

The Macroeconomic Impact of Climate Change: Global vs. Local Temperature*

Adrien Bilal[†]
Stanford University

Diego R. Känzig[‡]
Northwestern University

November, 2024

Abstract

This paper estimates that the macroeconomic damages from climate change are six times larger than previously thought. Exploiting natural global temperature variability, we find that 1°C warming reduces world GDP by 12%. Global temperature correlates strongly with extreme climatic events, unlike country-level temperature used in previous work, explaining our larger estimate. We use this evidence to estimate damage functions in a neoclassical growth model. Business-as-usual warming implies a 25% present welfare loss and a Social Cost of Carbon of \$1,367 per ton. These impacts suggest that unilateral decarbonization policy is cost-effective for large countries such as the United States.

JEL classification: E01, E23, F18, O44, Q54, Q56

Keywords: Climate change, macroeconomic damages, time-series variation, local and global shocks

*We thank Marios Angeletos, Marshall Burke, Gabriel Chodorow-Reich, Simon Dietz, Stephane Hallegatte, Jim Hamilton, Xavier Jaravel, Ben Jones, Eben Lazarus, Pooya Molavi, Ishan Nath, Ben Olken, Esteban Rossi-Hansberg, Toan Phan, Jón Steinsson, Jeffrey Shrader, Jim Stock, Chris Wolf, and numerous participants at conferences and seminars for helpful comments and suggestions. We thank Ramya Raghavan, Lilian Hartmann and Cathy Wang for outstanding research assistance. Adrien Bilal gratefully acknowledges support from the Chae Family Economics Research Fund at Harvard University.

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

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

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: they imply that a 1°C rise in temperature reduces world output at most by 1-2%. Under any conventional discounting, such effects seem hardly catastrophic. Why are perceptions of climate change misaligned with empirical estimates? Do existing estimates not account for the full impact of climate change, or are its costs truly small?

In this paper, we demonstrate that the macroeconomic impacts of climate change are six times larger than previously documented. We rely on a time-series local projection approach to estimate the impact of global temperature 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 damaging 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. We then use our reduced-form results to estimate structural damage functions in a simple neoclassical growth model. Climate change leads to a present-value welfare loss of 25% and a Social Cost of Carbon (SCC) of \$1,367 per ton of carbon dioxide.

In the first part of the paper, we develop our time-series approach. We assemble a new climate-economy dataset spanning the last 120 years from sources that are regularly updated 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. We define physical granular reanalysis measures of extreme temperature, droughts, 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 onward, and from the Jordà-Schularick-Taylor Macrohistory database for select countries since 1900.

Identification of the impact of temperature on GDP is complicated by their jointly trending behavior. We thus construct global and local (country-level) *temperature shocks*: innovations to the temperature process that are orthogonal to their long-run trends and persist for 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 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 onward. A 1°C global temperature shock leads to a gradual decline in world GDP that peaks at 12% after 6 years, is statistically significant at the 5% level in years 3 to 7, and does not fully mean-revert even after 10 years. Importantly, the temperature *shocks* have a persistent effect on the temperature *level*, that remains above 0.5°C for multiple years after the shock. Hence, the 12% impact partly reflects the accumulated effects of persistently elevated global temperature itself. These results remain unchanged for alternative de-trending approaches, such as one-step ahead forecast errors or a one-sided Hodrick-Prescott filter.

Four identification concerns may challenge the causal interpretation of our headline results. We address each of them in a series of robustness exercises. First, we account for omitted variable bias: global temperature shocks may coincide with the global economic and financial cycle. We control for rich measures of world economic performance: indicators for global economic recessions and global macro-financial variables (past world real GDP, commodity prices and interest rates). Our results remain unaffected by the specific set of controls and are not driven by any particularly influential years, indicating that temperature shocks are largely unrelated to economic shocks.

Second, we account for reverse causality: as output declines after a temperature shock, energy consumption and greenhouse gas emissions fall, lowering temperatures and increasing output going forward. Qualitatively, reverse causality thus leads us to underestimate the true impact of a global temperature shock. Quantitatively, it is likely negligible because short-run fluctuations in emissions imply small temperature variations. We confirm these arguments by explicitly adjusting for the impact of past greenhouse gas and aerosol emissions with a climate model and find virtually identical results.

Third, we verify that our estimates are likely externally valid and are stable across time periods and causes of temperature variation. We find remarkably similar estimates in three time periods (1900-2019, 1960-2019—our main sample—and 1985-2019) as well as when we exclude El Niño and volcanic eruptions from our identifying variation.

Fourth, we account for regional omitted variable bias: global temperature may be particularly driven by some countries while they also experience unrelated GDP growth. We

obtain virtually identical results when we project country-level GDP—rather than global GDP—on global temperature, controlling for country fixed effects and region-specific time trends. Collectively, our robustness exercises suggest that our specification captures the causal effect of global temperature on economic activity.

Our estimated effect of temperature shocks on world GDP stands in stark contrast to existing estimates of the economic impact of temperature. Nordhaus (1992), Dell et al. (2012), Burke et al. (2015) and Nath et al. (2022) find that a 1°C temperature increase reduces GDP by at most 1-2% in the medium run. Why do we find effects that are six times larger?

We focus on a different source of variation: changes in *global* mean temperature capture the comprehensive impact of climate change. By contrast, previous work exploits changes in *country-level*, *local* temperature. When we estimate the impact of *local* temperature on country-level GDP with the same empirical specification, we find similarly small and imprecise effects to previous studies: 1% per 1°C, not significant at the 5% level. Econometrically, panel analyses using *local* temperature net out common impacts of *global* temperature through time fixed effects. Instead, we focus on these common impacts.

Why, then, does global temperature depress economic activity so much more than local temperature? We show that global temperature is fundamentally different from local temperature. Global temperature shocks predict a large and persistent rise in the frequency of four extreme climatic events that cause economic damage: extreme temperature, droughts, extreme wind, and extreme precipitation. By contrast, local temperature shocks predict a much weaker rise in these extremes.

