario? Any plausible discount rate and 2100 temperature leads to welfare losses in excess of 20% and a SCC above \$500/tCO2. Discount rates below 1% lead to the SCC exceeding \$3,000/tCO2. Pessimistic scenarios with 2100 warming reaching 5°C above pre-industrial levels imply welfare losses larger than 60%.

We conclude by delineating the consequences of our results for decarbonization policy. Many decarbonization interventions cost between \$27 and \$95 per ton of CO2 abated (Bistline et al., 2023). A conventional SCC value of \$151/tCO2 implies that these policies

are cost-effective only if governments internalize benefits to the entire world, as captured by the SCC. However, a government that only internalizes domestic benefits values decarbonization benefits using a Domestic Cost of Carbon (DCC). The DCC is always lower than the SCC because damages to a single country are lower than at a global scale. For instance, under conventional estimates based on local shocks, the DCC of the United States is \$30/tCO2, making unilateral emissions reduction prohibitively expensive. Under our new estimates however, the DCC of the United States becomes \$211/tCO2 and thus largely exceeds policy costs. In that case, unilateral decarbonization policy becomes cost-effective for the United States.

Related literature. Our paper relates to the vast literature that measures economic damages from climate change. The canonical approach exploits variation in weather outcomes over time within a given spatial area to estimate the effects of *local* temperature on economic outcomes in a panel structure to achieve credible identification (see Dell et al., 2014 for a comprehensive review; and Dell et al., 2012; Burke et al., 2015; Newell et al., 2021 among many others). Nath et al. (2023) clarify how to translate short-run weather effects to long-run counterfactuals by distinguishing between the polar cases of transitory ("level") effects, or permanent ("growth") effects of temperature on GDP. Consistently across all these studies, medium-term effects range from 1% to 3% of GDP and rely exclusively on climatic variation within country or even smaller geographic units.

Our paper contributes to this literature by taking a fundamentally different approach: we directly exploit aggregate time-series variation in *global* mean temperature to capture climate change as closely as possible, instead of relying on within-country climatic variation that nets out common effects of global temperature. We obtain substantial damages associated with climate change despite using a *conservative* specification with persistent level effects rather than growth effects, similarly to Nath et al. (2023).

Perhaps surprisingly, few studies have explored time-series variation in temperature. Bansal and Ochoa (2011) find that the contemporaneous effect of a 1°C global temperature increase is to reduce growth by 1 percentage point. We show that accounting for the persistence of this response is crucial: the peak effect occurs six years out and is twelve times larger than the contemporaneous impact. Perhaps closest to our paper, Berg et al.

(2023) analyze the effects of global and idiosyncratic temperature shocks on GDP *dispersion* across countries. By contrast, our paper provides the first direct estimate of the *aggregate impact* of global mean temperature shocks. It is the relevant object to construct the SCC and is much more precisely estimated than individual country-level responses. Crucially, we use our macroeconomic estimates in a structural model to evaluate welfare losses and the SCC, and reconcile panel and time-series estimates by documenting the differential impact of global and local temperature shocks on extreme events.

As a result, our paper also contributes to the literature studying the economic impact of storms and heatwaves (Barro, 2006; Deschênes and Greenstone, 2011; Deryugina, 2013; Hsiang and Jina, 2014; Bilal and Rossi-Hansberg, 2023; Phan and Schwartzman, 2023; Tran and Wilson, 2023). We provide new evidence on the relationship between global temperature and extreme climatic events.

Our paper also connects to the literature assessing the welfare implications of climate change using Integrated Assessment Models surveyed in Nordhaus (2013). Recent work develops "bottom-up" models featuring rich regional heterogeneity, migration (Desmet and Rossi-Hansberg, 2015; Desmet et al., 2021; Cruz and Rossi-Hansberg, 2023; Rudik et al., 2022; Conte et al., 2022) and capital investment (Krusell and Smith, 2022; Bilal and Rossi-Hansberg, 2023) to match micro-level estimates and aggregate using the model. Our paper takes the reverse "top-down" approach: directly estimate and match the macroeconomic impact of changes in global temperature. Of course, our analysis remains necessarily silent about distributional effects.

Finally, our paper informs the long-lasting debate about whether Integrated Assessment Models are well-suited to represent the cost of climate change (Nordhaus, 2013; Stern et al., 2022). Our paper demonstrates that these models have historically delivered small costs of climate change not so much because they relied on incomplete foundations, but instead because they were calibrated to economic damages that did not represent the full impact of climate change.

Outline. The rest of this paper is organized as follows. Section 2 describes the data and estimates the macroeconomic effects of temperature shocks using our time series approach. Section 3 investigates how the effects of global and local temperature com-

pare. Section 4 introduces our dynamic model and describes our structural estimation approach. Section 5 evaluates the welfare implications of climate change. Section 6 concludes.

2 Global Temperature and Economic Growth

Climate change originates with a rise in *global* mean temperature. This change in global temperature affects the Earth's climate system as a whole—causing changes in weather patterns, ocean currents and atmospheric conditions, which in turn influence the frequency, intensity, and distribution of extreme weather events globally. Thus, we focus on the fundamental variability in global temperature to analyze the full impact of climate change on the world economy.

2.1 A Novel Climate-Economy Dataset

Our starting point is to construct a dataset covering 173 countries over the last 120 years to study the effects of temperature on the economy. We use world aggregates from this dataset in this section, and country-level outcomes in Section 3 below.

We obtain temperature data from the Berkeley Earth Surface Temperature Database. It provides temperature anomaly data at a spatial resolution of $1^\circ \times 1^\circ$ of latitude and longitude. Based on this gridded data, we construct population- and area-weighted temperature measures at the country level. We complement these local temperature measures with global mean temperature data from the National Oceanic and Atmospheric Administration (NOAA). As expected, aggregating the Berkeley Earth data to obtain a global temperature measure produces a series that is virtually perfectly correlated with the NOAA data series.

We rely on data from ISIMIP for information on extreme weather events. ISIMIP provides global, high-frequency datasets that record multiple atmospheric variables over the 20^{th} and early 21^{st} centuries. We use ISIMIP's observed climate dataset. It contains daily reanalysis measures of temperature, wind speed and precipitation, spanning the period 1901-2019 at the $0.5^\circ \times 0.5^\circ$ resolution. Based on this data, we compute indices of ex-

treme weather by recording the fraction of days within a country that experience extreme weather. We define extreme weather of each category as a realization above a fixed percentile of the daily weather distribution in 1901-1930 (for more details, see Appendix A.1).

We combine our climate dataset with economic information on GDP, population, consumption, investment, and productivity. We obtain a high-quality dataset for a comprehensive selection of countries around the world from the Penn World Tables. We also rely on data from the World Bank as an alternative. Given that both datasets only go back to the 1950s or 1960s, we also include data from the Jordà-Schularick-Taylor Macrohistory database, which features high-quality economic data for a selection of high-income countries starting in the late 19th century.

2.2 Global Temperature Shocks

How does global temperature affect economic growth? Figure 1 displays the evolution of global average temperature and world real GDP per capital since the post-World War II era in our dataset. In the mid-1950s to the mid-1970s, global average temperature remained relatively stable at around 14°C. However, from the late 1970s onward, global average temperature began to steadily rise again. At the same time, we observe relatively stable economic growth over the entire sample.

Figure 1: Global Average Temperature and Output Since 1950

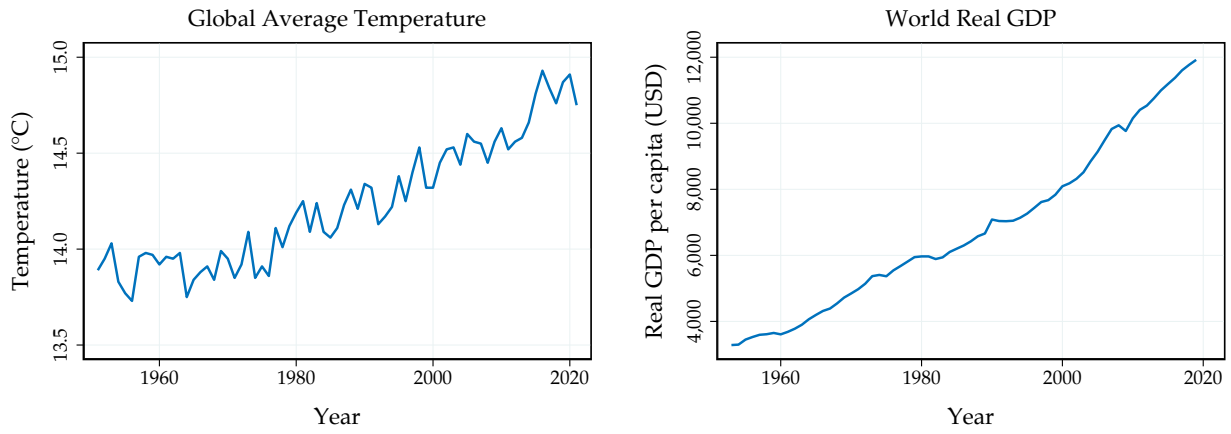

Notes: The figure shows the evolution of global average temperature, computed based on global temperature anomaly data and the corresponding climatology from NOAA, in the left panel, and the evolution of world real GDP per capita (in 2017 USD) computed based on PWT data in the right panel.

The trending behavior of the two series in Figure 1 complicates the identification of the economic effects of temperature increases. A simple regression of global GDP on temperature will yield a spuriously positive association between the two variables, as economic growth is associated with higher GHG emissions which eventually translates into higher temperature. Therefore, we do not focus on the level of temperature as the treatment in our projections, but instead focus on so-called *temperature shocks*. We define such shocks as potentially persistent deviations from the long-run trend in global mean temperature.

What drives these variations in temperature around the trend? The geoscience literature indicates two types of causes. First, external causes such as solar cycles and volcanic eruptions lead to short-run fluctuations in the Earth's mean temperature. Solar cycles have a typical period of 10 years and can warm the Earth by as much as 0.1°C (National Oceanic and Atmospheric Administration, 2009). Volcanic eruptions have shorter-lived cooling effects of up to 2 years due to sulphuric aerosols that increase albedo (National Oceanic and Atmospheric Administration, 2005). Second, internal climate variability—interactions within the climatic system itself that lead to irregularly recurring events—also affects temperatures. For instance, the El Niño-La Niña cycle varies unpredictably between 2 to 7 years and substantially affects global mean temperatures and weather extremes (Kaufmann et al., 2006; National Oceanic and Atmospheric Administration, 2023).

An important question is how to isolate the trend and transient components of temperature. To estimate the effects of temperature on future economic outcomes, it is critical to preserve the causality—in a time-series sense—of the data: we cannot rely on future values of temperature to identify the trend in the current period. In addition, the physical properties of natural climate variability require to allow for somewhat persistent deviations from trend.

One approach that satisfies our needs along both these dimensions is the method proposed by Hamilton (2018). The idea is to regress temperature h periods out on some of its lags as of period t and construct the temperature shock as the innovation in this regression:

$$\widehat{T^{\text{shock}}_{t+h}} = T_{t+h} - (\widehat{\alpha} + \widehat{\beta}_1 T_t + \ldots + \widehat{\beta}_{p+1} T_{t-p}), \tag{1}$$

where $\widehat{\beta}_i$ denotes the coefficient estimates of the regression of temperature on its lag i.

This exercise amounts to isolating shocks that persist typically for h periods. Selecting the horizon h is of course a crucial choice. Motivated by the fact that the climatic events that we consider can last for up to several years, we select a horizon of $h = 2$ and set the number of lags to $p = 2$ in our main specification. However, our results are virtually unchanged when we vary these values. In particular, the results are robust to identifying temperature shocks as one-step ahead forecast errors—an approach that is commonly used in the literature. We compare our approach to other ways of measuring the transient component in temperature in the robustness section below and in Appendix A.4.

Figure 2: Global Temperature Shock

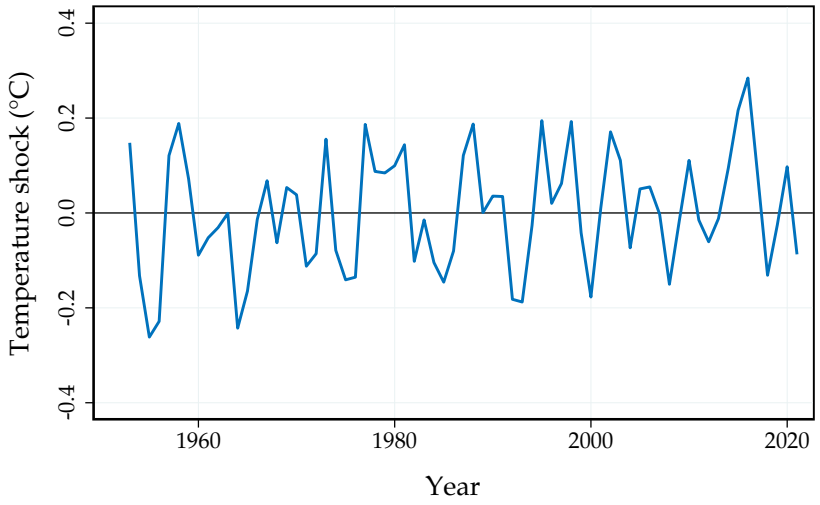

Notes: The figure shows the global temperature shocks, computed as in Hamilton (2018) with ($h = 2$, $p = 2$), over the post-World War II era.

Figure 2 shows the evolution of the resulting global temperature shocks over our sample of interest. As expected, the temperature shocks fluctuate around zero with an almost equal number of positive and negative shocks. The largest temperature shocks in our sample are around 0.3°C . The series is also weakly autocorrelated, reflecting the fact that we allow for relatively persistent deviations from the long-run temperature trend (see Figure A.2 in the Appendix). In our empirical specification, we therefore control for lagged temperature shocks as well; otherwise, serial correlation may bias the estimated impacts when not properly accounted for (see Nath et al., 2023, for an extended discussion of this point).

2.3 The Effect of Temperature Shocks in the Time Series

The economic effects of temperature shocks may take time to materialize. Therefore, we focus on the dynamic effects of temperature shocks up to 10 years out. Thus, we evaluate directly the long-run effects of temperature without the stringent assumptions required to extrapolate short-term temperature impacts. Of course, we would ideally trace out even longer-run effects, but our limited sample period prevents us from doing so consistently.

We estimate the dynamic causal effects to global temperature shocks using local projections à la Jordà (2005). This approach involves estimating the following series of regressions, one for each horizon $h = 0, \dots, 10$:

$$y_{t+h} - y_{t-1} = \alpha + \theta_h T_t^{\text{shock}} + \mathbf{x}_t' \boldsymbol{\beta} + \varepsilon_{t+h}, \tag{2}$$

where y_t is (log) world real GDP per capita, T_t^{shock} is the temperature shock and θ_h is the dynamic causal effect of interest at horizon h. We refer to the dynamic causal effects up to horizon h as the Impulse Response Function (IRF). $\mathbf{x}_t$ is a vector of controls and ε_t is a potentially serially correlated error term. To account for the serial correlation in GDP growth and temperature shocks, we include 2 lags of real GDP growth per capita and the global temperature shock. The confidence bands are computed based on the lag-augmentation approach (Montiel Olea and Plagborg-Moller, 2021).[1]

Omitted variable bias. A first threat to identification is that global temperature innovations may happen to be correlated with the global economic and financial cycle over time. For instance, if a severe El Niño event increases global average temperature at the same time that a global recession occurs, we will mistakenly attribute adverse economic impacts to climatic variations.

To account for this possibility, we include rich controls of the world economic performance. In particular, we control for global economic downturns, such as the large oil

[1] As in Nath et al. (2023), we currently do not take estimation uncertainty in the global temperature shock into account in our baseline specification. Alternatively, we use an local projection-instrumental variable approach, where we instrument changes in global temperature by the global temperature shock. Using the global temperature shock as an instrumental variable yields inference that is robust to estimation uncertainty in the shock. Reassuringly, the reduced-form and the instrumental variable approach yield very similar results both in terms of point estimates and inference. See Appendix A.2.3 for more details.

Figure 3: The Effect of Global Temperature Shocks on World Output

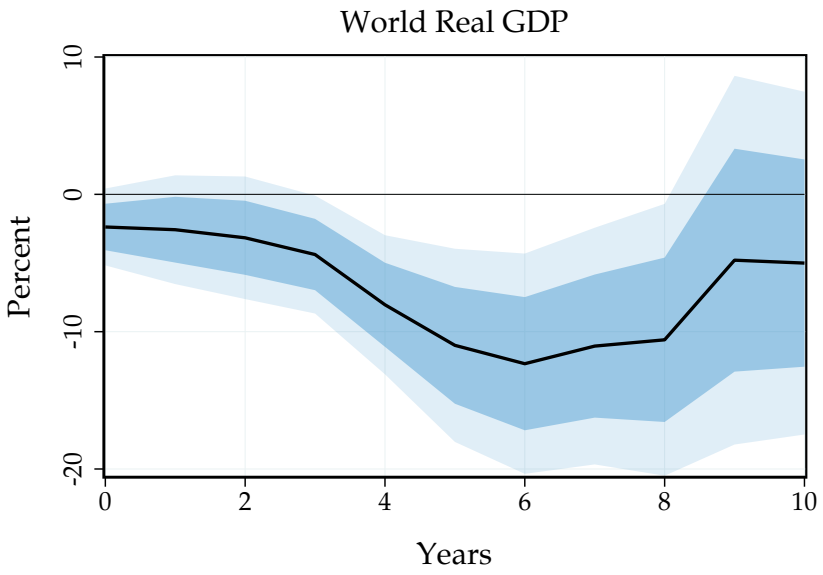

Notes: The figure shows the impulse responses of world real GDP per capita to a global temperature shock, estimated based on (2). The solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

shocks in the 1970s or the Great Recession, using a set of dummy variables.[2] Alternatively, we include a wider set of global macroeconomic and financial variables as additional controls.

Figure 3 shows the impulse response of world real GDP per capita to a global temperature shock of 1°C. The solid black lines are the point estimates and the shaded areas are 68 and 90% confidence bands, respectively. On impact, world real GDP falls by about 2%. However, the effect builds up over time. After about 6 years, world real GDP falls by more than 10%, and the adverse impact persists even 10 years after the shock. Our estimate represents major economic effects: it is of the same magnitude as the growth impacts that typically occur after severe banking or financial crises (Cerra and Saxena, 2008; Reinhart and Rogoff, 2009).

On the one hand, a 1°C temperature shock is a large shock that does not occur directly in our historical sample: we observe much smaller shocks throughout our sample. Our

[2]Our definition of global recession dates follows the World Bank (Kose et al., 2020). Specifically, we focus on the following episodes: 1973-1975, 1979-1983, 1990-1992, and 2007-2009. To allow for potential persistent effects of recessions, we also include 2 lags of the global recession indicator variable.

estimate for a 1°C shock scales up the linear effect of these smaller shocks. In effect, we abstract from potential non-linearities. However, in the presence of potential tipping points, we would expect even larger effects than predicted by our linear model.

On the other hand, a 1°C change is smaller than global warming going forward. The Intergovernmental Panel on Climate Change (IPCC) expects that under a business-as-usual scenario, global average temperature is expected to increase by 4.4°C by 2100 (Lee et al., 2023). We now demonstrate that our main estimate is robust to accounting for further threats to identification.

Reverse causality. Changes in economic activity may affect short-run variations in temperature: a decline in economic activity lowers emissions and temperature, and hence increases output going forward. This mechanism leads to a reverse causality threat.

There are three reasons why this concern is unlikely to substantially affect our interpretation. First, any reverse causality concern leads us to underestimate the effect of temperature on economic output. As temperature rises and economic activity initially declines, the resulting fall in emissions implies lower future temperatures and thus higher future output. Thus, true damages would be even larger than our estimates.