These conclusions are consistent with the geoscience literature: droughts, 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 local temperature realizations (Seneviratne et al., 2016; Wartenburger et al., 2017; Seneviratne et al., 2021; Domeisen et al., 2023). In line with this view, we find that ocean temperature, rather than land temperature, is responsible for the vast majority of the effects of global temperature on economic activity.

Quantitatively, including these four extreme events accounts for close to two thirds of our estimated global temperature impact. We reach this conclusion by estimating the impact of extreme events on GDP, which we combine with the dynamic correlation between global temperature shocks and extreme events to construct a counterfactual impact

of global temperature on GDP. Of course, our aggregation exercise is unlikely to account for the full effect of global temperature on GDP: we would need to specify and measure the universe of channels whereby global temperature affects the economy. Using global temperature directly bypasses this challenge.

Another possible explanation for the differential impacts of global and local temperature shocks is that general equilibrium linkages together with spatially correlated local temperature lead panel analyses to underestimate the true impact of local temperature. However, we find that general equilibrium linkages account for at most one fifth of our global temperature estimates. To reach this conclusion, we construct an external temperature shock by aggregating the local temperature of all trading partners of a given country, weighted by trade shares. External local temperature turns out to have a similarly small impact on economic activity to the direct effect of local temperature: given moderate levels of trade openness throughout the world, indirect effects cannot be substantially larger than direct effects of local temperature.

How and where do the worldwide GDP impacts of global temperature materialize? We document a significant fall in capital, investment and productivity after a global temperature shock. Warm and low-income countries appear to be more severely affected than cold and high-income countries, although these comparisons are somewhat noisy. Overall however, global temperature has more uniformly detrimental effects than local temperature.

In the second part of the paper, we develop a simple neoclassical growth model to translate our reduced-form estimates into welfare effects, similarly to Nordhaus (1992). We introduce productivity damages from temperature. Critically, we use our novel reduced-form effects to estimate the associated structural damage function.

We estimate productivity shocks that correspond to a global temperature shock by matching the estimated impulse response function of output. This mapping has a closed-form expression that guarantees identification. In doing so, we account for the internal persistence of global mean temperature in response to temperature shocks. We remain conservative and impose persistent level effects rather than growth effects. We find that a one-time transitory 1°C rise in global mean temperature leads to a 4% peak productivity decline. The estimated model also matches the untargeted impulse response of capital to a temperature shock.

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% rate of time preference. Climate change implies a precipitous decline in output of 46% by 2100. Capital shrinks by 37% and consumption drops by 37%, leading to a 25% welfare loss in permanent consumption equivalent in 2024. These magnitudes are comparable to the economic damage caused by the 1929 Great Depression, but experienced *permanently*.

If the economic effects of global temperature are so large, why were they not noticed after nearly 1°C of global warming since 1960? Because climate change occurs in small increments, its effects are hidden behind background economic variability. We show that since 1960, climate change caused a gradual reduction in the annual world growth rate that reaches one third of baseline by 2019. Because climate change is also persistent, its effects keep accumulating over time. Ultimately, world GDP per capita would be 19% higher today had no warming occurred between 1960 and 2019.

The estimated model lets us characterize the SCC using the global temperature response to a CO₂ pulse from Dietz et al. (2021) and Folini et al. (2024). We remain conservative and use the lower end of the range of temperature responses from Dietz et al. (2021), which is also consistent with historical emissions and warming data.

We obtain a SCC of \$1,367 per ton. This value is seven times larger than the high end of existing estimates (\$185 per ton, Rennert et al., 2022). The 95% confidence interval for the SCC ranges from \$469 per ton to \$2,264 per ton. While this range is non-trivial, even its lower bound is multiple times larger than conventional SCC values. Our focus on global temperature shocks accounts for this substantial difference. When we re-estimate our model based on the impact of local temperature shocks as in previous research, the welfare cost of climate change is 3% and the SCC is \$178 per ton, and neither of these values is statistically significant at the 5% level.

How sensitive are these results to specification choices? Any plausible discount rate and 2100 temperature results in welfare losses in excess of 15% and a SCC above \$500 per ton. Discount rates below 1% imply a SCC exceeding \$3,000 per ton. Scenarios with 2100 warming 5°C above pre-industrial levels lead to welfare losses larger than 40%. The mean climate sensitivity from Dietz et al. (2021) implies a SCC above \$1,700 per ton.

We conclude by delineating the consequences of our results for decarbonization policy. Decarbonization interventions cost \$80 per ton of CO₂ abated on average (Bistline et al., 2023). A conventional SCC value based on local temperature of \$178 per ton 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 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. Under conventional estimates, the DCC of the United States is \$36 per ton, making unilateral emissions reduction prohibitively expensive. Under our new estimates, the DCC of the United States becomes \$273 per ton and thus exceeds policy costs. In that case, unilateral decarbonization policy is cost-effective for the United States.

Related literature. Our paper contributes to an influential literature measuring economic damages from climate change based on historical temperature variation, surveyed in Burke et al. (2023) and Moore et al. (2024). The conventional approach estimates the effect of small, short-run *local* temperature shocks over time within a spatial area on economic outcomes in a panel structure to achieve credible identification (Dell et al., 2012; Dell et al., 2014; Burke et al., 2015; Newell et al., 2021; Kahn et al., 2021; Nath et al., 2022). Consistently across all these studies, medium-term effects range from 0.5% to 2% of GDP, on average, and rely exclusively on climatic variation within countries or smaller geographic units. Our paper takes a fundamentally different approach: we directly exploit aggregate time-series variation in *global* mean temperature instead of relying on within-country climatic variation that nets out common effects of global temperature.

Local temperature estimates can lead to large long-run economic effects if one assumes that medium-run growth impacts persist forever (Dell et al., 2012; Burke et al., 2015; Burke et al., 2023, “growth effects”). Our treatment of temperature trends and persistence builds on Nath et al. (2022), who clarify how to translate short-run local weather effects to long-run counterfactuals by distinguishing between the polar cases of “level effects” and “growth effects”. We purposefully do not impose permanent growth effects and instead use the degree of persistence implied by the data. If we did impose growth effects, global temperature impacts would be even larger.

Perhaps surprisingly, few studies have explored time-series variation in temperature. Bansal and Ochoa (2011) find that a 1°C global temperature increase reduces GDP by 1% *contemporaneously*. We show that the *persistence* of the GDP response is crucial: the peak effect is an order of magnitude larger than this contemporaneous impact. Berg et al. (2023) analyze the effects of global and idiosyncratic temperature shocks on GDP dispersion across countries. We directly estimate the *aggregate impact* of global temperature,

which is more precisely estimated than individual country-level responses. In contemporary work, Neal (2023) also emphasizes the role of global temperature, suggesting that correlated local temperature and spillover effects across countries may lead to underestimate the effect of local temperature (see also Zappalà, 2023). We show that spillovers fall short of rationalizing the gap between global and local temperature impacts. Instead, we highlight the distinct relationship between global temperature and extreme climatic events. Relative to these papers, we also use our macroeconomic estimates in a structural model to evaluate welfare and the SCC.

As such, our paper relates to the literature studying the economic impact of climatic phenomena such as storms, heatwaves or El Niño (Barro, 2006; Deschênes and Greenstone, 2011; Hsiang et al., 2011; Deryugina, 2013; Hsiang and Jina, 2014; Bilal and Rossi-Hansberg, 2023; Phan and Schwartzman, 2024; Tran and Wilson, 2023; Callahan and Mankin, 2023; Dingel et al., 2023). Kotz et al. (2024) show that local extreme temperature and precipitation increase damages beyond local average temperature. We evaluate the impact of global temperature directly and provide new evidence on the relationship between global temperature and a wide range of extreme climatic events.

Our paper also connects to the literature using Integrated Assessment Models surveyed in Nordhaus (2013). We take a “top-down” approach and directly estimate and match the macroeconomic impact of global temperature. Our analysis suggests that Integrated Assessment 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 (Nordhaus, 2013; Stern et al., 2022).

More recently, “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) match micro-level estimates and aggregate using the model. Our “top-down” approach is more holistic in that we do not need to specify and estimate all channels and general equilibrium effects mapping global temperature to damages, but remains necessarily limited in assessing distributional and adaptation effects.

In fact, both our global temperature approach and the conventional local temperature approach rely on moderate short-run temperature variation for identification. We thus remain as close as possible to the local temperature approach in our ability to capture

adaption. To the extent that adaptation may be more pronounced in response to large long-run temperature changes, our results represent an upper bound on realized economic damages from slowly unfolding climate change. Although assessing the role of adaptation is beyond the scope of this paper, existing evidence suggests a limited role for adaptation even for the United States (Burke and Emerick, 2016; Moscona and Sastry, 2022; Bilal and Rossi-Hansberg, 2023), in line with the stability of our estimates across time periods. Should there be an unprecedented uptake in adaptation in the future, our numbers would still represent society’s willingness to pay for such investments.

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 compares the effects of global and local temperature. 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

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 also construct area-weighted measures of global temperature, which includes land and ocean surface air temperature. Reassuringly, our measure correlates virtually perfectly with global temperature data from the National Oceanic and Atmospheric Administration (NOAA).

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 20th and early 21st 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. We compute exposure indices to extreme weather by recording the fraction of days within a country and year that experience a weather re-


alization above or below a fixed percentile of the daily weather distribution in 1950-1980. See Appendix A.1 for details.

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

Figure 1 displays the evolution of global average temperature and world real GDP per capita since the post-World War II era in our dataset. In the late-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 1960



Notes: Evolution of global average temperature, computed based on Berkeley Earth temperature anomaly data and the corresponding climatology, 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 greenhouse gas emissions which eventually translate 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 possibly 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).


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 of the data in a time-series sense: 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 suggest 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:

$$\widehat{T_{t+h}^{\text{shock}}} = T_{t+h} - (\widehat{\alpha} + \widehat{\beta}_1 T_t + \dots + \widehat{\beta}_p T_{t-p+1}), \quad (1)$$

where $\widehat{\beta}_i$ denotes the coefficient estimates of the regression of temperature on its lag i and $\widehat{\alpha}$ is the estimated intercept. This exercise amounts to isolating shocks that persist typically for h periods. Selecting the horizon h is of course an important 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. As we show in Section 2.4 below and in Appendix A.14, varying these values leaves our results essentially unchanged. In particular, Appendix A.14.1 reproduces all our main

Figure 2: Global Temperature Shock



Notes: Global temperature shocks, computed as in Hamilton (2018) with $(h = 2, p = 2)$, since 1960.

analyses under a one-step ahead forecast error $h = 1$ as commonly used in the literature and finds virtually identical results.

Figure 2 shows 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. Figure A.2 in Appendix A.2 indicates that the series is also weakly autocorrelated, because we allow for relatively persistent deviations from the long-run temperature trend. 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 (Nath et al., 2022).

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. We evaluate directly these medium-run effects of temperature without extrapolating 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 as in Jordà (2005). This approach involves estimating the following series of re-

gressions, one for each horizon $h = 0, \dots, 10$:

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h T_t^{\text{shock}} + \mathbf{x}_t' \boldsymbol{\beta}_h + \varepsilon_{t+h}, \quad (2)$$

where y_t is the outcome variable of interest, T_t^{shock} is the temperature shock and θ_h is the dynamic causal effect of interest at horizon h . We refer to the latter as the impulse response function. $\mathbf{x}_t$ is a vector of controls and ε_t is a potentially serially correlated error term. Our main outcome variable of interest is (log) world real GDP per capita. Because we are using the cumulative growth rate as the dependent variable, we are estimating a possibly persistent level effect. The estimation sample is 1960-2019.¹

We use local projections in our main analysis because they tend to be robust at longer horizons (Montiel Olea et al., 2024). Compared to Vector Autoregressions (VARs) or distributed lag models, local projections directly estimate the effects of interest rather than extrapolating from the first few autocovariances and allow for more flexible controls. Yet, we obtain similar results under alternative estimation models in Appendix A.5.