Second, and perhaps most importantly, annual fluctuations in emissions imply negligible temperature variations relative to the typical temperature shocks that we exploit. Typical year-to-year fluctuations in CO2 *emissions* are of the order of 2 gigatons. After accounting for oceanic and biosphere absorption, these annual fluctuations translate into 1 gigaton of *atmospheric* CO2. This magnitude corresponds to 0.15 part per million (ppm) in atmospheric CO2 *concentration*. Current CO2 atmospheric concentration is just above 400 ppm. Given a climate sensitivity between 2 and 4, year-to-year fluctuations in emissions thus imply year-to-year fluctuations in temperature of about 0.0005°C. This is an order of magnitude lower than natural climate variability which is of the order of 0.1°C.

Nevertheless, we perform two exercises to verify that reverse causality is unlikely to affect our results. First, we check whether our temperature shocks are forecastable by past macroeconomic or financial variables. To this end, we perform a series of Granger-causality tests. We find no evidence that macroeconomic or financial variables have any power in forecasting global temperature shocks (see Appendix A.2.1).

Second, we explicitly account for the feedback between output and temperature. We use estimates for the emissions-to-GDP elasticity and the sensitivity of temperature to an emissions impulse from Dietz et al. (2021). We then directly purge the GDP impulse response from the dynamic effects of past emissions changes. Figure 4(a) shows the results. Explicitly controlling for reverse causality has no meaningful effect on our results (see Appendix A.3 for more details).

Specification choices. Our main result—a significant fall in world output after global temperature shocks—is also robust to variations in our time-series specification. Figure 4 displays two additional sensitivity checks. First, our results are robust to changes in the definition of global temperature shocks. Panel (b) indicates that constructing temperature shocks as one-step ahead forecast errors following previous work (see e.g. Bansal and Ochoa, 2011; Nath et al., 2023), or using a one-sided HP filter, produces similar results.

Second, the effect remains when we expand the set of controls. In our baseline specification, we already control flexibly for recession periods and also include controls for past world GDP growth. Figure 4(c) reveals that we obtain similar results if we expand the set of controls to include global oil prices and the U.S. treasury yield. If anything, dropping the controls for recession periods tends to mildly attenuate the effect.

Overall, these results corroborate our interpretation that global temperature shocks are driven by external causes and internal climate variability and have a large causal effect on world GDP. We expand more flexibly on these robustness checks in the next section, where we study the effects of global temperature shocks in a panel of countries.

3 Temperature Shocks in the Panel of Countries

So far we have evaluated the impact of global temperature shocks directly on world GDP. We now exploit country-level data on GDP to achieve four distinct goals. Our first goal in Section 3.1 is to exploit the additional statistical power in the panel to further corroborate our results when controlling for possibly confounding trends at the country level and varying the span of our sample period. Our second goal in Sections 3.2 and 3.3 is to contrast the impact of global temperature shocks with existing work that has focused

Figure 4: Sensitivity of the Effect of Global Temperature Shocks in the Time Series

(a) Reverse causality

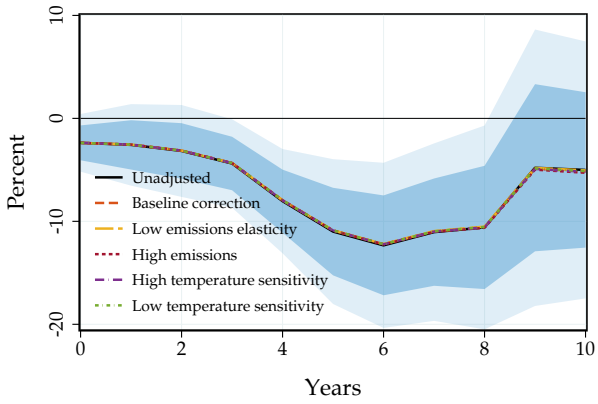

(b) Construction of temperature shock

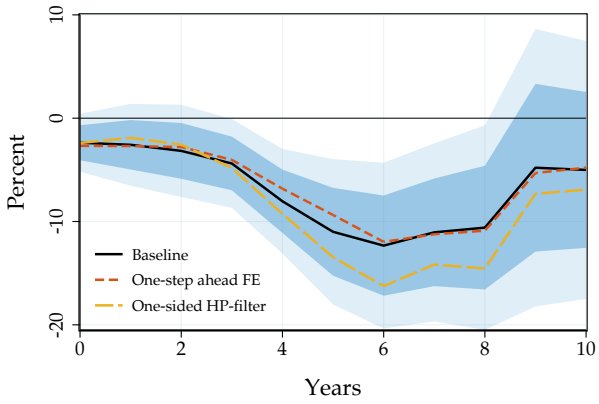

(c) Sensitivity with respect to controls

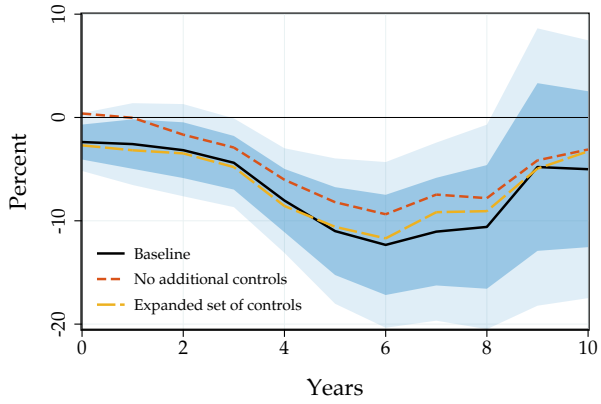

Notes: The figure shows the impulse responses of world real GDP per capita to a global temperature shock, estimated based on (2). Panel (a) shows how the GDP response is affected when we formally account for reverse causality, i.e. lower GDP leading to lower temperature, based on varying assumptions on the emissions elasticity, the level of emissions and the temperature sensitivity. Panel (b) illustrates the sensitivity with respect to the construction of the temperature shock. We compare our baseline, using the Hamilton (2018) approach, to the more commonly used one-step ahead forecast error and a shock obtained using the one-sided HP filter. Panel (c) shows the sensitivty with respect to the controls included. We compare our baseline to a specification that also controls for oil prices and the one-year US treasury yield, as well as to the case where we only control for lags of the temperature shock and GDP growth. In all subfigures, the lines correspond to point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively, corresponding to our baseline specification.

on country-level temperature shocks. Our third and fourth goals are to explore the margins through which GDP declines (Section 3.4) and the heterogeneity in country-level responses (Section 3.5).

3.1 Global Temperature Shocks in the Panel

To estimate the dynamic causal effects of temperature shocks in the panel, we employ the panel local projections approach (Jordà et al., 2020). In this section, we still estimate the effect of global temperature shocks, now averaged across 173 countries. However, the panel approach allows us to account for unobserved, time-invariant country characteristics using country fixed effects. We can also control for past GDP growth at the country level and regional trends. Specifically, we estimate the following series of panel regressions for horizons $h = 0, \dots, 10$:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \theta_h T_t^{\text{shock}} + \mathbf{x}_t' \boldsymbol{\beta} + \mathbf{x}_{i,t}' \boldsymbol{\gamma} + \varepsilon_{i,t+h}, \tag{3}$$

where $y_{i,t}$ is (log) real GDP per capita of country i in year t, T_t^{shock} is the (global) temperature shock and θ_h is the dynamic causal effect of interest at horizon h. $\mathbf{x}_t$ is a vector of global controls, $\mathbf{x}_{i,t}$ is a vector of country-specific controls and $\varepsilon_{i,t}$ is an error term. In our baseline specification, we use the same set of global controls as before but in addition also control for two lags of country-level GDP growth. We expand on these controls in further sensitivity checks below.

Because the temperature shock T_t^{shock} does not vary by country, the error term is potentially serially and cross-sectionally correlated. For inference, we thus rely on Driscoll and Kraay (1998) Standard Errors (SEs), which are robust to general forms of cross-sectional and serial dependence.[3]

By design, our specification is close to the specifications commonly used in the panel literature on the economic effects of local temperature shocks (e.g. Dell et al., 2012; Burke et al., 2015; Nath et al., 2023). Crucially however, the temperature shock T_t^{shock} does not vary by country in our case. As a result, we cannot control for time fixed effects as is

[3]Our results are robust to using two-way clustered SEs by country and year, or using bootstrapping techniques for inference. In fact, to construct the confidence bands for our estimated structural damage functions in Section 4.5, we rely on the boostrapped distribution, as estimated using a Wild bootstrap.

commonly done. Instead, we include a selection of global control variables as in our time-series specification (1).

Figure 5 shows the impulse responses to a global temperature shock, estimated in the panel of countries. Consistently with our aggregate time-series evidence, global temperature shocks lead to a significant fall in real GDP per capita, which is slightly larger than 10% at peak and persists even 10 years out. This estimated effect is strikingly similar to the estimates from the time series, indicating that our results are robust to accounting for unobserved fixed country characteristics.

Figure 5: The Average Effect of Global Temperature Shocks

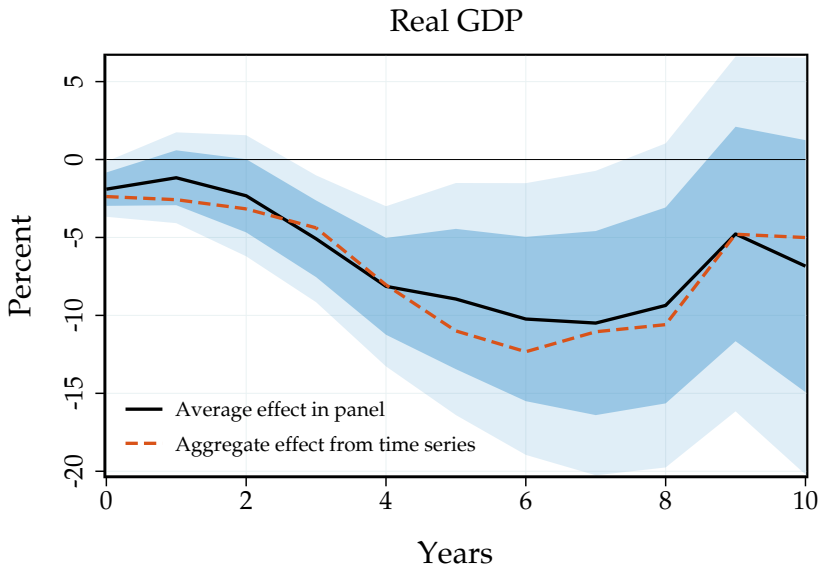

Notes: The figure shows the impulse responses of real GDP per capita to a global temperature shock, estimated in the panel using (3). The solid black line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively. The dashed red line is the aggregate effect of global temperature shocks on world GDP, estimated from the time series.

The increased statistical power in the panel lets us conduct a number of additional sensitivity checks. Our first sensitivity check evaluates whether our results depend on the choice of the sample period. Figure 6 displays our results. In Panel (a), the sample starts after the large oil shocks of the 1970s. We obtain remarkably similar results in this substantially shortened sample period. In Panel (b), we study the effects of global temperature shocks in a much longer sample, starting in 1900. For this analysis, we rely on a smaller set of countries for which we have consistent data. We use real GDP from 18

Figure 6: Sensitivity of the Average Effect of Global Temperature Shocks

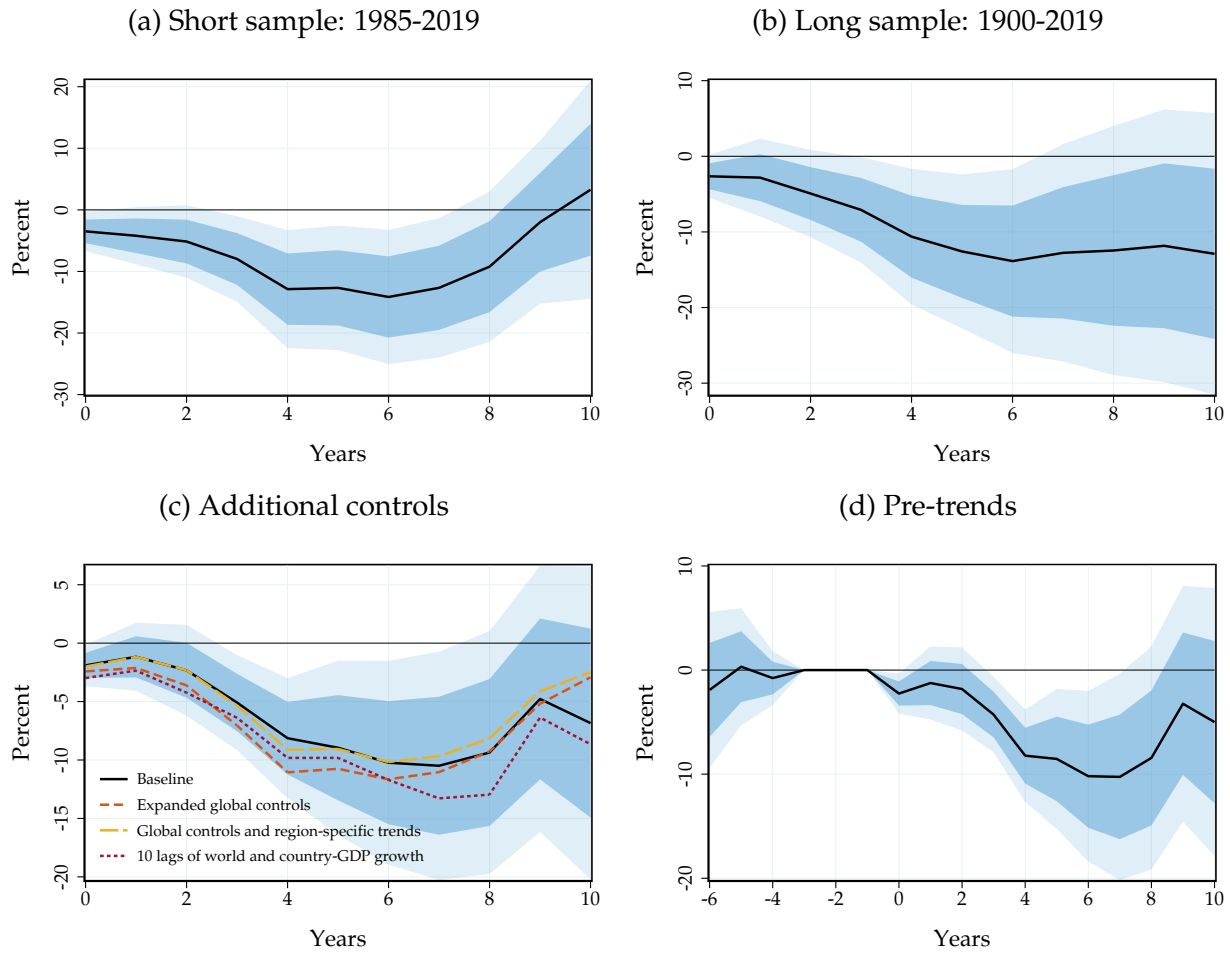

Notes: The figure shows the impulse responses of real GDP per capita to a global temperature shock, estimated in the panel using (3). The first two panels show the sensitivity to the sample period. Panel (a) displays our estimates when the sample is restricted to the most recent period, 1985-2019. Panel (b) presents the results when we use a longer sample starting in 1900. These results are based on a smaller set of developed countries. Panel (c) shows the sensitivity with respect to the controls included. We compare our baseline to a specification that controls for an expanded set of global variables, to a specification that adds subregion-specific time trends, and to a specification that controls for 10 lags of world and country-GDP growth. Panel (d) shows the pre-trends for our baseline response. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

advanced economies in the Jordà-Schularick-Taylor Macrohistory Database. In the longer sample, global temperature shocks are also associated with a significant fall in world real GDP: it reaches close to -15% at peak and turns out to be even more persistent than in our baseline sample.

Our second sensitivity check includes more flexible controls for potential confound-

ing effects. The main concern is that adverse global or regional shocks may coincide with temperature shocks, confounding our estimates. To this end, we add global oil prices, U.S. treasury yield, and, in our most restrictive specification, region-specific time trends. Figure 6(c) shows that our estimates turn out to be virtually invariant to the set of controls. To mitigate reverse causality concerns at longer horizons, we also consider a specification where we include up to 10 lags of world and country-GDP growth as controls. In this way, we flexibly control for potential delayed economic growth and the associated emissions on temperature levels. In Appendix A.2.4, we further establish that unobserved common shocks are not driving our results by exploiting an intermediate level of spatial aggregation of temperature shocks. The results from this specification that allows us to include time fixed effects turn out again to be very close to our baseline case.

Our last sensitivity check investigates whether our results may be due to pre-trends—although Table A.1 already suggests that Granger causality is unlikely to be a concern. Nevertheless, Figure 6(d) plots our main estimate together with estimates 6 years prior to the global temperature shock. Note that the effect in the three years immediately before the shock is zero by construction as we control for two lags of GDP growth. We do not detect any statistically significant nor economically meaningful effect up to 6 years prior to the shock. Overall, these results confirm the substantial and persistent negative effect of global temperature shocks on real GDP.

3.2 Global vs. Local Temperature

How do these effects compare to local temperature shocks? Conventional estimates in the literature imply that a 1°C rise in local temperature reduces GDP at most by 1-3% in the medium run (Nordhaus, 1992; Dell et al., 2012; Burke et al., 2015; Nath et al., 2023). To ensure that our findings are not driven by differences in the econometric specification or data choices, we reproduce the effects of local temperature shocks within our empirical framework.

To this end, we measure local temperature shocks using the Hamilton (2018) filter, as we do in Section 2.2 for global temperature, but now based on population-weighted country-level temperature data.

Figure 7 shows the local temperature shocks for the United States and South Africa

Figure 7: Local and Global Temperature Shocks

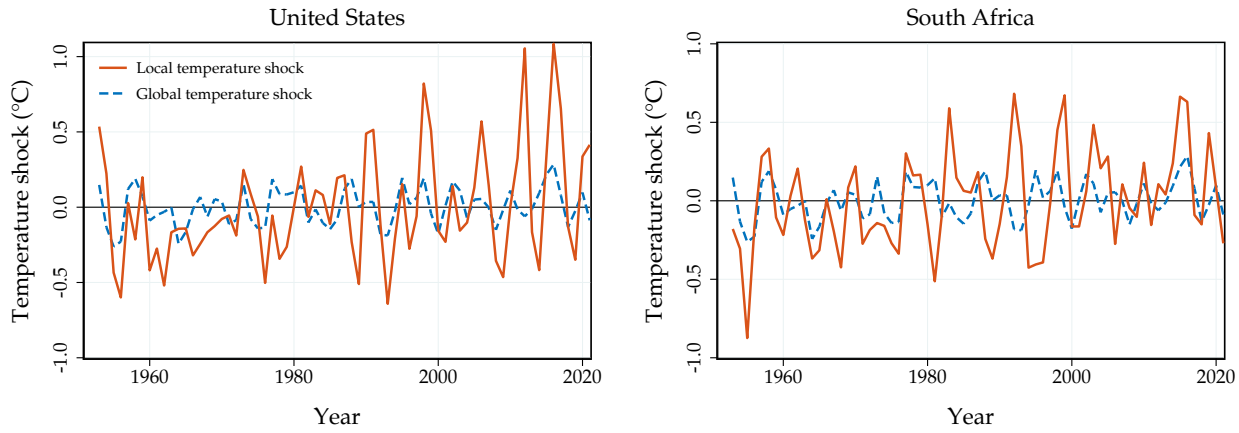

Notes: The figure shows the local temperature shocks for the United States (left panel) and South Africa (right panel) in red together with the global temperature shocks as the blue dashed line. All the shocks are computed based on the Hamilton (2018) approach with ($h = 2$, $p = 2$), over our sample from 1960. The local shocks are computed based on population-weighted country-level temperature data.

over our sample from 1960, as two illustrative examples. Local temperature shocks are larger and more volatile than global temperature shocks. The standard deviation of local shocks is about three to four times larger. Second, while local and global shocks are correlated—the correlation is 0.33—they frequently move in different directions. Thus, local shocks do not always translate into global shocks and vice-versa.