To account for the serial correlation in GDP growth and temperature shocks, we include 2 lags of real GDP growth per capita and of the global temperature shock. To control for the global business cycle more comprehensively, we also include 2 lags of global oil price changes and the U.S. treasury yield. Finally, we flexibly control for large economic shocks, such as the large oil shocks in the 1970s or the Great Recession, using a set of dummy variables.² We compute the confidence bands using the lag-augmentation approach (Montiel Olea and Plagborg-Møller, 2021).³


Figure 3 shows the impulse responses of global temperature and 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 90 and 95% confidence bands, respectively. Panel (a) indicates that global temperature increases by 1°C on impact. The effect of a global temperature shock on global temperature turns out to be highly persistent: after 10 years global tem-

¹Leveraging that temperature data is available for a longer period than GDP data, we estimate the temperature shock based on this longer sample (1950-2022) to mitigate the influence from observations at the beginning and the end of the sample.

²Our 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, 2007-2009, and 2011-2012. To allow for potential persistent effects of recessions, we also include 2 lags of the global recession indicator variable.

³As in Nath et al. (2022), we do not account for estimation uncertainty in the global temperature shock in our baseline specification. However, we alternatively conduct inference using bootstrapping techniques, and taking estimation uncertainty into account yields very similar inference. See Appendix A.3 for more details.

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



Notes: Impulse responses of global mean temperature and world real GDP per capita to a global temperature shock, estimated based on (2) on the period 1960-2019. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

perature is still elevated by about 0.5°C.

The persistent rise in global temperature leads to large economic effects. Panel (b) shows that, on impact, world GDP falls by about 2.5%. However, the effect builds up over time. After six years, world GDP falls by 12%, with effects that persist up to eight years out. Our estimate is of similar magnitude to growth impacts that typically occur after severe financial crises (Cerra and Saxena, 2008; Reinhart and Rogoff, 2009).

The gradual decline in world GDP reflects not only the direct impact of the initial temperature shock, but also the subsequent effects of persistently elevated temperature that accumulate over time. Figure A.5 in Appendix A.4 shows that these accumulated effects of persistently elevated temperature account for a substantial part of the peak effect and of its timing. We construct a counterfactual path of output that would correspond to a one-time fully transitory global temperature change using the method in Sims (1986). The peak impact is then 5% instead, and occurs already four years after the shock.

Of course, a 1°C temperature shock is a large shock that does not occur directly in our historical sample: we observe smaller shocks in practice. Our estimate for a 1°C shock scales up the linear effect of these smaller shocks. In effect, we abstract from potential non-linearities. Among the relatively small shocks we observe, we do not find much evidence for non-linearities. Figure A.9 in Appendix A.8 reports comparable impacts of small, larger, or only positive shocks. Of course, these results do not imply that non-linearities do not matter for larger shocks that have not yet materialized. However, in the

presence of potential future tipping points, one may expect larger effects than predicted by our linear model.

2.4 Robustness

The time-series nature of our identifying variation requires care in interpreting these conclusions. We now demonstrate that our main estimate is robust to accounting for various identification concerns.


Omitted variable bias. Global temperature innovations may happen to be correlated with the global economic cycle over time. For instance, if a severe El Niño event increases global temperature at the same time that a global recession occurs for unrelated reasons, we may mistakenly attribute adverse economic impacts to climatic variations.

To account for this possibility, we already include a rich set controls of the world economic performance in our main specification in equation (2). In Figure 4(a), we show that our results hold regardless of the particular set of macroeconomic controls. We consider two specifications based on more parsimonious sets of controls: a specification that does not include oil prices and the treasury yield and a specification that also excludes the global recession dates. Reassuringly, the point estimates are very similar across specifications, suggesting that global temperature shocks and economic shocks are largely unrelated in our sample. However, the additional controls help reduce sampling uncertainty and lead to more precise estimates.

We confirm that spuriously correlated economic shocks are unlikely to drive our results by examining how each year in the sample affects our estimates. For all years t , Figure 4(b) plots the change in GDP 5 years later at $t + 5$ against the temperature shocks at time t after residualizing both from our set of controls. The negative relationship turns out to be a robust one and is not driven by a specific set of outliers. Figure A.8 in Appendix A.6 displays a systematic jackknife exercise in which we censor one year at a time and find that our estimates are not driven by specific years. Overall, these results indicate that our estimates are unlikely to be driven by economic shocks spuriously correlated to temperature shocks.

Reverse causality. Changes in economic activity may affect short-run variations in temperature: a decline in economic activity lowers emissions and temperature, hence increases output going forward, and potentially biases our estimates.

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



Notes: Impulse responses of world GDP per capita to a global temperature shock, estimated from (2) over the period 1960-2019. Panel (a): sensitivity with respect to controls included: baseline (two lags of temperature shocks, world real GDP growth, oil prices, one-year US treasury yield and global recession dates); specification without oil prices and treasury yield; specification which only controls for two lags of the temperature shock and GDP growth. Panel (b): scatter plot of temperature shocks against the cumulative change in real GDP 5 years out, both after residualizing our set of controls. Panel (c): GDP response after adjusting for reverse causality. Panel (d): sensitivity with respect to the construction of the temperature shock: baseline with $h = 2$; one-step ahead forecast error $h = 1$; one-sided HP filter; one-step LP estimation with 4 lags of global temperature changes. Lines: point estimate. Dark and light shaded areas: baseline 90 and 95% confidence bands.

There are two reasons why reverse causality due to greenhouse gases is unlikely to substantially affect our interpretation. First, such reverse causality concerns typically lead us to underestimate the effect of temperature on economic output. As temperature rises and economic activity initially declines, the resulting fall in greenhouse gas emissions implies lower future temperatures and thus higher future output. Thus, true damages

would be even larger than our estimates.

Second, annual fluctuations in emissions imply negligible temperature variations relative to the typical temperature shocks that we exploit. For instance, 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.

Aerosol emissions can also lead to reverse causality, for instance due to sulfur dioxide (SO2). Aerosols have the opposite effect of greenhouse gases and reduce global temperature by reflecting incoming sunlight. Aerosols are shorter-lived than greenhouse gases in the atmosphere, which may amplify or dampen reverse causality concerns relative to greenhouse gases depending on the horizon of interest.

Two exercises verify that reverse causality is unlikely to affect our results. First, we test whether our temperature shocks are forecastable by past macro-financial variables with a series of Granger-causality tests in Table A.2, Appendix A.2. We find no evidence that global temperature shocks are forecastable, consistently with the substantial lag and small sensitivity between emissions and temperature changes.

Second, we explicitly account for the feedback between output and temperature through emissions. We consider the two most important greenhouse gases: carbon dioxide (CO2) and methane (CH4). We also include the main source of aerosol emissions: sulfur dioxide (SO2). We use standard estimates of the emissions-to-GDP elasticity and leading estimates of the dynamic sensitivity of temperature to an emissions impulse to construct our adjustment. We provide more details in Appendix A.7. Figure 4(c) confirms that explicitly adjusting for reverse causality has no meaningful effect on our results.

External validity and temperature variability. Different ways of constructing our temperature shocks or excluding specific sources of global temperature variation may lead to different results. We address this concern with a series of exercises.

We show that our results hold across a variety of definitions of temperature shocks. In our baseline specification, we measure temperature shocks using the Hamilton (2018) filter with a horizon $h = 2$. In Figure 4(d), we show that constructing temperature shocks

as one-step ahead forecast errors $h = 1$ following previous work (see e.g. Bansal and Ochoa, 2011; Nath et al., 2022) or using a one-sided HP filter produces similar results. In addition, Appendix A.14.1 reproduces all our main analyses under a one-step ahead forecast error $h = 1$ and finds virtually identical results.

We also show that our results are virtually unchanged when we directly estimate the effects of temperature on world GDP without highlighting the identifying variation through global temperature shocks. In that case, instead of estimating temperature shocks in a first step by projecting temperature on its lags, and then projecting world real GDP on temperature shocks in a second step, we directly project world real GDP on temperature with enough lags of temperature and GDP. Both approaches are numerically equivalent when we construct the shocks as one-step ahead forecast errors with the same controls (see Appendix A.5 for more details).

We include both positive and negative temperature shocks in our main analysis to maximize precision. To assess the effects of global warming, focusing on positive shocks may be more directly informative. Figure A.9 in Appendix A.8 shows that positive shocks alone have a similar—if anything, slightly larger—impact to our baseline analysis.

In addition, our results do not depend on specific sources of global temperature variation. We re-evaluate our results after netting out temperature variation generated by El Niño by controlling for an ENSO index in our main specification. The results are shown in Figure A.3 in Appendix A.2. The responses are close to our baseline estimates. Similarly, controlling for volcanic eruptions also yields virtually unchanged results. These exercises indicate that our main results capture a broad effect of global temperature on economic activity that is not specific to particular sources of temperature variation.

Together, these robustness exercises corroborate our interpretation that global temperature shocks are driven by various 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 three 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 with existing work that has focused on country-level temperature. Our third goal in Section 3.4 is to unpack the margins through which GDP declines and the heterogeneity in country-level responses.

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 in 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 $h = 0, \dots, 10$:

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

where $y_{i,t}$ is the outcome variable of interest for 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, and in addition control for two lags of country-level GDP growth and region-specific time trends. We assess the sensitivity with respect to these controls in further sensitivity checks below. Our main outcome variable of interest is country-level log real GDP per capita. Our sample is an unbalanced panel spanning 1960-2019.

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 rely on Driscoll and Kraay (1998) standard errors that are robust to cross-sectional and serial dependence.⁴

By design, our specification is close to the specifications commonly used in the panel

⁴Our results are robust to using two-way clustered standard errors 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 4.3, we rely on the distribution estimated using a Wild bootstrap. See Appendix A.3 for more details.

literature on the economic effects of local temperature shocks (e.g. Dell et al., 2012; Burke et al., 2015; Nath et al., 2022). 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 common when shocks are also country-specific. Instead, we include the same global control variables as in our time-series specification (2).


Figure 5(a) 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 substantial fall in real GDP per capita that exceeds 10% at peak, is statistically significant at the 5% level in years 3 to 8, and persists up to 10 years out. This effect is close to our time series analysis, indicating that our results are robust to accounting for unobserved fixed country characteristics.

The increased statistical power in the panel allows us conduct a number of additional sensitivity checks. Panel (b) evaluates whether our results depend on the sample period. We obtain similar results on a sample that starts in 1985 after the large oil shocks of the 1970s or on a sample ending before the 2008 Great Recession. We also consider a much longer sample starting in 1900. For this analysis, we rely on the 18 advanced economies in the Jordà-Schularick-Taylor Macroeconomic History Database for which we have consistent real GDP data. The results are again similar up to statistical precision. The stability of our estimates across time periods suggests a lack of adaptation to temperature shocks, at least historically.

Our second sensitivity check concerns the selection of controls, to account for potential confounding effects. We show that a specification excluding region-specific time trends produces very similar results. To account for potentially strongly persistent impacts of the global economic cycle, we consider specifications where we control for 4 lags of all our global controls, and a specification where we control for 10 lags of world and country-GDP growth. Figure 5(c) shows that our estimates turn out to be virtually invariant to the set of controls. In Appendix A.10, we further establish that unobserved global shocks are not driving our results by exploiting an intermediate level of spatial aggregation of temperature shocks. The results from this specification do again not change much when we control for time fixed effects.