To estimate the responses to local shocks, we rely on our panel specification (3), with the critical difference that the temperature shock is a *country-specific* temperature shock $T_{i,t}^{\text{shock}}$. In this first specification, we do not include time fixed effects to maximize comparability with (3) but include global controls. Alternatively, we also use a specification that includes time fixed effects:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \delta_t + \theta_h T_{i,t}^{\text{shock}} + \mathbf{x}_{i,t}'\gamma + \varepsilon_{i,t+h}, \tag{4}$$

which allows us to flexibly control for unobserved common shocks. In this case, the global controls are absorbed by the time fixed effects.

Figure 8 shows the estimated impulse responses to a local temperature shock of 1°C as the solid red line (global controls and no time fixed effect) and the dashed brown line (with time fixed effects). For comparison, we also include the impulse responses to a

global temperature shock (in black). With or without time fixed effects, local temperature shocks lead to a similar and significant fall in real GDP per capita. On impact, the effect stands at about -0.5% and reaches around -1.5% after 5 years. These estimates are close to previous findings in Dell et al. (2012), Burke et al. (2015), and Nath et al. (2023).

Figure 8: The Average Effect of Local Temperature Shocks

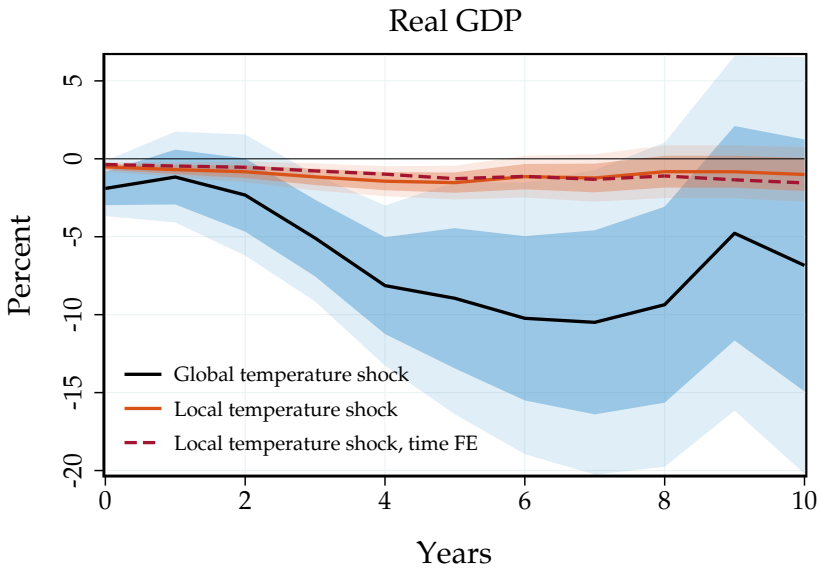

Notes: The figure shows the impulse responses of real GDP per capita to a local temperature shock, estimated in the panel using (3), (in red) against the effects of a global temperature shock (in blue). The solid lines are the point estimates and the dark and light shaded areas and dashed and dotted lines are 68 and 90% confidence bands, respectively. As an additional comparison, we also include the response to a local temperature shock from a specification with time fixed effects (brown dashed line).

This comparison reveals that *global* temperature have much more pronounced impacts on economic activity than *local* temperature. The estimated effects of global temperature shocks are about seven times larger than for local temperature shocks, based on the same empirical model and the same sample period.

Our analysis indicates that the key difference lies in the *nature* of the shock itself rather than in the set of global controls or time fixed effects. Climatic variation within country or even smaller geographic units may alleviate identification concerns, but misses any global effects of climate change—itself a global phenomenon. By contrast, our approach purposefully studies these global effects by focusing on climatic variation at the global level. We conclude that global temperature shocks lead to much larger economic effects

than local temperature shocks.

3.3 Reconciling the Impacts of Global and Local Temperature

Why, then, does global temperature cause more economic harm than local temperature? To shed light on this question, we study the climatic implications of local and global temperature shocks. Specifically, we investigate how these shocks impact the likelihood of extreme weather events, such as extreme temperature, extreme wind speed, and extreme precipitation.

Figure 9 displays our results. We first study the internal persistence of local temperature in response to either a global or a local temperature shock. Both global and local shocks lead to a persistent increase in temperature and the degree of persistence is broadly consistent across the two responses. Thus, the persistence of the shock cannot account for the differential impacts of global and local temperature shocks on GDP. In fact, Figure A.4 in Appendix A.2.2 shows that imposing the same internal persistence in response to global and local temperature shocks using the Sims (1986) method produces very similar results.

We then ask how the temperature shocks affect the occurrence of extreme weather events, such as extreme heat, extreme precipitation, and extreme wind. Local temperature shocks lead to an increase in the share of extreme heat days. However, global temperature shocks lead to a substantially larger increase in extreme heat days.

The contrast is even starker for extreme precipitation and extreme wind speed: global temperature shocks predict a large increase in their frequency, while local temperature shocks do not. These findings are consistent with the geoscience literature: wind speed and precipitation are outcomes of the global climate—through oceanic warming and atmospheric humidity—rather than outcomes of local temperature distributions. Given that extreme climatic events are known to cause economic damage (Deschênes and Greenstone, 2011; Hsiang and Jina, 2014; Bilal and Rossi-Hansberg, 2023), the differential effect of global versus local temperature shocks on extreme climatic events can explain the much larger economic effects of global temperature shocks.

Figure 9: The Impact on Extreme Weather Events

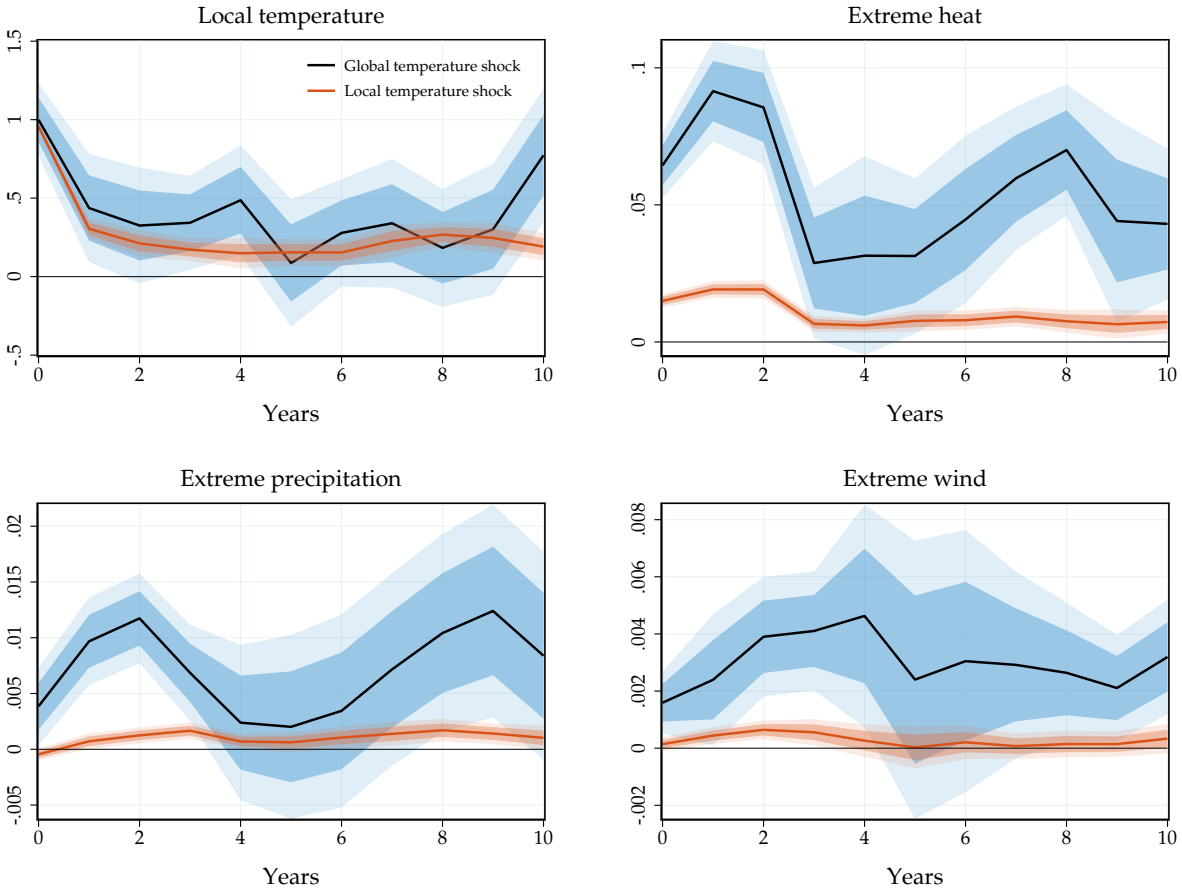

Notes: The figure shows the impulse responses of local temperature, extreme heat days, extreme precipitation days, and extreme wind days to global and local temperature shocks. The extreme weather variables record the share of days in a given year where temperature, precipitation, or wind speed are above a certain threshold. Specifically, we use percentiles of the daily weather distribution in 1900-1910. For temperature and precipitation, we use the 95 percentile, for wind we use the 99 percentile. To eliminate some of the noise inherent in the extreme weather data, we smooth these measures with a backward-looking (current and previous two years) moving average. However, our results are robust to using the raw extreme weather data. The global shock is depicted in blue, the local shock is shown in red. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

3.4 Mechanisms

Through which mechanisms do global temperature shocks transmit to the world economy? We have documented that extreme events rise after such a shock, but which margins of the economy respond most? We answer these questions by evaluating the dynamic causal effects of global temperature shocks on economic variables such as capital, investment and productivity in our panel of countries.

Figure 10: Transmission of Global Temperature Shocks

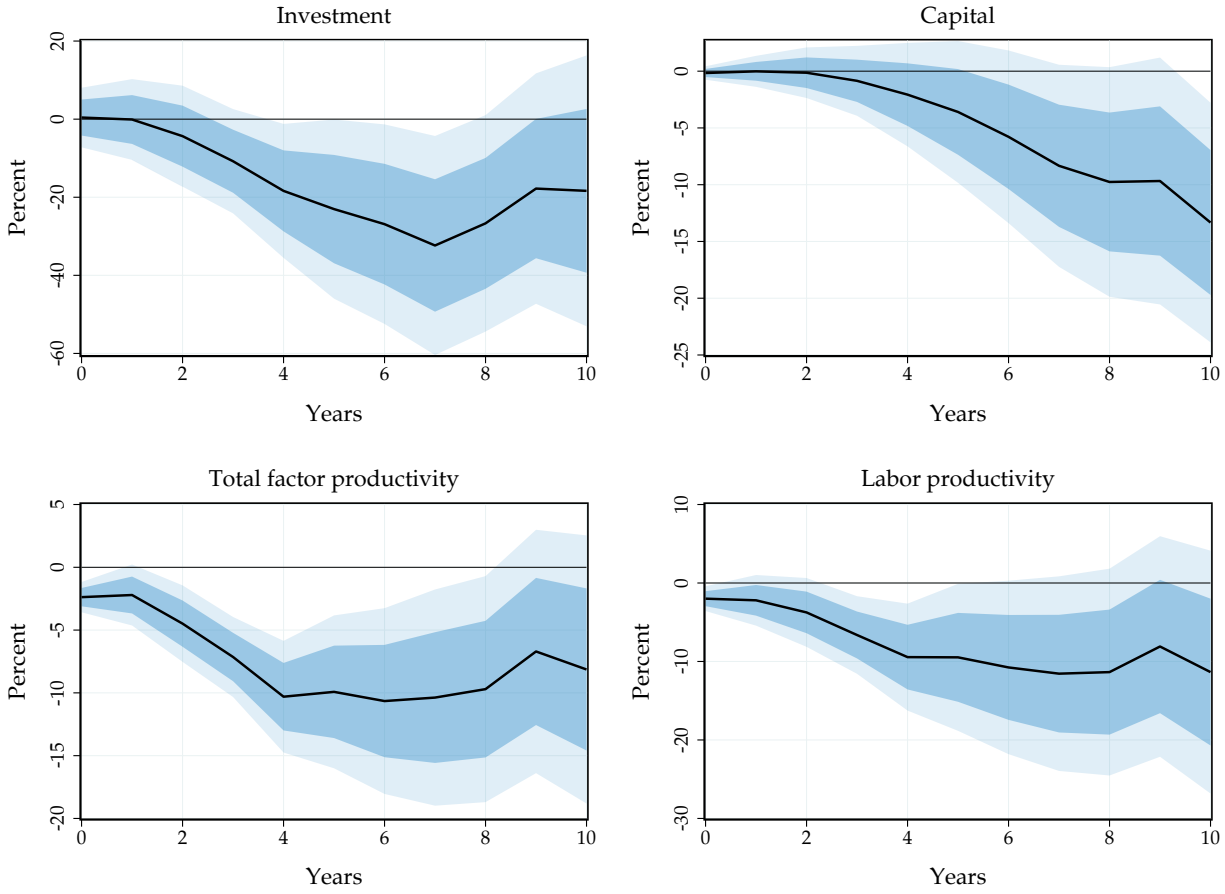

Notes: The figure shows the impulse responses of investment, the capital stock, total factor productivity and labor productivity to a global temperature shock, estimated in the panel based on (3). Investment and the capital stock are expressed in per capita terms. Labor productivity is measured as output over employment. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

Figure 10 displays our results. Global temperature shocks lead to a substantial and significant fall in the capital stock and in investment. The sluggish fall in the capital stock is consistent with the adverse impact of future extreme weather events such as storms that materialize as a sequence of capital depreciation shocks. Consistent with Hsiang and Jina (2014), we find that disasters associated with global warming do not stimulate growth. Instead, national income declines, productive capital dwindles and investment falls.

We also find evidence that productivity falls significantly after global temperature shocks. This is true for Total Factor Productivity (TFP) as estimated in the Penn World Tables and for labor productivity. The impact effect, which stands at about -2%, is consistent with experimental studies on the impact of temperature on productivity (Seppanen et al., 2003). However, these effects tend to build up over time, reaching around -10% after about four years.

3.5 Regional Heterogeneity

We have documented that global temperature shocks lead to a substantial fall in economic activity, on average. How are these effects distributed across countries? Are poorer countries more affected? Are the effects attenuated in countries located in colder climates? And how do the effects vary across different regions?

We start by studying how the effects vary by average temperature and income. To this end, we bin countries into different groups based on temperature and income data. Specifically, we bin countries into three temperature and income groups, based on data from 1957-1959 to ensure that group characteristics are not influenced by the effects of the global temperature shocks.

Figure 11 displays our results. Panel (a) shows the effects to a global temperature shock for cold countries (average temperature below 10°C), temperate climate countries (average temperature between 10°C and 20°C) and hot countries (average temperature above 20°C). Hot countries display the strongest adverse effects of temperature shocks. This result is *qualitatively* consistent with previous evidence on local temperature shocks (Dell et al., 2012; Burke et al., 2015; Nath, 2022). *Quantitatively*, global temperature shocks have larger effects across all countries, especially for hot countries. Consistently, we find that temperate countries also display a response that is economically large. Only for

Figure 11: Heterogeneous Effects of Global Temperature Shocks

(a) By average temperature

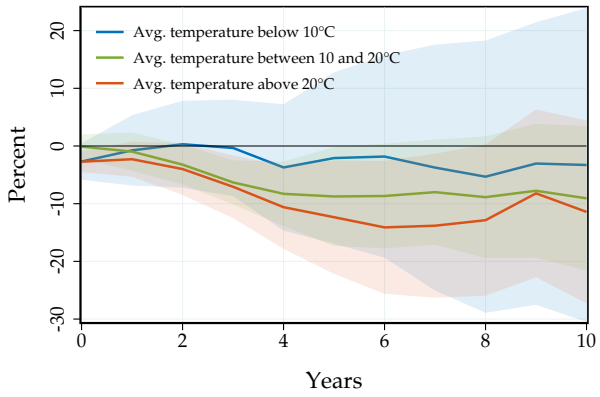

(b) By income per capita

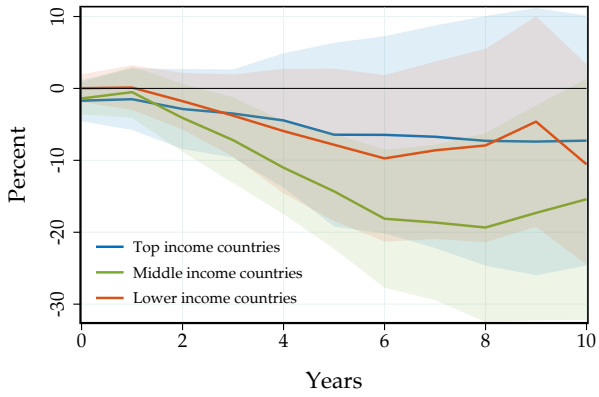

Notes: The figure shows the impulse responses of real GDP per capita to a global temperature shock, for different groups of countries. In Panel (a), we group countries by their average temperature in 1957-1959. In Panel (b), we group countries by their per capita income (in PPP terms) in 1957-1959. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 90% confidence bands, respectively.

colder countries, we observe a somewhat smaller effect that is also not statistically significant.

Figure 11(b) shows the responses by income per capita. Specifically, we consider effects on poorer countries (real GDP per capita below 3,000 USD), middle income countries (real GDP per capita between 3,000 and 12,000 USD), and high income countries (real GDP per capita above 12,000 USD). We find that real GDP per capita falls for all income groups. Interestingly, however, poorer and rich countries display a comparable response. We estimate the largest adverse effects for middle income countries, with a peak effect of close to -20%. We caution that the *relative* effects of global temperature by country temperature and income groups are not precisely estimated and should be interpreted with some caution.

In Figure 12, we study the impact of global temperature shocks on different regions. We document significantly negative effects in most regions. We estimate the strongest negative effects in relatively hot regions such as Southeast Asia and Sub-Saharan Africa. Contrary to local temperature shocks, we document that global temperature shocks lead to adverse economic effects even in higher-income, colder regions. The peak effect in North America is around -10%, and in Europe around -7%, even though the response is

not as precisely estimated. The only region that gains from global temperature shocks is Central and East Asia. We conjecture that this result is driven by the relatively large number of cold countries in this region that may benefit from warmer temperatures.