Our third sensitivity check investigates whether our results may be due to pre-trends. Although Table A.2 in Appendix A.2 already suggests that Granger causality is unlikely to be a concern, Figure 5(d) plots our main estimate together with estimates 6 years prior

Figure 5: The Average Effect of Global Temperature Shocks and Sensitivity




Notes: Impulse responses of real GDP per capita to a global temperature shock estimated in the panel using (3) over the period 1960-2019. Panel (a): baseline panel specification together with time-series response. Panel (b): results under shorter (1985-2019), excluding the Great Recession (1960-2007), and longer (1900-2019, with restricted set of countries) samples. Figure A.17 in Appendix A.14 reports individual impulse response functions for each sample together with confidence intervals. Panel (c): sensitivity with respect to controls: baseline (two lags of temperature shocks, country-level GDP, growth, world real GDP growth, oil prices, one-year US treasury yield, global recession dates and subregion-specific time trends); specification excluding subregion-specific time trends; specification with expanded set of global controls (four lags); specification that controls for 10 lags of world and country-GDP growth. Panel (d): baseline response with pre-trends. Lines: point estimate. Dark and light shaded areas: 90 and 95% confidence bands based on Driscoll and Kraay (1998) errors.

to the global temperature shock. The effect in the three years before the shock is zero by construction since 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.

Finally, we show in Appendix A.14 that our results are robust with respect a range of alternative choices. We find very similar results with GDP data from the World Develop-

Figure 6: Local and Global Temperature Shocks



Notes: 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 ($h = 2$, $p = 2$). Local shocks computed based on population-weighted country-level temperature data.

ment Indicators, with temperature data from NOAA and NASA, with the datasets in Dell et al. (2012) and Burke et al. (2015), and with different numbers of lags included in our local projections. 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 imply that a 1°C rise in local temperature reduces GDP at most by 1-2% in the medium run (Dell et al., 2012; Burke et al., 2015; Nath et al., 2022). To ensure that our findings are not driven by differences in econometric specifications or data choices, we reproduce the effects of local temperature shocks in our empirical framework. 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 6 shows local temperature shocks for the United States and South Africa over our sample from 1960, as two illustrative examples. The standard deviation of local temperature shocks is about three times larger than that of global temperature shocks. While local and global shocks have a correlation of 0.31, they frequently move in different directions. Thus, local shocks do not always correspond to 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). However, we also consider two alternative specifications:

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \theta_h^{\text{global}} T_t^{\text{shock}} + \theta_h^{\text{local}} T_{i,t}^{\text{shock}} + \mathbf{x}'_t \boldsymbol{\beta}_h + \mathbf{x}'_{i,t} \boldsymbol{\gamma}_h + \varepsilon_{i,t+h} \quad (4a)$$

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \delta_{t,h} + \theta_h T_{i,t}^{\text{shock}} + \mathbf{x}'_{i,t} \boldsymbol{\gamma}_h + \varepsilon_{i,t+h} \quad (4b)$$

In specification (4a), we estimate the impacts of global and local temperature shocks jointly. This provides a straightforward way to assess whether the two responses are statistically significant from each other. Specification (4b) includes time fixed effects, which allows us to flexibly control for any unobserved common shocks. In that case, the time fixed effects absorb the global temperature shocks and any other global controls.


Figure 7(a) shows the estimated impulse responses to a local temperature shock of 1°C, together with the responses to a global temperature shock from Section 3.1. Local temperature shocks lead to a fall in real GDP, even though the response is not statistically significant at the 5% level. On impact, the effect stands at -0.5% and reaches -1% after 4 years. These point estimates and associated uncertainty are similar to previous findings in Dell et al. (2012), Burke et al. (2015), and Nath et al. (2022) when aggregated across all countries, and of course mask a substantial degree of underlying heterogeneity. Controlling for time fixed effects does not make much of a difference. If at all, the inclusion of time fixed effects attenuates the impacts of local temperature somewhat.

The comparison reveals that *global* temperature has much more pronounced impacts on economic activity than *local* temperature. The estimated effects of global temperature shocks are larger than of local temperature shocks by an order of magnitude, based on the same empirical model (3) and the same sample period. This difference is not only *economically* but also *statistically* significant. Panel (b) reveals that the responses of GDP to global and local temperature shocks over our horizon of interest based on (4a) are statistically different at the 5% level in years 3 to 8.⁵

One possible explanation for the differential impact of global and local temperature shocks is *statistical*: global temperature shocks may lead to a more pronounced increase

⁵Estimating the effect of both shocks simultaneously does not change the univariate impulse responses materially, reflecting that different variation identifies the impact of global and local temperature shocks. See Appendix A.9 for details.

Figure 7: The Effect of Global vs. Local Temperature Shocks



Notes: Impulse responses of GDP per capita, estimated over the period 1960-2019. Panel (a): solid lines: GDP responses to global and local temperature shocks based on (3); dashed red: GDP response to local temperature shock from model with time fixed effect, (4b). Panel (b): difference between GDP responses to global and local temperature shocks based on joint model (4a). Panel (c): local temperature response to temperature shocks. Unweighted regression, implying that the time 0 impact of global temperature is close to 1°C. When area-weighted, the global temperature time 0 impact is larger than 1°C, as land warms more than oceans, see Figure A.14 in Appendix A.13. Panel (d): GDP responses imposing the persistence of the global temperature response on local temperature. Dark and light shaded areas: 90 and 95% confidence bands.

in local temperature. Figure 7(c) shows the response of local temperature to a local and a global temperature shock, respectively. On impact, both local and global temperature shocks lead to an increase in local temperature of about 1°C. Yet, the increase in local temperature is somewhat more persistent after a global temperature shock. To account for this difference in persistence, we construct a counterfactual local temperature shock, imposing the same internal persistence as for the global shock, using again the Sims (1986)

method. Figure 7(d) shows that the difference in persistence cannot account for the differential impact of global and local temperature shocks. Imposing the same persistence increases the impacts of local temperature somewhat, but the cumulative effects of global temperature shocks are still six times larger.

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: changing the set of controls or fixed effects does not affect the local temperature results meaningfully. Climatic variation within country or even smaller geographic units may help alleviate identification concerns, but misses any global effects of climate change—itsself a global phenomenon. By contrast, our approach purposefully studies these common effects by focusing on climatic variation at the global level.