Figure 12: Regional Impacts of Global Temperature Shocks

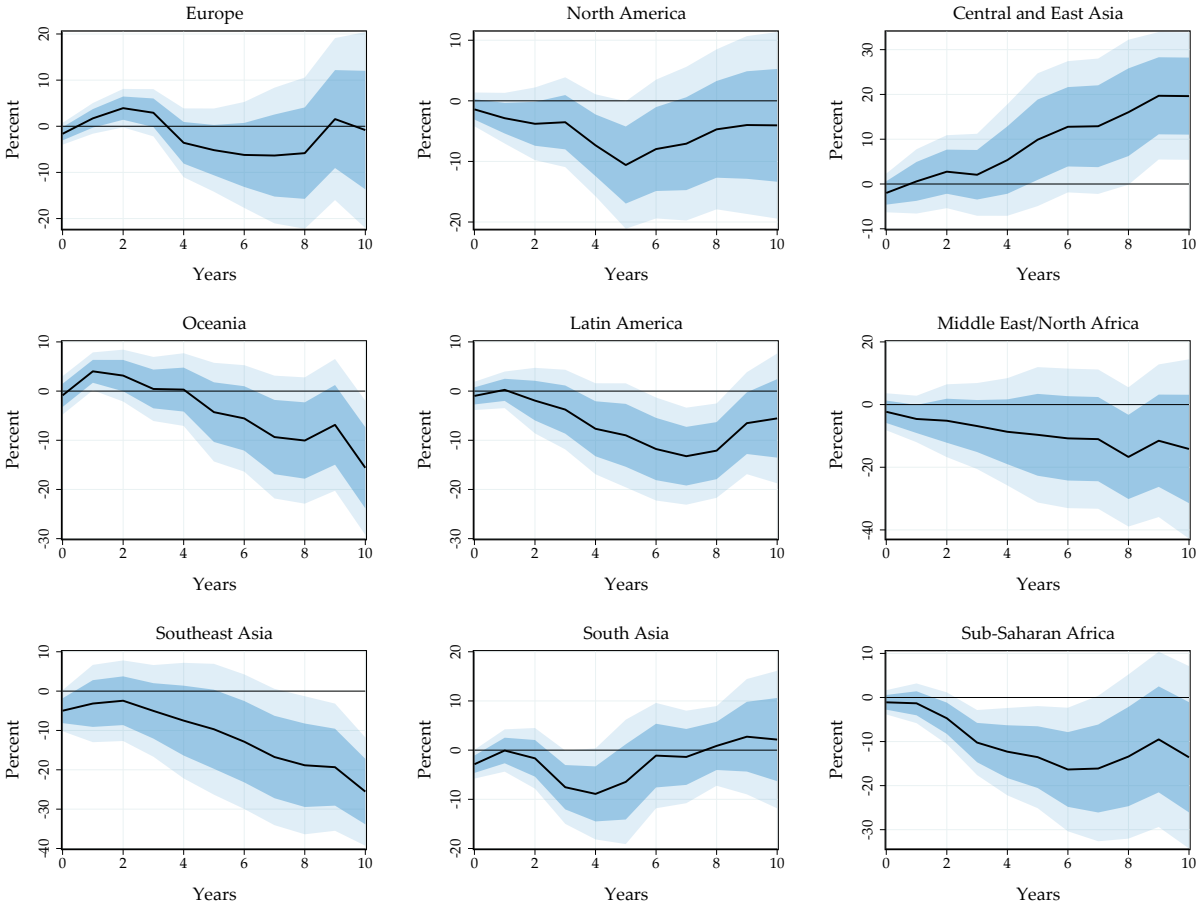

Notes: The figure shows the impulse responses of real GDP per capita to a global temperature shock, for different regions across the world. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

Our results indicate that there are meaningful differences in the effects of global temperature shocks. Overall, these effects are more uniformly detrimental than for local temperature shocks (see e.g. Burke et al., 2015).

So far we established the reduced-form impact of global temperature shocks on economic activity at the world and country level. We now turn to our structural model to convert these estimates into welfare losses and a value of the Social Cost of Carbon.

4 A Model of Climate Change Across the World

Our framework closely follows the standard neoclassical growth model. As such, it mirrors the backbone of the Dynamic Integrated Climate Economy (DICE) model introduced by Nordhaus (1992). Our key innovations are to introduce capital depreciation damages and to use our new reduced-form estimates of the impact of global temperature shocks to structurally estimate the damage functions in the model.

4.1 Setup

Agents and preferences. Time is continuous and runs forever. There is a unit continuum of infinitely-lived identical households who populate the world economy. Households have Constant Relative Risk Aversion (CRRA) flow preferences: $U(C) = \frac{C^{1-\gamma}-1}{1-\gamma}$. Labor supply is exogenous and set to $L_t = 1$. Households discount the future at rate ρ.

Technology. Firms produce according to a Cobb-Douglas production function in capital K_t and labor L_t with time-dependent TFP Z_t: $Y_t = Z_t K_t^{\alpha} L_t^{1-\alpha}$. They hire labor and rent capital from households in competitive factor markets. Production implies a time-dependent capital depreciation rate Δ_t. Firms cover depreciation. The paths of Z_t, Δ_t are perfectly foreseen.

Budgets. Households earn wages w_t, hold capital K_t and rent it out to firms for production. The net interest rate is r_t. Firms make zero profits given constant returns to scale, so we omit profits in the budget constraint of the household, which writes:

$$C_t + \dot{K}_t = w_t + r_t K_t.$$

Households are endowed with an initial capital stock K_0.

Equilibrium. A competitive equilibrium of our economy is a collection of sequences $\{C_t, K_t, L_t, r_t, w_t\}_{t=0}^{\infty}$ such that households optimize given prices $\{r_t, w_t\}_t$:

$$\max_{\{C_t, K_t\}_t} \int_0^{\infty} e^{-\rho t} U(C_t) dt \quad \text{subject to} \quad C_t + \dot{K}_t = w_t + r_t K_t \quad \text{given } K_0;$$

firms optimize given prices $\{r_t, w_t\}_t$: $\max_{K_t^D, L_t^D} Z_t (K_t^D)^\alpha (L_t^D)^{1-\alpha} - (r_t + \Delta_t) K_t^D - w_t L_t^D$; and factor markets clear: $K_t = K_t^D$ and $1 = L_t^D$.

4.2 Climate Change

We model climate change as changes in TFP Z_t and capital depreciation Δ_t over time. We take the path of global mean temperature T_t relative to a reference level T_0 as given. Global mean temperature affects TFP and capital depreciation through structural damage functions $\{\zeta_s,\ \delta_s\}_{s \geq 0}$:

$$Z_t = Z_0 \exp\left(\int_0^t \zeta_s \widehat{T}_{t-s} ds\right) \qquad \Delta_t = \Delta_0 \exp\left(\int_0^t \delta_s \widehat{T}_{t-s} ds\right), \tag{5}$$

where we denoted $\widehat{T}_t = T_t - T_0$ excess temperature relative to the reference level. Only shocks since $t = 0$ affect TFP and capital depreciation.

ζ_s and δ_s govern the persistence of the effect of transitory global temperature shocks on TFP and capital depreciation. When $\zeta_s,\ \delta_s$ are Dirac mass points at $s = 0$, global temperature shocks have purely transitory level effects. When $\zeta_s,\ \delta_s$ are positive functions that asymptote to zero, global temperature shocks have persistent level effects. When $\zeta_s,\ \delta_s$ are positive functions that asymptote to a positive value, global temperature shocks have growth effects.

When temperature $T_t \equiv \overline{T}$ is constant, the economy converges to its steady-state with the corresponding values of TFP and capital depreciation rate:

$$\overline{Z} = Z_0 \exp\left((\overline{T} - T_0) \int_0^\infty \zeta_s ds\right) \qquad \overline{\Delta} = \Delta_0 \exp\left((\overline{T} - T_0) \int_0^\infty \delta_s ds\right). \tag{6}$$

The steady-state expression (6) highlights that the cumulative damage functions $\int_0^\infty \zeta_s ds$ and $\int_0^\infty \delta_s ds$ determine the long-run impact of global temperature changes. In that case, we need $\zeta_s,\ \delta_s$ to be integrable to obtain a well-defined steady-state. This requirement rules out growth effects which would imply an economy that asymptotes to zero. In any case, we do not find any evidence supporting growth effects.

Because we focus on climate damages, we do not model emissions and associated externalities. Thus, the competitive equilibrium is efficient as is standard in the neoclassical

growth model.

4.3 The Social Cost of Carbon

In our framework, we define the Social Cost of Carbon as the one-time dollar amount $\mathcal{C}$ that households would pay at time 0 that would make them indifferent between a world with an additional ton of CO2 emitted at time 0, and a world starting in steady-state, without emissions, but having paid $\mathcal{C}$.

Given that we do not model emissions directly, we must map a one-time CO2 pulse into a temperature path in order to calculate the SCC. We follow Folini et al. (2024) and use the temperature response of global mean temperature to a CO2 pulse from Dietz et al. (2021), itself based on Joos et al. (2013). Dietz et al. (2021) report the temperature response in multiple state-of-the-art atmospheric circulation and radiative forcing models. These temperature responses let us map welfare losses into the SCC consistently with state-of-the-art goescience.

We denote by $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$ the path of excess warming implied by a one-time pulse of a single ton of CO2 emitted at time 0. The multi-model mean response in Dietz et al. (2021) indicates that temperature rise steadily and approximately exponentially. Temperature stabilizes at 0.002°C 15 years out after a short and moderate overshoot.

We purposefully remain *conservative* and use the lower end of the range of temperature responses from Dietz et al. (2021): we define $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$ as half of the the multi-model mean in Dietz et al. (2021). Doing so additionally ensures that historical emissions are consistent with historical warming data. When we use the multi-model mean, our SCC numbers double. Our welfare numbers would remain unchanged as they do not depend on the temperature response to a CO2 pulse.

We then construct productivity and capital depreciation paths $\{Z_t^{\text{SCC}}, \Delta_t^{\text{SCC}}\}_{t\geq 0}$ according to equation (5) in which we use the temperature path $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$ rather than a global warming scenario. The model delivers a path of value functions $\{V_t^{\text{SCC}}(K)\}_{t\geq 0}$, equilibrium capital stocks $\{K_t^{\text{SCC}}\}_{t\geq 0}$ with initial condition $K_0^{\text{SCC}} = K^{ss}$, leading to a path of realized values $\{V_t^{\text{SCC}}(K_t^{\text{SCC}})\}_{t\geq 0}$, in response to this CO2 pulse-induced warming.

Our definition requires that the SCC $\mathcal{C}$ be defined implicitly as:

$$V^{ss}(K^{ss} - \mathcal{C}) = V_0^{\text{SCC}}(K^{ss}), \tag{7}$$

where *ss* superscripts denote steady-state quantities.

To gain intuition, consider the case when the SCC is not too large. Then, a first order perturbation implies that the SCC satisfies:

$$\mathcal{C} = \int_0^\infty e^{-\rho t} u'(C^{ss})(C^{ss} - C_t^{\text{SCC}})dt = \frac{1}{\rho}\frac{C^{ss} - \overline{C}^{\text{SCC}}}{C^{ss}}, \tag{8}$$

where $\frac{C^{ss} - \overline{C}^{\text{SCC}}}{C^{ss}}$ is defined as the consumption-equivalent welfare loss from the warming implied by the CO2 pulse.

The first condition states that the SCC is the present discounted value of consumption losses due to the warming implied by the CO2 pulse. Using the definition of consumption-equivalent welfare losses, the second condition indicates that the SCC is equal to the present discounted value of the consumption-equivalent welfare losses warming induced by the CO2 pulse. While these conditions are useful to gain intuition, in our quantification we always use the nonlinear definition (7) that accounts for a time-varying marginal rate of substitution.

4.4 Estimation Strategy

Our next step is to estimate the structural damage functions ζ_s, δ_s. To do so, we match the reduced-form impulse response functions of GDP and capital to global temperature shocks from Figures 8 and 10. We proceed in two steps.

In the first step, we calibrate our model based on standard values from the literature, with the exception of our damage functions. We set risk-aversion to $\gamma = 2$. The capital share is $\alpha = 0.33$. The baseline annual capital depreciation rate is $\Delta_0 = 0.08$. Our choice of annual discount rate $\rho = 0.02$ follows Rennert et al. (2022). However, we will assess the robustness of our results with respect to the discount rate.

In the second step, we invert our model to estimate our structural damage functions ζ_s, δ_s: the sequence of TFP and depreciation shocks that correspond to a temperature

shock. We leverage that the actual temperature shocks that arise during our sample are small and therefore imply output and capital fluctuations of the order of 1% (see Figure 7). Therefore, we can use a first-order perturbation of the model around the initial steady-state. Specifically, we consider a sequence of temperature shocks $\widehat{T}_t$. We denote by $\widehat{z}_t$ the resulting log deviation in TFP and by $\widehat{\Delta}_t$ the resulting level deviation capital depreciation rates. We denote by $\widehat{y}_t$, $\widehat{k}_t$ the log deviations in output and capital along the transition. We emphasize that we use log-linearization for *estimation* only, *not* for *counterfactuals*.

Proposition 1. *(Model inversion)*
There exists $\mathcal{K}_t(\widehat{z})$, $\mathcal{J}_{t,s}$ *given in Appendix B.3 such that, to a first order in* $\widehat{T}_t$*:*

$$\widehat{y}_t = \widehat{z}_t + \alpha \widehat{k}_t \qquad\qquad \widehat{k}_t = \mathcal{K}_t(\widehat{z}) + \int_0^\infty \mathcal{J}_{t,s}\widehat{\Delta}_s ds$$

Proof. See Appendix B.3. □

Proposition 1 delivers an identification result. Given observed output and capital responses $\widehat{y}_t$, $\widehat{k}_t$, we can recover the underlying sequence of productivity shocks $\widehat{z}_t$ and capital depreciation shocks $\widehat{\Delta}_t$.

The first equation of Proposition 1 lets us recover the sequence of productivity shocks directly from the observed output and capital responses—this relationship is immediate from the production function.

The main content of Proposition 1 lies in the second equation. By log-linearizing the equilibrium conditions of the model and solving explicitly for the equilibrium sequence of capital, we relate capital deviations to the sequence of capital depreciation rates through the sequence-space Jacobian $\mathcal{J}_{t,s}$ (Auclert et al., 2021; Bilal and Goyal, 2023) given productivity shocks embeded in $\mathcal{K}_t(\widehat{z})$. In the context of the neoclassical growth model, we can solve in closed form for this Jacobian as a function of parameters and steady-state objects. When $\mathcal{J}_{t,s}$ is invertible, the capital depreciation shocks are identified. We use Proposition 1 to obtain the sequence of TFP and depreciation rates $\widehat{z}_t$, $\widehat{\Delta}_t$ that correspond to any sequence of temperature shocks $\widehat{T}_t$.

We use these observations to estimate ζ_s, δ_s. However, we cannot yet directly use the output and capital impulse response functions from Figures 8 and 10. These impulse response functions correspond to a *persistent* underlying global temperature shock, i.e.

a shock that increases global mean temperature persistently as shown in Figure 9. The structural damage functions ζ_s, δ_s correspond to the impact of a *transitory* temperature shock. This observation is critical: omitting to account for the internal persistence of the temperature shock would overstate the impact of global warming (Nath et al., 2023). Thus, we deconvolute the data before using Proposition 1.

We construct the impulse response function to a one-time transitory temperature shock with linear combinations of the impulse response function to the observed, persistent temperature shock. This approach follows Sims (1986). It is equivalent to using a recursive approach. Indeed, denote by $\widetilde{y}_t$ the unknown impulse response function of output to a transitory temperature shock. In discrete data and under linearity: $\widehat{y}_t = \sum_{s=0}^{t} \widehat{T}_{t-s}\widetilde{y}_s$. We then obtain $\widetilde{y}_t = \left(\widehat{y}_t - \sum_{s=0}^{t-1} \widehat{T}_{t-s}\widetilde{y}_s\right) / \widehat{T}_0$ recursively.

With the deconvoluted impulse response functions of output and capital to a one-time unit transitory temperature shock at hand, we use Proposition 1 and obtain the corresponding shocks $\widehat{z}_t$, $\widehat{\Delta}_t$. We then identify $\zeta_s = \widehat{z}_s$ and $\delta_s = \widehat{\Delta}_s / \Delta_0$.

In practice, we face two additional challenges. We address both of them by imposing a smooth functional form for our structural damage function. We constrain ζ_s, δ_s to be of the form $A\left(e^{-Bs} - e^{-Cs}\right)$.

The first challenge that our constrained estimation addresses is that we can only estimate the impulse response functions $\widehat{y}_t$, $\widehat{k}_t$ up to a finite horizon. By contrast, Proposition 1 requires the entire impulse response function. We cannot simply set the capital impulse response to 0 from year 11 onwards, as this would imply a large underlying capital windfall gain for the economy. By constraining the shape of the structural damage functions, we use our 10 data points to estimate 3 parameters per damage function.

The second challenge is to discipline the long-run effects of temperature shocks. By constraining the structural damage functions, we ensure that the effects of transitory temperature changes vanish in the very long run. If we estimated the structural damage functions entirely unconstrained and with a longer horizon, temperature shocks could potentially have longer-ranging but extremely imprecisely estimated effects. Therefore, our approach is *conservative* in that it limits the long-run impact of a one-time transitory temperature shock. In any case, as Figure 13 below makes clear, the *deconvoluted* GDP impulse response fully mean-reverts to 0 after 10 years.

Hence, instead of exactly inverting the model, we estimate A, B and C for ζ_s, δ_s separately using Ordinary Least Squares (OLS) to minimize the squared deviations from the equations in Proposition 1 for the first 10 years only.

4.5 Estimation Results

Figure 13 shows our estimation results. The panels in column (a) display the output (i) and capital (ii) responses to internally persistent temperature shocks, in the model and in the data. By construction, these responses account for the persistent increase in global temperature levels in response to global temperature shocks as estimated in the data (see Figure 9). The dashed lines are the impulse responses as estimated in Section 3. The solid lines show the corresponding responses in the estimated model. Our model closely tracks the empirical responses. Of course, the fit of the model relies on our constrained functional form: if we did not constrain the damage function, the fit would be one-to-one.

The dashed lines in column (b) show the deconvoluted responses of output and capital that we use for estimation. These are the responses to a one-time transitory global temperature shock of 1°C. As expected, the output and capital responses are smaller, given the considerable degree of internal persistence of the estimated global temperature shock. However, they remain sizeable and peak at around -5%, respectively. The solid lines show again the model fit under our constrained functional form.

Finally, the panels in column (c) depict the estimated structural damage functions, ζ_s and δ_s. The solid lines with circles represent the responses to a one-time transitory global temperature shock of 1°C. It implies a short-run productivity loss of nearly 2.5% and an increase in the capital depreciation rate of 0.3 p.p. Despite the temperature shock being transitory, the impact on productivity and capital depreciation decays only slowly and largely persists for up to 10 years. The bootstrapped confidence bands reflect the confidence intervals around our empirical output and capital. The response of productivity is more precisely estimated as the response of capital depreciation.

We illustrate the importance of isolating the output and capital responses to a transitory shock by showing the damage functions if one mistakenly targeted the responses to a persistent temperature shock (the dashed lines with circles). In that case, the productivity effect exceeds 5% and capital depreciation approaches 0.6 p.p. in the medium run. Thus,

Figure 13: Productivity and Capital Depreciation after Global Temperature Shocks

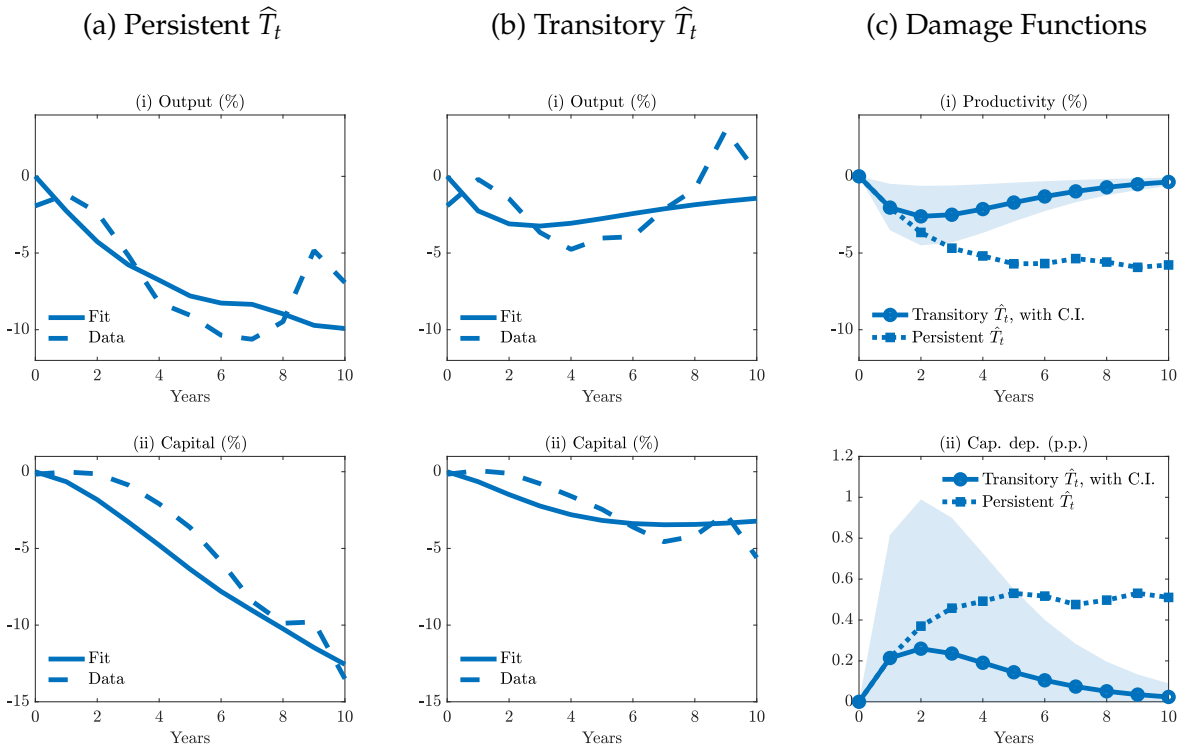

Notes: The figure shows our estimation results from matching the model impulse responses to the empirical responses to global temperature shocks. The four left panels show the output and capital responses in the data and the model. Column (a) shows the responses to persistent temperature shocks. Column (b) shows the responses to transitory temperature shocks used in the estimation. Column (c) plots the implied productivity and capital depreciation shocks, together with 68% confidence intervals (shaded area) for the transitory case based on 1000 bootstrap draws from the empirical output, capital and temperature IRFs.

when constructing counterfactual paths due to global warming, it is crucial to correctly cumulate the effects to *transitory* rather than *persistent* temperature shocks. Otherwise, the impacts of climate change would be overstated.