3.3 Reconciling the Impacts of Global and Local Temperature

Why, then, does global temperature cause more economic harm than local temperature? We consider two possible *economic* explanations. The first explanation is that global temperature shocks are inherently different from local temperature shocks and capture potentially damaging climatic implications that local temperature does not. The second explanation is that local temperature is the true determinant of damages but compounds through economic spillovers that are however netted out in the panel specification.

Extreme Climatic Events. We start by investigating whether global temperature predicts meaningful shifts in climatic phenomena. We ask how temperature shocks correlate with the likelihood of extreme weather events: extreme temperature, drought, extreme precipitation, and extreme wind speed. As detailed in Section 2.1, we define an exposure index for each of these events by counting the fraction of cell-days within each year and country that exceed a given threshold. This exposure index can thus be interpreted as a probability. We use the panel local projection specification (3) and denote by θ_h^X the impact of a 1°C temperature shock on the exposure index of event X at horizon h . Figure 8 displays our results.

Local temperature shocks lead to an increase in the share of extreme heat and drought days. However, global temperature shocks lead to a substantially larger increase in these extremes. Our extreme heat and drought indices have a baseline probability of 0.05 and 0.25 in 1950-1980, respectively. Thus, a 1°C global temperature shock correlates with a five-fold rise in the frequency of extreme heat and a 15% increase of the frequency of

Figure 8: Extreme Weather Events and Temperature



Notes: Impulse responses θ_h^X of extreme temperature, drought, extreme precipitation, and extreme wind exposures to global and local temperature shocks, estimated based on (3) with the expanded set of global controls. Extreme weather exposure indices record the share of cell-days in a given year and country where temperature, precipitation, or wind speed are above or below a threshold. We define thresholds using the daily weather distribution in 1950-1980. Temperature: above 95th percentile. Drought: below the 25th percentile. Precipitation: above the 99th percentile. Wind: above the 99th percentile. Though not necessary for our results, we smooth the precipitation and wind measures with a backward-looking (current and previous two years) moving average to remove their inherent noise. Solid lines: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

drought, an order of magnitude more than for local temperature shocks. 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 have no significant effect. We construct the extreme precipitation and wind index to have a baseline probability of 0.01 in 1950-1980. Thus, a 1°C global temperature shock correlates with an increase of the frequency of extreme precipitation of over 50% and extreme wind of about 40%.


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 (Seneviratne et al., 2016; Wartenburger et al., 2017; Seneviratne et al., 2021; Domeisen et al., 2023). 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 correlation of global versus local temperature shocks on extreme climatic events may rationalize the larger economic effects of global temperature shocks.

To gauge the quantitative importance of this channel, we start by estimating the impact of local extreme events in a panel local projection specification similar to (3). We denote by ϕ_h^X the impact of extreme event X 's exposure index on GDP at horizon h . Figure A.13 in Appendix A.12 reveals that these events are associated with substantial economic damages. A five-fold rise in extreme heat exposure at the country level lowers GDP by 2% at peak. A 15% rise in drought exposure lowers GDP by 1%. A 50% increase in extreme precipitation lowers GDP by 0.5% and a 40% increase in wind exposure lowers output by 0.4%.

Next, we aggregate the local impacts of extreme events. We interact the increase in extreme event exposure following a global temperature shock θ_h^X from Figure 8 with the GDP loss associated with these extreme events from Figure A.13 in Appendix A.12. To do so, we adjust the estimates ϕ_h^X to correspond to a one-time fully transitory rise in exposure using again the method in Sims (1986). This persistence adjustment transforms the initial estimates ϕ_h^X into new estimates ψ_h^X . In practice, this adjustment has minor consequences because extreme events have low unconditional internal persistence. We then aggregate these impacts according to $\Theta_h = \sum_X \sum_{t=0}^h \theta_t^X \psi_{h-t}^X$, where the sum over X includes the four extreme events and local temperature. Thus, the aggregate impact Θ_h now factors in the persistent response of extreme events to a global temperature shock $\{\theta_h^X\}_h$.

Figure 9(a) displays our results. The rise in local temperature and extremes leads to a significantly larger economic impact than for local temperature alone. The peak effect on GDP is in excess of 6% and the cumulative impact is close to two thirds of the cumulative effect of a global temperature shock. This result indicates that global temperature has a larger impact on economic activity than local temperature because the *physical nature* of the shock is different: it captures the broader implications of warming and in particular the rise in damaging extreme events.

Figure 9: The Role of Ocean Temperature and Extreme Events



Notes: Panel (a): aggregated effect on GDP based on local temperature and extreme events impacts Θ_h (dashed red) together with the impulse responses to a global temperature shock based on our baseline empirical model (3). Panel (b): impulse responses of GDP to an ocean temperature and a land temperature shock, estimated jointly using (2). Dark and light shaded areas: 90 and 95% confidence bands.

We confirm our transmission channel of global temperature shocks through extreme events by considering ocean and land temperature. Since oceanic warming is critical for the formation of some of our extreme events, we expect it to account for a large part of our global temperature impacts. We jointly estimate the impact of ocean and land temperature on GDP in Figure 9(b). The impact of ocean temperature on GDP aligns with the overall effect of global temperature—if anything, it is somewhat larger—suggesting that ocean temperatures are key to understand the impact of warming on economic activity. The impact of land temperature is smaller than ocean temperature, though more pronounced than the impact of local temperature, suggesting that spatially correlated changes in local temperature may predict more extreme events than idiosyncratic temperature fluctuations. Yet, these comparisons are noisy and so should be interpreted with caution.

Our results highlight that it is critical to consider climatic outcomes beyond local temperature in panel approaches (Kotz et al., 2024), but also illustrates the challenges associated with such “bottom-up” aggregation exercises. Capturing all relevant local impacts individually is challenging: researchers need to know ex-ante which variables to consider, be able to measure them consistently throughout the world, and accurately estimate their degree of internal persistence. Even then, Figure 9(a) suggests that this “bottom-up” aggregation approach still underestimates the full impact of global temperature even with

four measures of extreme events. A key advantage of our time-series approach is that it directly encompasses all relevant local impacts that are predictable by global temperature.