How do the productivity and capital depreciation effects of global temperature shocks compare to previous estimates? Answering this question is challenging because little work directly estimates the impact of global temperature shocks. We can compare our estimates to outcomes of structural models that build up from micro-level estimates of the impact of extreme events. For instance, Bilal and Rossi-Hansberg (2023) find that a permanent 1°C rise in global mean temperature implies a 1% decline in TFP and a 0.3 p.p. rise in the capital depreciation rate for the United States. Our estimates for persistent temperature shocks—closer to a permanent shock—are of the same order of magnitude

though larger for productivity, likely reflecting that the United States is more resilient to extreme events than lower-income countries.

How do the productivity and capital depreciation effects of global temperature shocks compare to those associated with local temperature shocks? Given that the empirical responses are substantially larger in the data for global temperature shocks (Figure 7), such shocks likely also imply larger damages. To answer this question quantitatively, we repeat our estimation but targeting the impulse response functions following local temperature shocks.

Figure 14: Productivity and Capital Depreciation after Local Temperature Shocks

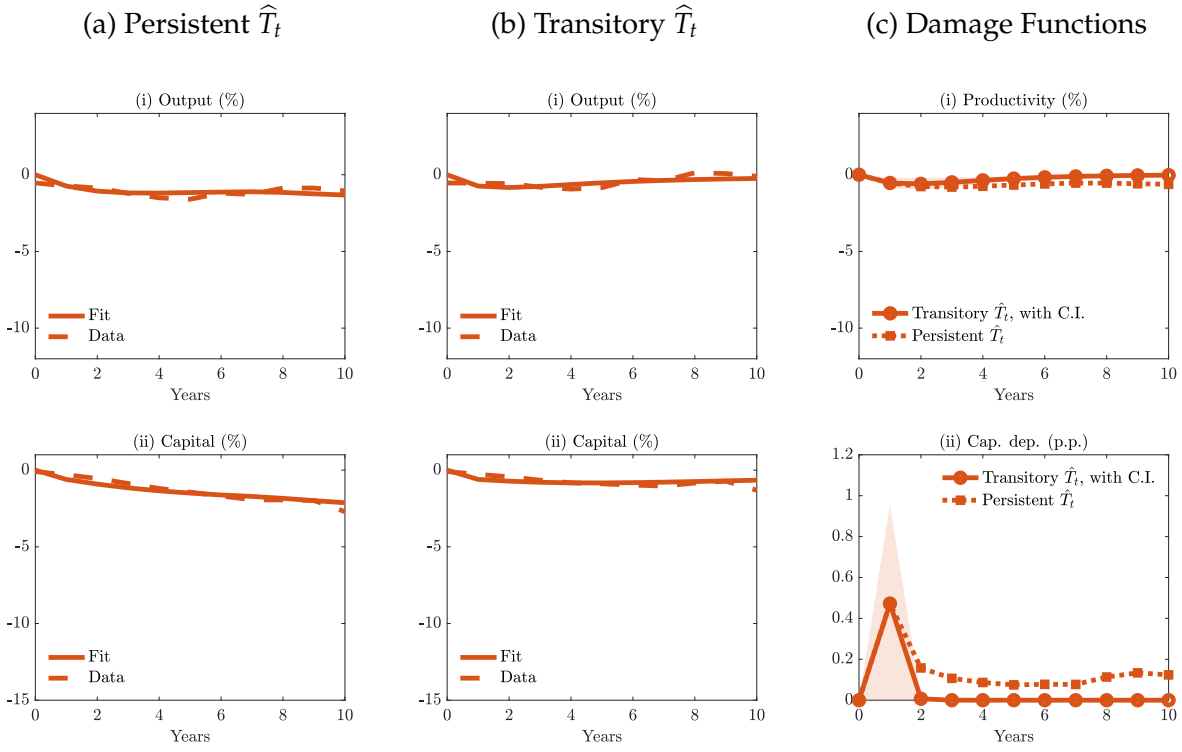

Notes: The figure shows our estimation results from matching the model impulse responses to the empirical responses to local temperature shocks. The four left panels show the output and capital responses in the data and the model. Column (a) shows the responses to persistent temperature shocks. Column (b) shows the responses to transitory temperature shocks used in the estimation. Column (c) plots the implied productivity and capital depreciation shocks, together with 68% confidence intervals (shaded area) for the transitory case based on 1000 bootstrap draws from the output, capital and temperature IRFs.

Figure 14 displays the productivity and capital depreciation effects of local temperature shocks. The productivity effect of local temperature shocks is five times smaller than under global temperature shocks. The impact response of capital depreciation is some-

what larger but vanishes immediately, so that the cumulated impact is substantially lower than under global shocks. We conclude that global temperature shocks have much larger effects on economic fundamentals.

5 The Welfare Impact of Climate Change

In this section, we use our estimated model to evaluate the consequences of climate change for welfare and the SCC.

5.1 Representing Climate Change

To evaluate the consequences of climate change, our first step is to construct a path for global mean temperature. We choose 2024 as our baseline year $t = 0$. Our central warming scenario is one where the world warms by 3°C above pre-industrial levels by 2100, after which temperature asymptotes to 3.3°C above pre-industrial levels in the very long-run. This scenario is *conservative* since the IPCC considers business-as-usual to imply over 4°C of warming by 2100 (Lee et al., 2023). We calibrate an exponential convergence such that the warming path matches these two targets and denote by $\widehat{T}_t = T_t - T_0$ the corresponding path. Crucially, given that the world has warmed by approximately 1°C since pre-industrial times, such a scenario implies 2°C of *additional* warming since $t = 0$ (2024) by year $t = 76$ (2100).

To highlight the central role of global temperature shocks, we construct two counterfactuals. In the first counterfactual, we start in steady-state and use the structural damage functions estimated under *global* temperature shocks to construct changes in productivity and capital depreciation:

$$Z_t^{\text{global}} = Z_0 \exp\left(\int_0^t \zeta_s^{\text{global}} \widehat{T}_{t-s} ds\right) \qquad \Delta_t^{\text{global}} = \Delta_0 \exp\left(\int_0^t \delta_s^{\text{global}} \widehat{T}_{t-s} ds\right),$$

where the estimates for ζ_s^{global}, δ_s^{global} correspond to the impact of a one-time, transitory global temperature shock (Figure 13(c), solid line with circles). In the second counterfactual, we start in steady-state and use the structural damage functions estimated under

local temperature shocks to construct changes in productivity and capital depreciation:

$$Z_t^{\text{local}} = Z_0 \exp\left(\int_0^t \zeta_s^{\text{local}} \widehat{T}_{t-s} ds\right) \qquad \Delta_t^{\text{local}} = \Delta_0 \exp\left(\int_0^t \delta_s^{\text{local}} \widehat{T}_{t-s} ds\right),$$

where the estimates for ζ_s^{local}, δ_s^{local} correspond to the impact of a one-time, transitory local temperature shock (Figure 14(c), solid line with squares).

We then compare allocations and welfare in an economy that warms according to $\widehat{T}_t$, to allocations and welfare in an economy in which $\widehat{T}_t \equiv 0$. We represent welfare losses from climate change as an equivalent percent decline in steady-state consumption. That is, a 1% welfare loss under climate change means that households would be as well off if there was no climate change, but they permanently gave up 1% of their steady-state consumption. To solve for counterfactuals, we emphasize again that we use standard global numerical methods to obtain the global solution—we only use log-linearization for estimation.

To calculate the SCC, we construct analogues of these quantities for global and local damage functions following a one-time pulse of a single ton of CO2 as described in Section 4.3. Importantly, the SCC calculations are independent from the global warming scenario because they rely on the temperature response to a given CO2 pulse $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$. Conversely, the welfare calculations are independent from $\{\widehat{T}_t^{\text{SCC}}\}_{t\geq 0}$.

5.2 Welfare and the SCC

Figure 15 presents our main results. Panel (a) depicts the path of global mean temperatures. Panel (b) reveals that output starts dropping rapidly when global mean temperatures rise, relative to a world that is not warming. By 2050, output declines by 30%. In 2100, output is 50% below what it would have been without climate change. This substantial decline in output reflects accumulated productivity losses that eventually reach 30% as well as a steep 4 p.p. rise in the capital depreciation rate, representing a 50% increase.

Panel (c) highlights the combined adverse impact of lower productivity and higher depreciation rates on capital accumulation. Initially, investment rises as households anticipate lower income going forward and therefore save, following standard permanent income logic. Rapidly however, capital starts decumulating under the combined pressure

Figure 15: Transitional Dynamics Under Climate Change

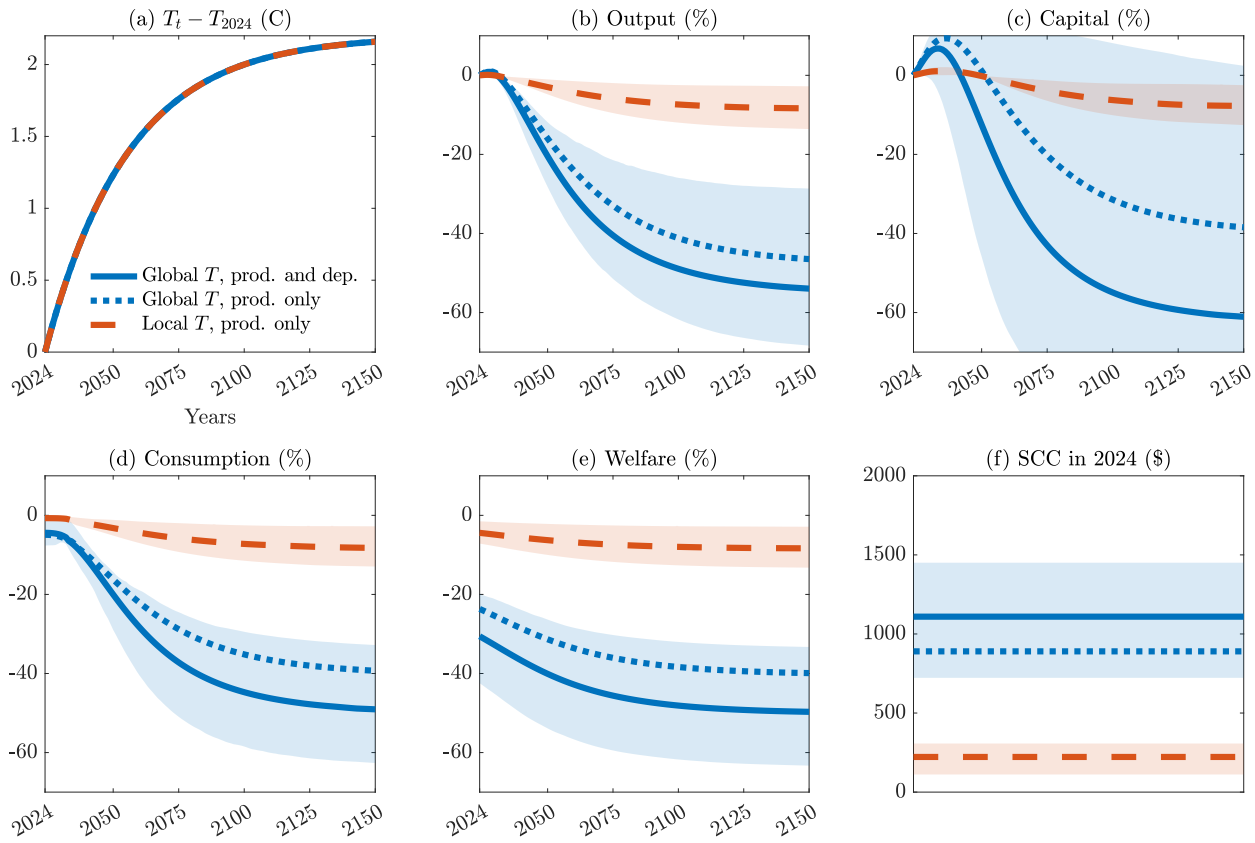

Notes: The figure shows the transitional dynamics of our estimated model under our scenario where the worlds warms by 3°C above pre-industrial levels by 2100. The blue solid lines represent transitional dynamics when we estimate the model based on global temperature shocks, together with 68% confidence intervals (shaded blue). The dotted blue lines represent transitional dynamics when we only use productivity damages under global temperature shocks. The dashed red lines represent transitional dynamics when we use only productivity shocks estimated under local temperature shocks, together with 68% confidence intervals (shaded red). Confidence intervals based on 1000 bootstrap draws from output, capital and temperature IRFs.

of lower output and higher depreciation. By 2100, capital is 60% below what it would have been without climate change. Panel (d) reveals that consumption declines as much as output, eventually exceeding a 50% loss in the long run.

This substantial decline in consumption translates into a large welfare loss. Panel (e) shows that the welfare impact of climate change amounts to a 31% welfare loss in consumption equivalent percent. This welfare loss exceeds the consumption impact as households discount but value future declines in consumption as well. As temperature

keeps rising, welfare continues to decline and reaches a 52% loss.

Our results indicate that the impact of climate change is substantial. In welfare terms, the cost of climate change is 640 times the cost of business cycles, or ten times the cost of moving from current trade relations to complete autarky. Perhaps most strikingly, in terms of output, capital, consumption, and thus welfare, climate change is comparable in magnitude to the effect of fighting a major war domestically. However, climate change is *permanent*. Thus, the losses from living in a world with climate change relative to a world without it are comparable to fighting a major war domestically, *forever*.

Our results also shed light on how much economic growth was missed because of past climate change. When we start the economy in 1960 and feed in the historical path of warming until 2024, our counterfactuals indicate that world GDP per capita would be 37% higher today had no warming occurred between 1960 and 2024 instead of the actual 0.75°C increase in global mean temperatures. This difference translates into a 29% reduction in the annual growth rate of the world economy since 1960 (half a percentage point).

Panel (f) uses our structural damage function to construct the SCC. We obtain a SCC of \$1,056/tCO2. This value is more than six times larger than the \$185/tCO2 value in Rennert et al. (2022). There are two possible reasons why we obtain a large SCC and substantial welfare costs of climate change. The first possible reason is our focus on global temperature shocks. The second possible reason is that we include damages to productivity and capital depreciation, rather than productivity alone as in most previous work.

We demonstrate that our focus on global temperature shocks is the main driver of our conclusions. We do so by re-estimating our model based on the impact of local temperature shocks on productivity only, consistently with previous research. In that case, and consistently with previous estimates (Nordhaus, 1992; Dell et al., 2012; Burke et al., 2015; Nath et al., 2023), Panels (b)-(e) show that climate change then implies present value welfare costs of 4% and the SCC is \$151/tCO2, just below the value in Rennert et al. (2022). When we estimate our model based on the impact of global temperature shocks on productivity only, we obtain a welfare loss of 22% and a SCC of \$842/tCO2, five to six times larger than with local shocks. Including damages to capital depreciation further increases these values to our main results. Our bootstrapped confidence intervals highlight the un-

certainty around these point estimates. The 68% confidence interval for the SCC ranges from \$723 to \$1,451/tCO2. Despite non-trivial uncertainty, even the lower bound of that confidence interval is several times larger than conventional SCC estimates.

5.3 Sensitivity

Given the sizeable magnitude of our results, we investigate which parameters may be driving them. Figure 16 displays a sensitivity analysis with respect to two key parameters: the discount rate ρ and 2100 global mean temperature.

Panel (a) shows the welfare losses as a function of the discount factor ρ, and panel (b) shows the corresponding SCC. The solid line depicts losses when using global temperature shocks to estimate our model, and the dashed line depicts losses when using local temperature shocks. As expected, a higher discount rate lowers welfare losses and the SCC: households then value damages that are far in the future by less. Our baseline discount rate value of $\rho = 0.02$ is consistent with Rennert et al. (2022) and with the secular decline in interest rates. However, even at much higher discount rates—up to 0.04 or 0.05—we still obtain sizable losses in excess of 20% in consumption equivalent. The corresponding SCC remains two to three times as large as the higher end of previous estimates. By contrast, as we approach very low discount rates consistent with Stern (2006), welfare losses exceed 40% and the SCC rises above \$3,000/tCO2. Welfare losses are less sensitive to the discount rate than the SCC because welfare losses represent an annualzed flow of losses, while the SCC is a discounted stock valuation as shown in equation (8).

Panels (c) and (d) show welfare losses and the SCC when we vary 2100 temperature relative to pre-industrial levels. Of course, in Panel (f), the SCC is independent from the warming scenario because it only depends on the temperature response to a CO2 pulse.

Welfare losses under 20% materialize only at very low warming scenarios of 1.5°C since pre-industrial levels by 2100. The IPCC evaluates that the world is on track for 4°C above pre-industrial levels under business as usual: global mean temperatures already largely exceed 1°C since pre-industrial levels, and some estimates indicate that 2023 reached 1.48°C since pre-industrial levels. By contrast, more pessimistic scenarios under which global mean temperatures reach 5°C since pre-industrial levels in 2100 lead to present value welfare losses of 60%.

Figure 16: Welfare and the Social Cost of Carbon under Alternative Choices

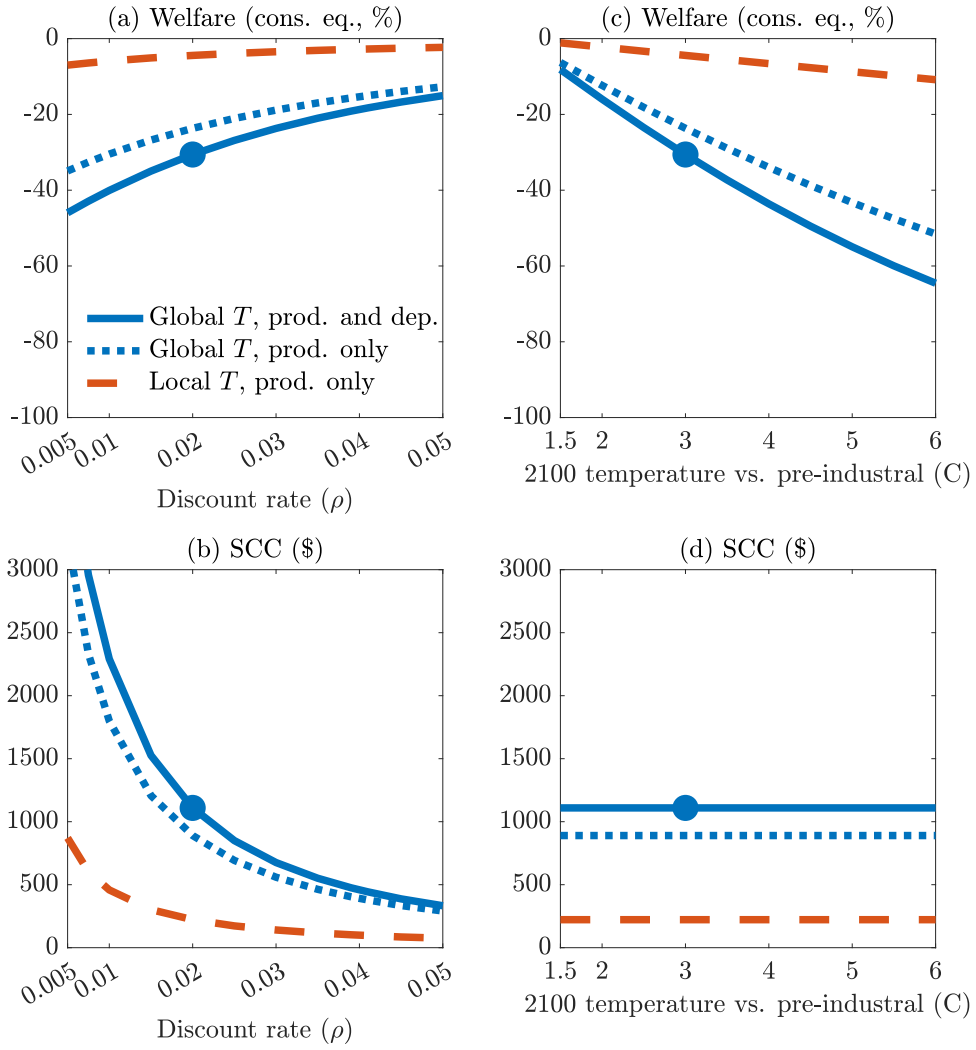

Notes: The figure shows the sensitivity of our model-implied welfare costs and social cost of carbon, both in 2024, with respect to the discount rate (ρ) and the 2100 global mean temperature. The solid blue line depicts the effects when estimating the model estimated using global temperature shocks (baseline). The dashed red line depicts the effects when estimating the model with local temperature shocks.

Our sensitivity analysis indicates that substantial climate damages occur over a wide range of specification choices. We conclude that climate change poses a substantial threat to the world economy.

6 Conclusion

In this paper, we demonstrate that the impact of climate change on economic activity is substantial. We leverage natural climate variability in global mean temperature to obtain time-series estimates that are representative of the overall impact of global warming. We find that a 1°C rise in global temperature causes global GDP to persistently decline, with a peak loss at 12%. This large effect is due to an associated surge in extreme climatic events. By contrast, local temperature shocks used in the traditional panel literature lead to a minimal rise in extreme events and to much smaller economic effects. Together, our results imply a SCC of $1,056/tCO2 and a 31% welfare loss from a moderate warming scenario. These effects are comparable to having a major war fought domestically, forever.

Not only do our results indicate that climate change represents a major threat to the world economy, they also have salient consequences for decarbonization policy. Many decarbonization interventions cost between $27 and $95 per ton of CO2 abated (Bistline et al., 2023). A conventional SCC value of $151/tCO2 implies that these policies are cost-effective only if governments internalize benefits to the entire world, as captured by the SCC. However, a government that only internalizes domestic benefits values mitigation benefits using a Domestic Cost of Carbon (DCC). The DCC is always lower than the SCC because damages to a single country are less than to the entire world. For instance, under conventional estimates based on local shocks, the DCC of the United States is $30/tCO2, making unilateral emissions reduction prohibitively expensive. Under our new estimates however, the DCC of the United States becomes $211/tCO2 and thus largely exceeds policy costs. In that case, unilateral decarbonization policy is cost-effective for the United States.

References

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models". *Econometrica* 89.5, pp. 2375–2408.

Bansal, Ravi and Marcelo Ochoa (2011). "Temperature, Aggregate Risk, and Expected Returns". *NBER Working Paper Series 17575*.

Barro, Robert J. (Aug. 2006). "Rare Disasters and Asset Markets in the Twentieth Century". *The Quarterly Journal of Economics* 121.3, pp. 823–866.

Berg, Kimberly A., Chadwick C. Curtis, and Nelson Mark (2023). *GDP and temperature: Evidence on cross-country response heterogeneity*. Tech. rep. National Bureau of Economic Research.

Bilal, Adrien and Shlok Goyal (2023). "Some Pleasant Sequence-Space Arithmetic in Continuous Time". *Working Paper*.

Bilal, Adrien and Esteban Rossi-Hansberg (2023). "Anticipating Climate Change Across the United States". *National Bureau of Economic Research Working Paper 31323*.

Bistline, John, Neil R. Mehrotra, and Catherine Wolfram (2023). "Economic Implications of the Climate Provisions of the Inflation Reduction Act". *Brookings Papers on Economic Activity*.

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

Cerra, Valerie and Sweta Chaman Saxena (2008). "Growth Dynamics: The Myth of Economic Recovery". *American Economic Review* 98.1, pp. 439–57.

Conte, Bruno, Klaus Desmet, and Esteban Rossi-Hansberg (2022). *On the Geographic Implications of Carbon Taxes*. Working Paper 30678. National Bureau of Economic Research.

Cruz, José-Luis and Esteban Rossi-Hansberg (2023). "The Economic Geography of Global Warming". *Review of Economic Studies*, forthcoming.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2012). "Temperature shocks and economic growth: Evidence from the last half century". *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

— (2014). "What Do We Learn from the Weather? The New Climate–Economy Literature". *Journal of Economic Literature* 52.3, pp. 740–798.

Deryugina, Tatyana (2013). "The role of transfer payments in mitigating shocks: Evidence from the impact of hurricanes". *Available at SSRN 2314663*.

Deschênes, Olivier and Michael Greenstone (2011). "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US". *American Economic Journal: Applied Economics* 3.4, pp. 152–85.

Desmet, Klaus, Robert E. Kopp, Scott A. Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H. Strauss (2021). "Evaluating the Economic Cost of Coastal Flooding". *American Economic Journal: Macroeconomics* 13.2, pp. 444–86.

Desmet, Klaus and Esteban Rossi-Hansberg (2015). "On the Spatial Economic Impact of Global Warming". *Journal of Urban Economics* 88, pp. 16 –37.

Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans (2021). "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists* 8.5, pp. 895–921.

Driscoll, John C. and Aart C. Kraay (1998). "Consistent covariance matrix estimation with spatially dependent panel data". *Review of Economics and Statistics* 80.4, pp. 549–560.

Folini, Doris, Aleksandra Friedl, Felix Kübler, and Simon Scheidegger (Jan. 2024). "The Climate in Climate Economics*". *The Review of Economic Studies*, rdae011.

Hamilton, James D. (2018). "Why you should never use the Hodrick-Prescott filter". *Review of Economics and Statistics* 100.5, pp. 831–843.

Hsiang, Solomon M. and Amir S. Jina (2014). "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones". *National Bureau of Economic Research Working Paper 20352*.

Joos, F., R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver (2013). "Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis". *Atmospheric Chemistry and Physics* 13.5, pp. 2793–2825.

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". *American Economic Review* 95.1, pp. 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2020). "The effects of quasi-random monetary experiments". *Journal of Monetary Economics* 112, pp. 22–40.

Kaufmann, Robert K., Heikki Kauppi, and James H. Stock (2006). "Emissions, concentrations, & temperature: a time series analysis". *Climatic Change* 77, pp. 249–278.

Kose, M. Ayhan, Naotaka Sugawara, and Marco E. Terrones (2020). "Global recessions".

Krusell, Per and Anthony A. Smith (2022). *Climate change around the world*. Tech. rep. National Bureau of Economic Research.

Lee, Hoesung, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. (2023). "IPCC, 2023: Climate Change 2023: Synthesis Report, Summary for Policymakers. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland."

Montiel Olea, José Luis and Mikkel Plagborg-Moller (2021). "Local projection inference is simpler and more robust than you think". *Econometrica* 89.4, pp. 1789–1823.

Nath, Ishan (2022). "Climate Change, The Food Problem, and the Challenge of Adaptation through Sectoral Reallocation". *Working Paper*.

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (2023). "How Much Will Global Warming Cool Global Growth?" *NBER Working Paper Series*.

National Oceanic and Atmospheric Administration (2005). *Volcanos and Climate*. https://gml.noaa.gov/grad/volcano.html#:~:text=Volcanos%20that%20pose%20the%20greatest,up%20to%20a%20few%20years..

— (2009). *Climate Change: Incoming Sunlight*. https://www.climate.gov/news-features/understanding-climate/climate-change-incoming-sunlight#:~:text=During%20strong%20solar%20cycles%2C%20the,0.1%20degrees%20Celsius%20or%20less..

— (2023). *What are El Niño and La Niña?* https://oceanservice.noaa.gov/facts/ninonina.html.

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

Nordhaus, William D. (1992). "An optimal transition path for controlling greenhouse gases". *Science* 258.5086, pp. 1315–1319.

— (2013). "Integrated economic and climate modeling". *Handbook of Computable General Equilibrium Modeling*. Vol. 1. Elsevier, pp. 1069–1131.

Phan, Toan and Felipe F. Schwartzman (2023). "Climate defaults and financial adaptation".

Reinhart, Carmen M. and Kenneth S. Rogoff (2009). "The aftermath of financial crises". *American Economic Review* 99.2, pp. 466–472.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. (2022). "Comprehensive evidence implies a higher social cost of CO2". *Nature* 610.7933, pp. 687–692.

Rudik, Ivan, Gary Lyn, Weiliang Tan, and Ariel Ortiz-Bobea (2022). "The Economic Effects of Climate Change in Dynamic Spatial Equilibrium".

Seppanen, Olli, William J. Fisk, and David Faulkner (2003). "Cost benefit analysis of the night-time ventilative cooling in office building".

Sims, Christopher A. (1986). "Are forecasting models usable for policy analysis?" *Quarterly Review* 10.Win, pp. 2–16.

Stern, N. (2006). *Stern Review: The Economics of Climate Change*.

Stern, Nicholas, Joseph Stiglitz, and Charlotte Taylor (2022). "The economics of immense risk, urgent action and radical change: towards new approaches to the economics of climate change". *Journal of Economic Methodology* 29.3, pp. 181–216.

Tran, Brigitte Roth and Daniel J. Wilson (2023). "The local economic impact of natural disasters". Federal Reserve Bank of San Francisco.

Online Appendix

Unveiling the Macroeconomic Impact of Climate Change

Adrien Bilal[†]

Diego R. Känzig[‡]

Contents

[†]Harvard University, CEPR and NBER. E-mail: adrienbilal@fas.harvard.edu. Web: www.sites.google.com/site/adrienbilal

[‡]Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com

A Empirics

A.1 Data

A.1.1 Economic Data

We source economic information on GDP, population, consumption, investment and productivity for a comprehensive selection of countries around the world from the Penn World Tables (PWT; Feenstra et al., 2015). Our main output measure is real GDP per capita from the national accounts (rgdpna/pop). For our country comparisons by income, we use (expenditure-side) real GDP per capita at chained PPPs (rgdpe/pop). For capital, we use the capital stock from national accounts (rnna). Investment, we compute using data on capital and capital depreciation (delta) based on the capital accumulation equation $I_t = K_t - (1 - \delta_t)K_{t-1}$. For total factor productivity, we also use the measure based on national accounts (rtfpna). We compute a measure of labor productivity based on output and employment data (rgdpna/emp).[1]

The PWT data set is commonly used in the literature and of high quality. However, as an alternative, we also use data from the World Bank. One limitation of both of these data sets is that they only go back to the 1950s or 1960s. To extend our analysis to a longer historical sample period, we therefore also include data from the Macro-history Database (Jordà et al., 2017), which features high-quality economic data for 18 developed countries starting in the late 19th century.

A.1.2 Climate Data

Gridded temperature datasets. Our primary gridded temperature dataset is Berkeley Earth, due to its geographic coverage, temporal coverage, and update frequency.

We obtain gridded temperature anomalies (using air temperatures at sea ice) at a daily and monthly frequency between 1850 and 2022 from Berkeley Earth (2023), at a resolution of $1^\circ \times 1^\circ$ latitude-longitude grid. Temperature anomalies are deviations from the climatology, which is measured as the 1951-1980 mean temperature (Rohde and Hausfather,

[1] We use employment as a proxy for the labor input because the data on average hours is not very well populated.

2020). Grid-level temperature levels are constructed by adding the grid-level climatology to the grid-level anomaly series.

We also obtain gridded estimates of temperature, wind, and precipitation at a daily frequency between 1901 and 2019 from the Inter-Sectoral Impact Model Intercomparison Project (ISIMIP), at a 0.5° spatial resolution (Lange et al., 2023).

To assess the sensitivity of the results to the gridded temperature data used, we obtain alternate, prominent datasets used in the literature. We obtain gridded temperature levels (surface air temperature) at a monthly frequency between 1948 and 2014 from the Princeton Global Forcing Dataset (version 2) constructed by Sheffield et al. (2006), a later version of which was used, for instance, by Nath et al. (2023). Additionally, we obtain the gridded temperature levels (surface air temperatures) at a monthly frequency between 1900 and 2014 from the Willmott and Matsuura, University of Delaware Dataset (version 4.01) (Matsuura and National Center for Atmospheric Research Staff, 2023), earlier versions of which were used, for instance, by Dell et al. (2012) and Burke et al. (2015).

Aggregation of gridded temperature datasets. To aggregate the gridded temperature datasets to the global or country level we consider two different type of weights. One approach is to use area weights. Specifically, we use the area of the grid, calculated using the latitude and longitude. Alternatively, we use population weights. In that case, we use the grid-level population count in 1000 as weights, obtained from the Center for International Earth Science Information Network (CIESIN), Columbia University (2018).

Global temperatures. We obtain land and ocean surface temperature anomalies (in degrees Celsius) at an annual frequency between 1850 and 2022 from NOAA National Centers for Environmental Information (2023a). Temperature anomalies are deviations from the climatology, which is measured as the 1901-1000 mean temperature, 13.9 degree Celsius (NOAA National Centers for Environmental Information, 2023b). Temperature levels are constructed by adding the climatology to the anomaly series.

We also obtain the combined land-surface air and sea-surface water temperature anomalies (in degrees Celsius) at an annual frequency between 1880 and 2022 from Lenssen et al. (2019) and NASA Goddard Institute for Space Studies (2023). Tempera-

ture anomalies are deviations from the climatology, which is measured as the 1951-1980 mean temperature, approximately 14 degree Celsius (NASA Earth Observatory, 2020). Temperature levels are similarly constructed by adding the climatology to the anomaly series.

As a quality check of the gridded temperature data, we compute population- and area-weighted global temperature measures and compare them to the official measures from NOAA and NASA. Note that both official measures follow an area-weighted aggregation scheme. Reassuringly, aggregating the Berkeley Earth gridded temperature data using area weights to obtain a global temperature measure produces a series that is virtually perfectly correlated with both the NOAA and NASA global temperature series: we find that the measures based on all these different data sets align very well, as shown in Figure A.1.

Figure A.1: Global Average Temperature Since 1950

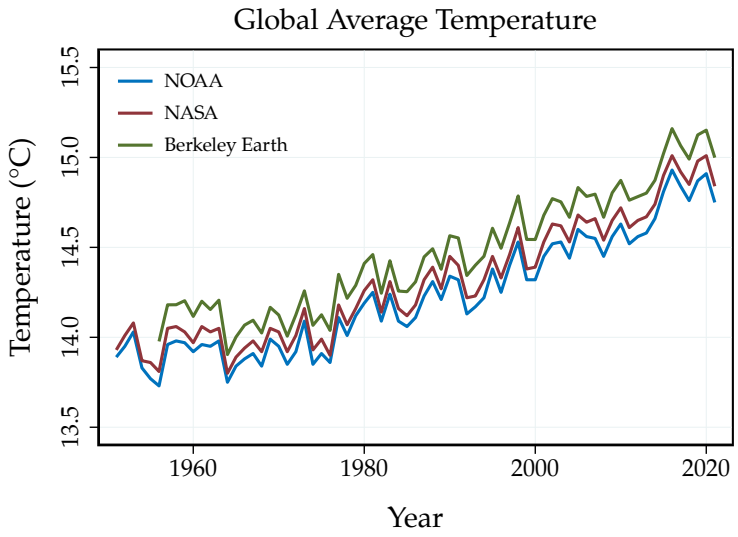

Notes: The figure shows the evolution of global average temperature. The NOAA and NASA measures are constructed by adding the climatology to the official anomaly series. The Berkeley Earth measure is constructed by first, obtaining grid-level temperature levels by adding the grid-level climatology to the grid-level anomaly series, and second, aggregating the grid-level temperature levels using area weights. We plot the Berkeley Earth series starting 1956, following which the percentage of monthly grid-level missing observations is consistently below ≈2%.

Country-level temperatures. We use the Berkeley Earth gridded temperature data to construct population- and area-weighted country-level mean temperatures. In our anal-

yses, we use population-weighted temperature as the baseline, however, using area-weighted measures produces very similar results. To assess the sensitivity of the results with respect to the gridded temperature data used, we similarly compute the population- and area-weighted country-level mean temperatures using the Princeton Global Forcing Dataset and the University of Delaware Dataset. We find that the results are consistent across different temperature datasets.

Extreme climatic events. We use the ISIMIP gridded estimates of temperature, wind, and precipitation at a daily frequency between 1901 and 2019 to construct extreme events indicators for each latitude-longitude grid. To define a threshold for extreme events, we use the percentiles of the distribution of the variables between 1901 and 1930, and define an extreme event as one where the realization of a variable was above a given percentile of its distribution. Specifically, we use the percentiles of the worldwide distribution to construct "absolute" extreme events indicators, and the percentiles of a country's distribution for "relative" indicators. We use the relative indicators as our baseline, however, our results are robust to using the absolute indicators.

To aggregate the variables across the grids to construct country-level measures, we use two methods. First, we construct the daily average of the variable for the country, and then compute the fraction of days in the year when the variable was above the threshold percentile (i.e., "country-level" extreme events indicator). Alternatively, we also compute the fraction of days in the year when the variable was above the threshold percentile *at the grid-level*, and then aggregate this indicator for the country (i.e., "cell-level" extreme events indicator). Of course, the threshold percentile changes across the definitions: for the former, we use the distribution of daily country-level averages, and for the latter, the distribution of daily grid-level observations between 1901 and 1930. Note that similar to the aggregation of gridded temperature datasets, we consider both area- and population-weights in both methods above. We use the country-level, area-weighted indicators as our baseline. However, the results are robust to using our alternative measures (cell-level and/or population-weighted).

A.2 Additional Figures and Tables

In this appendix, we present some additional figures and tables to complement the analysis in the main text.

A.2.1 Statistical Properties of Global Temperature Shocks

Serial correlation. Figure A.2 shows the autocorrelation function of the global temperature shock. The shocks are weakly autocorrelated. This is not too surprising, given that we construct the shocks as multi-step forecast errors. To account for this serial correlation, we therefore include two lags of the global temperature shock in our local projections. However, as we show in Appendix A.4, our results are robust with respect to the number of lags for the temperature shock.

Figure A.2: Autocorrelation of Global Temperature Shock

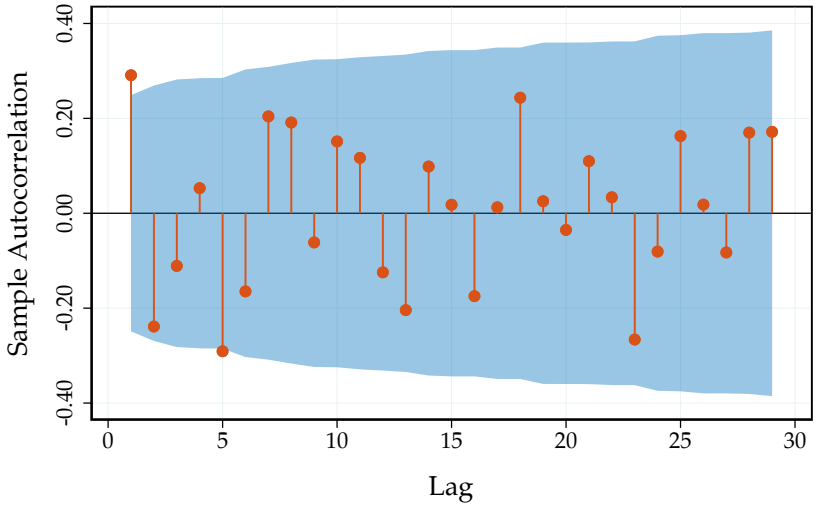

Notes: The figure shows the autocorrelation function of global temperature shocks, together with the 95% confidence bands, computed based on Bartlett's formula for MA(q).

Forecastablility. A desirable feature of "shocks" is that they should not be forecastable by past information (Ramey, 2016). In our context, if global tempreature shocks were forecastable by economic variables, this could point to reverse causality or other endogeneity threats. Thus, we check whether our temperature shocks are forecastable, considering a wide set of past macroeconomic or financial variables in a series of Granger-causality

tests. To account for the long and variable lags between emissions and warming, we conservatively include up to 8 years worth of lags.[2] Table A.1 reports the results. We find no evidence that macroeconomic or financial variables have any power in forecasting global temperature shocks. None of the selected variables Granger cause the series at conventional significance levels. The joint test is also insignificant.

Table A.1: Granger-causality Tests

Variable	p-value
Real GDP	0.494
Population	0.801
Brent price	0.756
Commodity price index	0.664
Treasury 1Y	0.830
Overall	0.825

Notes: The table shows the p-values of a series of Granger causality tests of the global temperature shock series using a selection of macroeconomic and financial variables. Non-stationary variables are transformed to growth rates. We allow for up to 8 lags.

A.2.2 Accounting for the Persistence of Temperature Shocks

As discussed in the paper, global temperature shocks lead to a persistent increase in temperature levels, that is in fact much more persistent than the shock itself: Figure 9 in the main text shows that global temperature shocks lead to an increase in average local temperature that tends to persist over our entire impulse horizon, albeit at a lower level than the initial shock of 1°C . An interesting question is then how global mean temparature responds. To this end, we estimate the response of global mean temperature to global temperature shocks using a simple local projection. The results are shown in Figure A.3. Global mean temperature increases persistently after the shock and the response turns out to be quite similar to the average response of local temperature. If at all, the response turns out to be slightly more pronounced, which could reflect the fact that global mean temperature also includes sea surface temperature.

[2]We would like to ideally include 10 lags (= our impulse horizon) but unfortunately in our baseline sample we do not have enough degrees of freedom to do so.

Figure A.3: The Effect of Global Temperature Shocks on Temperature Levels

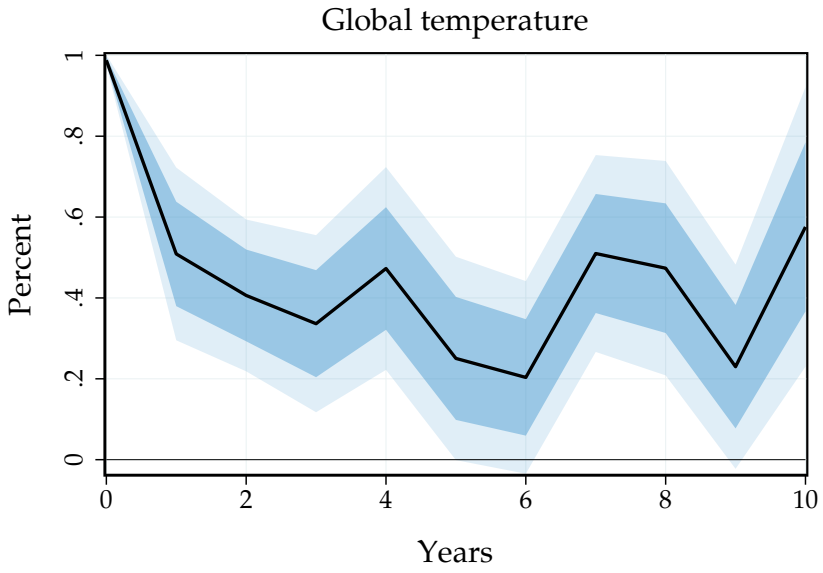

Notes: The figure shows the impulse responses of global mean temperature to a global temperature shock, estimated based on (2). The solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

Recall from Figure 9 in the main text, also local temperature shocks lead to persistent increases in local temperature. However, the increase turns out to be slightly less persistent than for global temperature shocks. To better compare the effects of global and local temperature shocks, we thus estimate the effects of local temperature shocks on real GDP, imposing the same persistence of the local temperature response as for global temperature shocks. We do so using the Sims (1986) approach.

The results are shown in Figure A.4. The effects of local temperature shocks are in this case slightly more pronounced than in our baseline, depicted in Figure 8. Importantly, however, the effects of the global temperature shocks are still by a magnitude larger that for local temperature shocks. Thus, the slight difference in persistence cannot account for the differential impacts of global and local temperatures shocks.

Figure A.4: The Effects of Local and Global Temperature Shocks

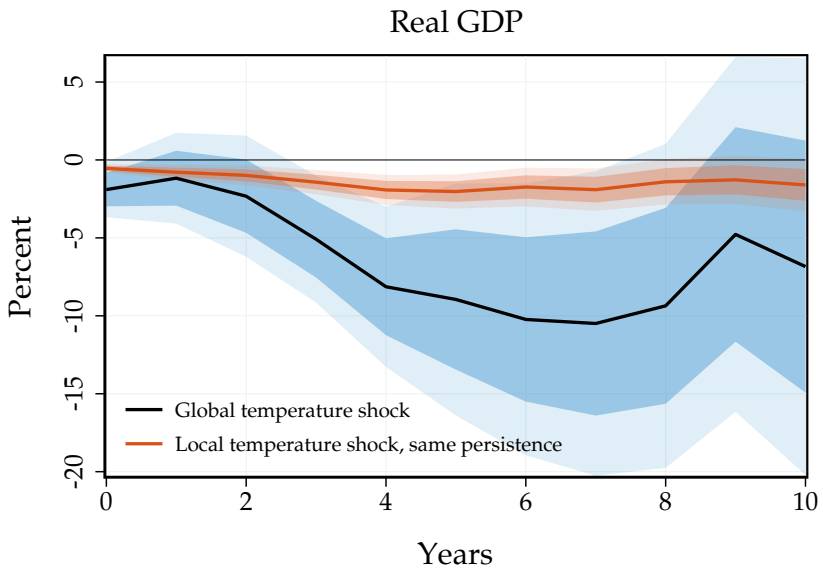

Notes: The figure shows the impulse responses of real GDP per capita to a local temperature shock, estimated in the panel using (3), against the effects of a global temperature shock. To make the shocks more comparable, we impose that the local temperature shock has the same effect on local temperature levels as global temperature shocks using the Sims (1986) method. The solid lines are the point estimates and the dark and light shaded areas and dashed and dotted lines are 68 and 90% confidence bands, respectively.

A.2.3 An LP-IV Approach

Our baseline specifications take the global temperature shock as given and do not take estimation uncertainty in the shock into account. To assess the potential role of estimation uncertainty in the shock, we alternatively consider a local-projection-instrumental variable approach. Specifically, we use the global temperature shock as an instrument for changes in global temperature. The specification then reads:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \theta_0^h \Delta T_t + \mathbf{x}_t' \boldsymbol{\beta} + \mathbf{x}_{i,t}' \boldsymbol{\gamma} + \varepsilon_{i,t+h},$$

where we instrument ΔT_t with T_t^{shock}. To account for the serial correlation in temperature changes, we also control for two lags of temperature changes, on top of our baseline controls.

Importantly, as discussed in Wooldridge (2002), generated instruments do not suffer

from the inference problem associated with generated regressors.

The results are shown in Figure A.5. The results are very similar to our baseline estimates, both in terms of point estimates and coverage of the confidence bands. This suggests that accounting for estimation uncertainty in the global temperature shock is not that consequential in our application.

Figure A.5: Local Projections-Instrumental Variable Approach

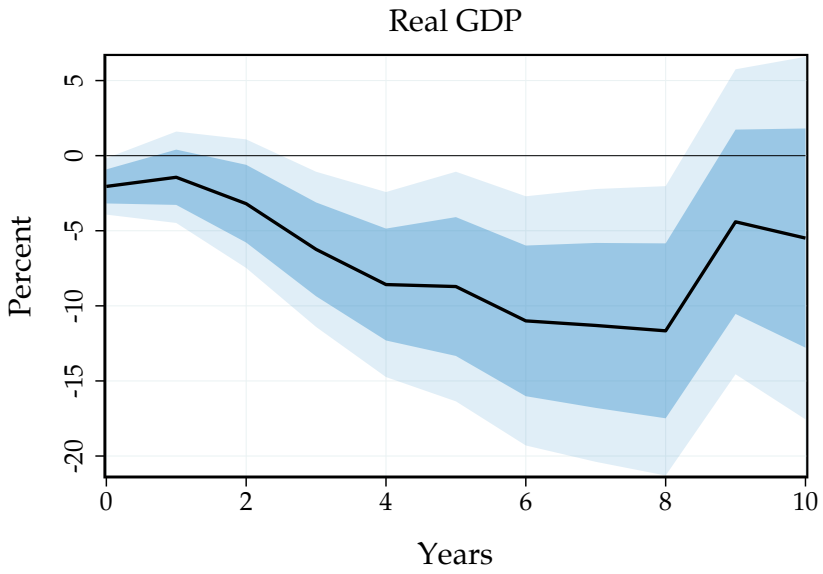

Notes: The figure shows the impulse responses of real GDP per capita to a global temperature shock, estimated based on the panel local projection-instrumental variable approach, instrumenting global temperature changes with the global temperature shock. The solid black line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

A.2.4 Time Fixed Effects and Correlated Temperature Shocks

In this appendix, we shed further light on the role of time fixed effects. In Figure A.6(a), we zoom in on the comparison between the impulse responses of local temperature shocks from the specification with time fixed effects to the baseline specification without time fixed effects. The responses from the local temperature shock model with time fixed effects are strikingly close to the baseline with global controls. Furthermore, the coverage of the confidence bands is also comparable. Overall, these results indicate that our controls successfully account for common shocks.

Figure A.6: The Role of Time Fixed Effects

(a) Global controls versus time fixed effects

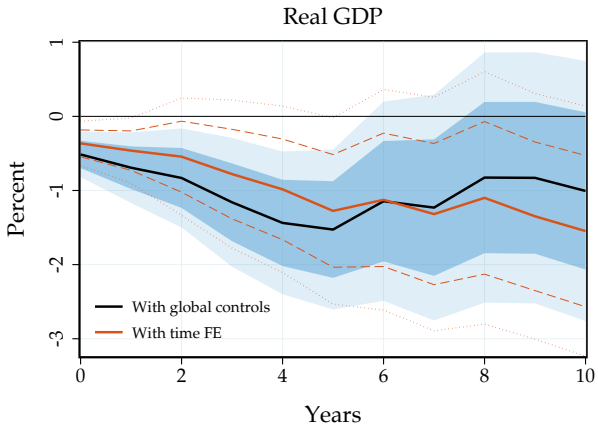

(b) Correlated temperature shocks

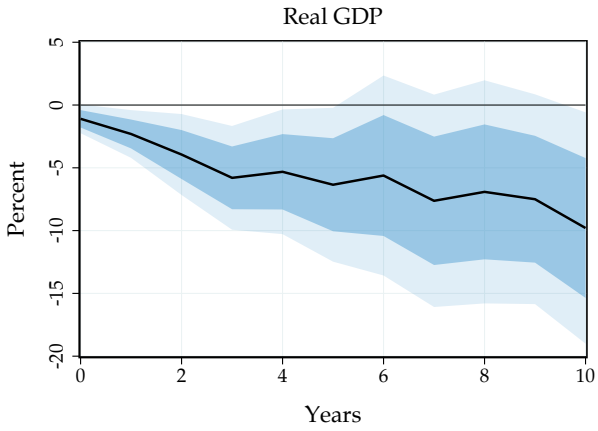

Notes: The figure shows the impulse responses of real GDP per capita. Panel (a) shows the responses of a local temperature shock and compares the specification with global controls (3) to the specification with time FE (4). The black line corresponds to the specification with global controls, the red line to the specification with time fixed effects. Panel (b) shows the impulse responses to correlated temperature shocks from a specification controlling for time fixed effects. In all subfigures, the solid lines are the point estimates and the dark and light shaded areas and dashed and dotted lines are 68 and 90% confidence bands, respectively.

To further mitigate concerns that other unobserved global factors may confound our results, we exploit regional variation in temperature. We construct country-level temperature shocks that also incorporate external temperature. For each country, we compute a shock that is a weighted average of its own temperature shock and all other temperature shocks in the world, weighted by country distance with closer countries getting a higher weight. In this way, we are able to obtain correlated temperature shocks that still vary by country. This allows us to control for time fixed effects—something we cannot do in the specification with global temperature shocks.

The results are shown in Figure A.6(b). Real GDP per capita falls substantially after such correlated temperature shocks, approaching -10% at its peak. Thus, the effects turn out to be again much larger than for local temperature shocks and the responses are overall quite similar to the estimated effects for global temperature shocks. Importantly, we reach this conclusion based on a specification with time fixed effects, flexibly controlling for any unobserved common shocks. We conclude that global temperature shocks lead to much larger economic effects than local temperature shocks.

A.3 Accounting for Reverse Causality

In this appendix, we describe how we account for reverse causality. We build on the deconvolution ideas in Sims (1986). We start by describing deconvolution without reverse causality. Then we show how to use deconvolution to handle reverse causality.

For expositional purposes, we omit additional controls and residuals from the notation to focus on the core idea. Adding controls and residuals would only add notation without changing any of the results. We estimate the relationship between a temperature shock $\widehat{T}_t$ at time t and output:

$$\Delta y_{t+h} = \theta_h \widehat{T}_t,$$

where T_t denotes the temperature shock, and $\Delta x_{t+h} \equiv x_{t+h} - x_{t-1}$. Temperature T_t has some internal persistence too:

$$T_{t+h} - T_{t-1} = \Gamma_h \widehat{T}_t.$$

A.3.1 No Reverse Causality

We are interested in the structural coefficients Θ_h that describe the effect of a one-time temperature change $\widehat{T}_t$ on GDP y_{t+h} periods ahead. The structural relationship is thus:

$$\Delta y_{t+h} = \Theta_h \widehat{T}_t.$$

However, in the data temperature does not revert back to 0 after period 1. Given this internal persistence in temperature, we thus observe the path of GDP after damages cumulate over time:

$$\Delta y_{t+h} = \sum_{k=0}^{h} \Theta_{h-k} \Delta T_{t+k} = \left(\sum_{k=0}^{h} \Theta_{h-k} \Gamma_k \right) \widehat{T}_t.$$

Hence, we measure θ, which is related to Θ through:

$$\theta_h = \sum_{k=0}^{h} \Theta_{h-k} \Gamma_k.$$

We can solve for Θ recursively:

- **Initial horizon.** We have $\theta_0 = \Theta_0 \Gamma_0$, and so $\Theta_0 = \frac{\theta_0}{\Gamma_0}$.
- **Subsequent horizons.** By recursion, we are given $(\Theta_k)_{k \leq h-1}$. We have:

$$\theta_h = \Theta_h \gamma_0 + \sum_{k=1}^{h} \Theta_{h-k} \Gamma_k = \Theta_h \gamma_0 + \sum_{k=0}^{h-1} \Theta_k \Gamma_{h-k} \quad \Longrightarrow \quad \Theta_h = \frac{\theta_h - \sum_{k=0}^{h-1} \Theta_k \Gamma_{h-k}}{\Gamma_0}.$$

A.3.2 Reverse Causality

We now introduce a feedback of GDP on temperature through emissions. Denote by $\overline{E}$ the baseline level of emissions. Consistently with standard production function structures, we posit, in log changes:

$$\Delta e_t = \eta \Delta y_t,$$

where η is the so-called "Okun elasticity" of emissions to GDP over the business cycle, and ranges from 0.5 to 1.

The relationship between the level of emissions and temperature going forward is known (see for instance Dietz et al., 2021). After an exegenous increase in emissions, but in the absence of an exogenous temperature shock, climate models imply:

$$\Delta T_{t+h} = \phi_h \Delta E_t = \eta \overline{E} \phi_h \Delta y_t \equiv \psi_h \Delta y_t,$$

where $E_t = \overline{E} \times e_t$.

Thus, we write the full temperature process as the outcome of both exogenous temperature shocks and past economic activity. Γ_h is the temperature response to an own shock absent any emissions feedback. Θ_h denotes economic damages absent any emissions feedback. We obtain the joint system:

$$\Delta y_{t+h} = \sum_{k=0}^{h} \Theta_{h-k} \Delta T_{t+k} \qquad \Delta T_{t+h} = \Gamma_h \widehat{T}_t + \sum_{k=0}^{h} \psi_{h-k} \Delta y_{t+k}$$

As before, we measure the local projection θ_h of output Δy_{t+h} on the initial temperature

shock $\widehat{T}_t$: $\Delta y_{t+h} = \theta_h \widehat{T}_t$. We also measure the local projection γ_h of temperature ΔT_{t+h} on the initial temperature shock $\widehat{T}_t$: $\Delta T_{t+h} = \gamma_h \widehat{T}_t$.

Thus, we obtain as system of equations with unknowns (Θ, Γ) given (θ, γ) and ψ:

$$\theta_h = \sum_{k=0}^{h} \Theta_{h-k} \gamma_k \widehat{T}_t \qquad\qquad \gamma_h = \Gamma_h \widehat{T}_t + \sum_{k=0}^{h} \psi_{h-k} \theta_k \widehat{T}_t.$$

As before, we solve for (Θ, Γ) recursively:

- **Initial horizon.** We have $\theta_0 = \Theta_0 \gamma_0$ and $\gamma_0 = \Gamma_0 + \psi_0 \theta_0$. Therefore, we obtain: $\Theta_0 = \frac{\theta_0}{\gamma_0}$ and $\Gamma_0 = \gamma_0 - \psi_0 \theta_0$.
- **Subsequent horizons.** By recursion, we are given $(\Theta_k, \Gamma_k)_{k \leq h-1}$ and $(\Theta_k, \Gamma_k)_{k \leq h}$. We have:

  $$\theta_h = \Theta_h \gamma_0 + \sum_{k=1}^{h} \Theta_{h-k} \gamma_k = \Theta_h \gamma_0 + \sum_{k=0}^{h-1} \Theta_k \gamma_{h-k} \qquad \gamma_h = \Gamma_h + \sum_{k=0}^{h} \psi_{h-k} \theta_k$$

  Therefore, we obtain:

  $$\Theta_h = \frac{\theta_h - \sum_{k=0}^{h-1} \Theta_k \gamma_{h-k}}{\gamma_0} \qquad\qquad \Gamma_h = \gamma_h - \sum_{k=0}^{h} \psi_{h-k} \theta_k.$$

  The deconvoluted GDP response is the same as without reverse causality. The only difference lies in the implied GDP response after a temperature shock that follows its own internal persistence.
- **Implied GDP response.** The implied GDP response after correcting for reverse causality is thus:

  $$\theta_h^{\text{true}} = \sum_{k=0}^{h} \Theta_{h-k} \Gamma_k = \sum_{k=0}^{h} \Theta_k \Gamma_{h-k}.$$

Practical implementation. In practice, we need a sequence ψ. We construct a central case, and some alternatives for robustness.

The central case uses the following parameters. We use $\eta = 1$: emissions move one-for-one with output. Average world emissions during our 1960-2019 sample are $\overline{E} =$

22.5 Gt/y.[3] The temperature response in Celsius to a 100 Gt pulse in Dietz et al. (2021) is well-approximated by:

$$100 \times \phi_h = a_{100} \times (e^{-b \times h} - e^{-c \times h}) + d_{100} \times (1 - e^{-f \times h})$$
$$a_{100} = 0.1878,\ b = 0.083,\ c = 0.2113,\ d_{100} = 0.1708,\ f = 0.2113.$$

Then: $\psi_h = \eta \times \overline{E} \times \phi_h$.

We consider combinations of the following other possible cases (including the baseline) that correspond to alternative Okun elasticities, and the lower and upper end of the temperature response in Dietz et al. (2021): $\eta' = 0.5$; $(a', d') = 0.5 \times (a, d)$, $(a', d') = 2 \times (a, d)$.

A.4 Additional Robustness Checks

In this appendix, we perform a number of additional sensitivity checks on the effect of global temperature shocks based on our panel local projections.

Figure A.7 shows the results. In Panels (a)-(b), we assess the sensitivity with respect to the GDP and temperature data used. Using real GDP per capita from the PWT or from the WDI produces very similar results. Similarly, using aggregated global mean temperature data from the Berkeley Earth dataset or off-the-shelf measures from NASA or NOAA produces virtually identical results.

In Panels (c)-(d), we study the sensitivity with respect to the number of lags included for real GDP and temperature shocks. When varying the lag order of the dependent variable, we keep the lag order of our temperature shock at the basline value and vice versa. Our results turn out to be very robust with respect to the lag order. Recall, in the main text, we show that our results even survive when we control up to 10 lags of real GDP.

In Panel (e), we perform a more extensive assessment of how constructing the temperature shocks affects the results. Using simple one-step ahead forecast errors, using the one-sided HP filter or the simple 2-year difference proposed in Hamilton (2018) produces qualitatively very similar results.

[3] See `https://ourworldindata.org/co2-emissions`.

Overall, these results further illustrate the robustness of our finding that global temperature shock lead to a sizeable, persistent and statistically significant fall in economic output that is by a magnitude larger than the estimates in the literature for local temperature shocks.

Figure A.7: Sensitivity of the Average Effect of Global Temperature Shocks

(a) GDP data

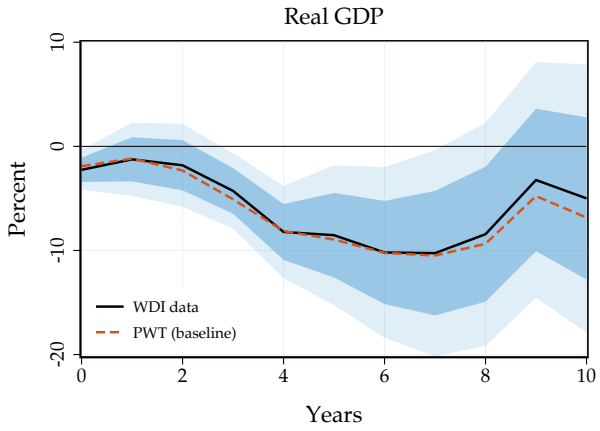

(b) Temperature data

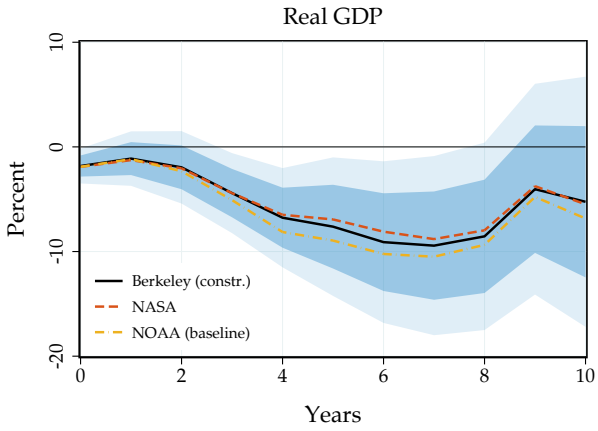

(c) Lag order dependent variable

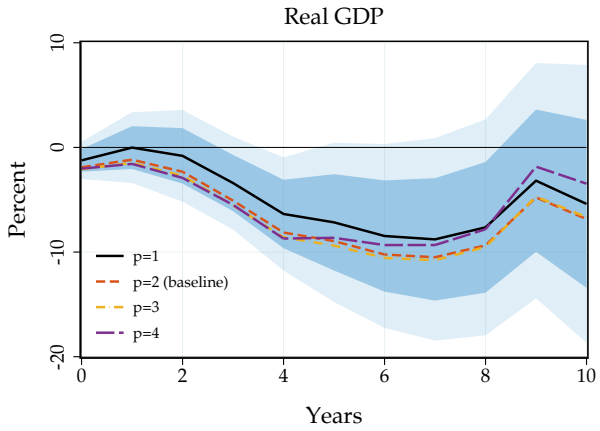

(d) Lag order temperature shock

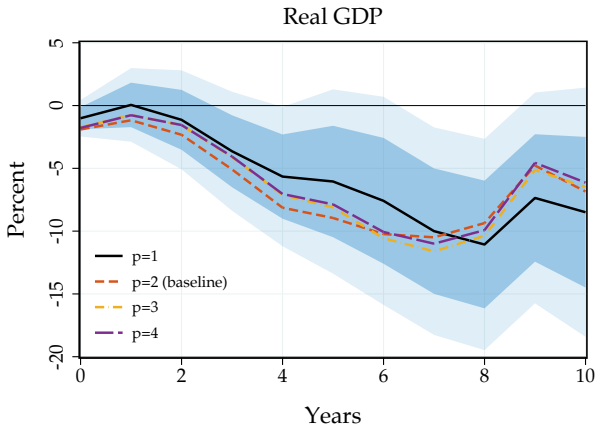

(e) Construction of temperature shock

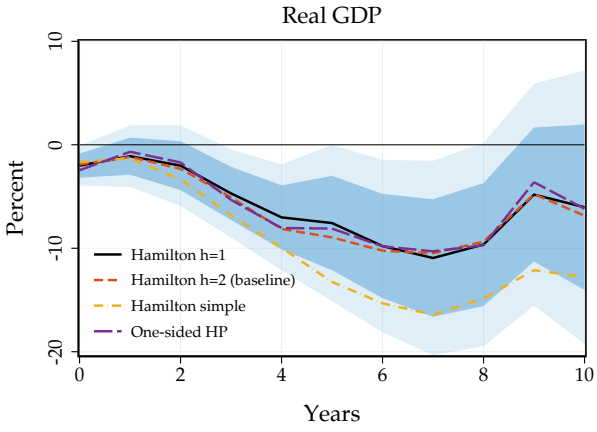

Notes: The figure assesses the sensitivity of the effects of global temperature shocks on real GDP per capita to a global temperature shock, with respect to data choices, the number of controls included, and the construction of the temperature shock. In all subfigures, the solid line is the point estimate and the dark and light shaded areas are 68 and 90% confidence bands, respectively.

B Model

Our solution to the neoclassical growth model is entirely standard and we present it for completeness.

B.1 Equilibrium

The resource constraint is:

$$\dot{K}_t = Z_t K_t^{\alpha} - C_t - \Delta_t K_t.$$

Firm behavior and market clearing implies $r_t + \Delta_t = \alpha Z_t K_t^{\alpha-1}$ and $w_t = (1-\alpha)K_t^a$. The Euler equation is:

$$\dot{C}_t = \gamma^{-1}(\alpha Z_t K_t^{\alpha-1} - \Delta_t - \rho)C_t.$$

In steady-state,

$$\begin{aligned} r = \alpha Z K^{\alpha-1} &= \rho + \Delta \implies K = \left(\frac{\alpha Z}{\rho+\Delta}\right)^{\frac{1}{1-\alpha}} \\ C &= ZK^{\alpha} - \Delta K. \end{aligned}$$

B.2 Linearization

We denote steady-state variables without time subscripts. We denote deviations from steady-state with hats. We linearize the resource constraint:

$$\begin{aligned} \frac{d\widehat{K}_t}{dt} &= (\alpha Z K^{\alpha-1} - \Delta)\widehat{K}_t - \widehat{C}_t + \widehat{Z}_t K^{\alpha} - \widehat{\Delta}_t K \\ &= \rho\widehat{K}_t - \widehat{C}_t + Y\widehat{z}_t - K\widehat{\Delta}_t. \end{aligned}$$

where we denoted $\widehat{z}_t = \widehat{Z}_t/Z$. Next, we linearize the Euler equation:

$$\begin{aligned}\frac{d\widehat{C}_t}{dt} &= \frac{C}{\gamma}\left(-\alpha(1-\alpha)ZK^{\alpha-2}\widehat{K}_t + \alpha K^{\alpha-1}\widehat{Z}_t - \widehat{\Delta}_t\right) \\ &= \frac{C}{\gamma}\left(-\frac{(1-\alpha)r}{K}\widehat{K}_t + r\widehat{z}_t - \widehat{\Delta}_t\right).\end{aligned}$$

We define:

$$X_t = \begin{pmatrix}\widehat{K}_t \\ \widehat{C}_t\end{pmatrix} \ , \ s_t = \begin{pmatrix}\widehat{z}_t \\ \widehat{\Delta}_t\end{pmatrix}.$$

We can summarize the linearized resource constraint and Euler equation as:

$$\dot{X}_t = AX_t + S_t,$$

where:

$$A = \begin{pmatrix}\rho & -1 \\ -\frac{(1-\alpha)rC}{\gamma K} & 0\end{pmatrix} \ , \ S_t = Bs_t \ , \ B = \begin{pmatrix}Y & -K \\ \frac{rC}{\gamma} & -\frac{C}{\gamma}\end{pmatrix}.$$

We have an initial condition $\widehat{K}_0$, and a terminal condition $\widehat{C}_t \to 0$. We now apply standard Blanchard-Kahn arguments. Let $A = M^{-1}DM$, with D diagonal. For determinacy we require that parameters are such that D has a positive eigenvalue in the top left position, and a negative eigenvalue in the bottom right position. We denote by $\mathcal{X}_t = MX_t$, so that

$$\dot{\mathcal{X}}_t = D\mathcal{X}_t + MS_t.$$

We then solve explicitly for $\mathcal{X}_t$:

$$\mathcal{X}_t = e^{tD}\left[\mathcal{X}_0 + \int_0^t e^{-sD}(MS_t)ds\right].$$

Hence, long-run stability requires the top entry of the bracket to be zero as time grows. That is:

$$0 = \mathcal{X}_{0,1} + \int_0^\infty e^{-sD_1}(MS_s)_1 ds.$$

Therefore,

$$M_{1\bullet}X_0 = -\int_0^\infty e^{-sD_1} M_{1\bullet} S_s ds.$$

We can thus solve for initial consumption:

$$\widehat{C}_0 = -\frac{1}{M_{12}}\left[M_{11}\widehat{K}_0 + \int_0^\infty e^{-sD_1} M_{1\bullet} S_s ds\right].$$

We denote $\varepsilon_K = -\frac{M_{11}}{M_{12}}$, $\varepsilon_S = -\frac{1}{M_{12}} M_{1\bullet}$ and $\varepsilon_{S,s} = e^{-sD_1}\varepsilon_S$. We can write more compactly:

$$\widehat{C}_0 = \varepsilon_K \widehat{K}_0 + \int_0^\infty \varepsilon_{S,s} S_s ds.$$

Of course, this condition must hold at all times:

$$\widehat{C}_t = \varepsilon_K \widehat{K}_t + \int_0^\infty \varepsilon_{S,s} S_{t+s} ds.$$

B.3 Model Inversion: Proof of Proposition 1

We substitute the solution for linearized consumption into the law of motion of capital:

$$\frac{d\widehat{K}_t}{dt} = (L_{11} - \varepsilon_K)\widehat{K}_t + S_{1t} - \int_0^\infty \varepsilon_{S,s} S_{t+s} ds.$$

Denote $\kappa = -(L_{11} - \varepsilon_K)$ and $\mathcal{S}_t = S_{1t} - \int_0^\infty \varepsilon_{S,s} S_{t+s} ds$ so that:

$$\frac{d\widehat{K}_t}{dt} = -\kappa \widehat{K}_t + \mathcal{S}_t.$$

Assuming we start in steady-state, we obtain:

$$\widehat{K}_t = e^{-\kappa t} \int_0^t e^{\kappa s} \mathcal{S}_s ds.$$

In percentage deviations:

$$\frac{\widehat{K}_t}{K} = \frac{e^{-\kappa t}}{K} \int_0^t e^{\kappa s} \mathcal{S}_s ds.$$

We can directly back out productivity shocks from the production function given movements in output and capital:

$$\frac{\widehat{Y}_t}{Y} = \widehat{z}_t + \alpha \frac{\widehat{K}_t}{K}.$$

We then use the capital accumulation equation to recover capital depreciation shocks. To do so, we express:

$$\begin{aligned}
\int_0^t e^{\kappa s} \mathcal{S}_s ds &= \int_0^t e^{\kappa s} \left(S_{1s} - \int_0^\infty \varepsilon_{S,r} S_{s+r} dr \right) ds \\
&= \int_0^t e^{\kappa s} S_{1s} ds - \iint_0^\infty \mathbb{1}[s \leq t] \varepsilon_{S,r} S_{s+r} e^{\kappa s} ds dr \\
&= \int_0^t e^{\kappa s} S_{1s} ds - \iint_0^\infty \mathbb{1}[s \leq t] \varepsilon_{S} S_{s+r} e^{\kappa s - D_1 r} ds dr.
\end{aligned}$$

Changing variables to $\tau = s + r$ over r, we obtain

$$\begin{aligned}
\int_0^t e^{\kappa s} \mathcal{S}_s ds &= \int_0^t e^{\kappa s} S_{1s} ds - \varepsilon_S \iint_0^\infty \mathbb{1}[s \leq t, s \leq \tau] S_\tau e^{\kappa s - D_1(\tau - s)} ds d\tau \\
&= \int_0^t e^{\kappa s} S_{1s} ds - \varepsilon_S \int_{\tau=0}^\infty e^{-D_1 \tau} S_\tau \int_{s=0}^{\min\{t,\tau\}} e^{(D_1+\kappa)s} ds d\tau \\
&\equiv \int_0^t e^{\kappa s} S_{1s} ds - \varepsilon_S \int_{\tau=0}^\infty J_{t,\tau} S_\tau d\tau,
\end{aligned}$$

where we defined:

$$J_{t,\tau} = e^{-D_1 \tau} \int_{s=0}^{\min\{t,\tau\}} e^{(D_1+\kappa)s} ds = e^{-D_1 \tau} \frac{e^{(D_1+\kappa)\min\{t,\tau\}} - 1}{D_1 + \kappa}.$$

Having estimated the productivity shocks, we can express:

$$S_t = \overline{S}_t + \widehat{\Delta}_t \overline{S}_\Delta \qquad \overline{S}_t \equiv \widehat{z}_t \begin{pmatrix} Y \\ \frac{rC}{\gamma} \end{pmatrix} \qquad S_\Delta \equiv \begin{pmatrix} -K \\ -\frac{C}{\gamma} \end{pmatrix}.$$

Then, we write

$$\int_0^t e^{\kappa s} \mathcal{S}_s ds = \int_0^t e^{\kappa s} \overline{S}_{1s} ds - \varepsilon_S \int_{\tau=0}^{\infty} J_{t,\tau} \overline{S}_\tau d\tau + + \overline{S}_{\Delta,1} \int_0^t e^{\kappa s} \widehat{\Delta}_s ds - \left(\varepsilon_S \overline{S}_\Delta\right) \int_0^\infty J_{t,s} \widehat{\Delta}_s ds.$$

Hence, we have obtained that:

$$\frac{\widehat{K}_t}{K} = \mathcal{K}_t(\widehat{z}) + \mathcal{J}_{t,\bullet} \widehat{\Delta}_\bullet,$$

where

$$\begin{aligned} \mathcal{K}_t(\widehat{z}) &= \frac{e^{-\kappa t}}{K} \left[\int_0^t e^{\kappa s} \overline{S}_{1s} ds - \varepsilon_S \int_0^\infty J_{t,s} \overline{S}_s ds \right] \\ \mathcal{J}_{t,s} &= \frac{e^{-\kappa t}}{K} \left[\overline{S}_{\Delta,1} \mathbb{1}[s \leq t] e^{\kappa s} ds - \left(\varepsilon_S \overline{S}_\Delta\right) J_{t,s} ds \right]. \end{aligned}$$

This concludes the proof of Proposition 1.

References Appendix

Berkeley Earth (2023). *Data Overview*. https://berkeleyearth.org/data/. (Visited on 08/09/2023).

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

Center for International Earth Science Information Network (CIESIN), Columbia University (2018). *Gridded Population of the World, Version 4.11 (GPWv4): Population Count*. NASA Socioeconomic Data and Applications Center (SEDAC). https://sedac.ciesin.columbia.edu/data/set/gpw-v4-population-count-rev11. (Visited on 09/12/2023).

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2012). "Temperature Shocks and Economic Growth: Evidence from the Last Half Century". *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans (2021). "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists* 8.5, pp. 895–921.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (Oct. 2015). "The Next Generation of the Penn World Table". *American Economic Review* 105.10, pp. 3150–3182.

Hamilton, James D. (2018). "Why you should never use the Hodrick-Prescott filter". *Review of Economics and Statistics* 100.5, pp. 831–843.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (Jan. 2017). "Macrofinancial History and the New Business Cycle Facts". *NBER Macroeconomics Annual* 31. Publisher: The University of Chicago Press, pp. 213–263.

Lange, Stefan, Matthias Mengel, Simon Treu, and Matthias Büchner (2023). *ISIMIP3a atmospheric climate input data (v1.2)*. ISIMIP Repository. https://doi.org/10.48364/ISIMIP.982724.2. (Visited on 12/04/2023).

Lenssen, Nathan J. L., Gavin A. Schmidt, James E. Hansen, Matthew J. Menne, Avraham Persin, Reto Ruedy, and Daniel Zyss (2019). "Improvements in the GISTEMP Uncertainty Model". *Journal of Geophysical Research: Atmospheres* 124.12, pp. 6307–6326.

Matsuura, Kenji and National Center for Atmospheric Research Staff (2023). *The Climate Data Guide: Global (land) precipitation and temperature: Willmott & Matsuura, University of Delaware*. https://climatedataguide.ucar.edu/climate-data/global-land-precipitation-and-temperature-willmott-matsuura-university-delaware. (Visited on 08/09/2023).

NASA Earth Observatory (Jan. 2020). *World of Change: Global Temperatures*. https://earthobservatory.nasa.gov/world-of-change/global-temperatures. (Visited on 08/08/2023).

NASA Goddard Institute for Space Studies (2023). *GISS Surface Temperature Analysis (GISTEMP v4)*. https://data.giss.nasa.gov/gistemp/. (Visited on 08/08/2023).

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (2023). "How Much Will Global Warming Cool Global Growth?" *NBER Working Paper Series*.

NOAA National Centers for Environmental Information (2023a). *Climate at a Glance: Global Time Series*. https://www.ncei.noaa.gov/access/monitoring/climate-at-a-glance/global/time-series. (Visited on 08/01/2023).

— (2023b). *Global Surface Temperature Anomalies: Mean Temperature Estimates*. https://www.ncei.noaa.gov/access/monitoring/global-temperature-anomalies/mean. (Visited on 08/01/2023).

Ramey, Valerie A. (2016). "Macroeconomic shocks and their propagation". *Handbook of Macroeconomics* 2, pp. 71–162.

Rohde, Robert A. and Zeke Hausfather (Dec. 2020). "The Berkeley Earth Land/Ocean Temperature Record". *Earth System Science Data* 12.4, pp. 3469–3479.

Sheffield, Justin, Gopi Goteti, and Eric F. Wood (July 2006). "Development of a 50-Year High-Resolution Global Dataset of Meteorological Forcings for Land Surface Modeling". *Journal of Climate* 19.13. Publisher: American Meteorological Society Section: Journal of Climate, pp. 3088–3111.

Sims, Christopher A. (1986). "Are forecasting models usable for policy analysis?" *Quarterly Review* 10.Win, pp. 2–16.

Wooldridge, Jeffrey M. (2002). *Econometric analysis of cross section and panel data*. MIT press.