Economic spillovers. Our analysis of extreme events suggests that there is limited scope left for economic spillovers to rationalize the large gap between global and local temperature impacts. We now confirm this argument quantitatively.

When the trading partners of a given country are hit by adverse local temperature realizations, some of the resulting economic consequences may also be felt domestically as hypothesized by Neal (2023) and shown in Dingel et al. (2023) and Zappalà (2023). How does explicitly accounting for such spillovers affect our global and local temperature shock estimates?

To gauge the relevance of economic spillovers, we exploit an intermediate level of spatial aggregation of local temperature shocks. We first construct an external temperature measure for each country that averages local temperature in surrounding countries, weighted by their respective trade share: $T_{i,t}^{\text{ex, trade}} = \sum_{j \neq i} \pi_{ij} T_{j,t}$, where π_{ij} denote trade shares based on imports plus exports between countries i and j in 1960. Next, we apply the same Hamilton (2018) filter to $T_{i,t}^{\text{ex, trade}}$ to construct an external temperature shock.


We expect the trade-weighted external temperature shock to have a substantial impact on GDP if economic spillovers explain the difference between local and global temperature impacts. In that case, the effect of global temperature shocks would also be largely absorbed by the external temperature shock when estimating the impacts jointly. We operationalize these ideas by estimating the impact of trade-weighted external temperature, both in isolation and jointly with global temperature and with local temperature.⁶

Trade-weighted external temperature turns out to moderately affect GDP. At short horizons, the effect is comparable to local temperature. At longer horizons, the effect even turns positive even though the response is not statistically significant (Figure A.12(a) in Appendix A.11). Figure 10(a) shows the responses to a trade-weighted external temperature shock and a global temperature shock, jointly estimated in the same local projection model. The response to the trade-weighted temperature shock from the joint model is very close to the estimated impact from the individual model. Global temperature continues to have a substantial adverse impact on GDP, even when controlling for trade-weighted external temperature.⁷

⁶Because we largely leverage regional variation here, we exclude the region-specific time trends from our set of controls. Omitting region-specific trends leaves our baseline results virtually unchanged.

⁷The shape of the response to the global temperature shock is slightly different from Figure 5 because

Figure 10: The Role of Economic Spillovers



Notes: Impulse responses of world real GDP per capita to external temperature shocks. Left panel: response to a global and trade-weighted shock, estimated jointly in the same local projection specification. Right panel: response to a local and trade-weighted shock, estimated jointly in the same local projection specification. Solid lines: point estimates. Dark and light shaded areas: 90 and 95% confidence bands. Sample of countries differs from main analysis due to availability of trade data at the beginning of the sample.

Figure 10(b) indicates that the effects of local temperature shocks slightly rise when controlling for trade-weighted temperature shocks. This change is consistent with exported spillover effects and spatially correlated temperature shocks. However, quantitatively, neither the direct effect of local temperature nor the indirect effect of trade-weighted external temperature are able to bridge the gap with global temperature impacts. Overall, these results suggest that economic spillovers play a limited role in accounting for the difference between global and local temperature.


3.4 Margins of GDP and Regional Impacts

We have documented that global temperature shocks lower world GDP, but how and where does GDP respond most?

We evaluate the effects of global temperature shocks on capital, investment and productivity in our panel of countries in Figure 11. Global temperature shocks lead to a substantial and significant fall in investment and in the capital stock. Consistently with Hsiang and Jina (2014), we find that disasters associated with global warming do not stimulate growth. Instead, national income, productive capital and investment all dwindle. Total Factor Productivity (TFP) as estimated in the Penn World Tables and labor produc-

we cannot obtain trade information for all countries at the beginning of our sample and must thus rely on a different set of countries.

Figure 11: Transmission of Global Temperature Shocks




Notes: Impulse responses of investment per capita, the capital stock per capita, total factor productivity and labor productivity to a global temperature shock, estimated based on (3) over the period 1960-2019. Labor productivity: output over employment. Total factor productivity: Penn World Tables. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

tivity fall significantly after global temperature shocks. The effects strengthen from -2% on impact to -10% after four years.

In addition to unpacking the margins of world GDP, we analyze how the impact of global temperature varies across different regions. Are warmer or lower-income countries more affected? Figure 12 displays the impact of global temperature shocks on nine regions of the world. No region gains. We estimate the strongest negative effects—close to -20% at peak—in hot regions such as Southeast Asia and Sub-Saharan Africa. Contrary to local temperature, global temperature leads to adverse economic effects even in higher-income, colder regions. The peak effect in North America is -10%, and in Europe is -7%, albeit not very precisely estimated.

Figure 12: Regional Impacts of Global Temperature Shocks



Notes: Impulse responses of GDP per capita to global temperature shocks for different regions across the world based on (3) over the period 1960–2019, conditioning on the different regions. Solid lines: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

We evaluate whether the impact of global temperature shocks systematically varies by country baseline temperature and income level in Figure A.24, Appendix A.15. Although somewhat imprecisely estimated, we find suggestive evidence that warm and low-income countries display the strongest adverse effects of global temperature shocks, while cold and high-income countries are less sensitive to global temperature shocks. This result is qualitatively consistent with previous evidence on local temperature (Dell et al., 2012; Burke et al., 2015; Nath, 2022). Quantitatively however, global temperature shocks have larger and more uniformly detrimental effects than local temperature shocks.

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 innovation is to use our reduced-form estimates of the impact of global temperature shocks to structurally estimate damage functions in the model.

4.1 Model Description

Setup. 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 flow preferences: $U(C) = \frac{C^{1-\gamma}-1}{1-\gamma}$. Labor supply is exogenous and set to $L_t = 1$. The pure rate of time preference of households is ρ .

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. Capital depreciates at rate δ , which is constant over time and covered by firms. The path of productivity Z_t is perfectly foreseen.

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 .

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

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

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

Climate change. We model climate change as changes in TFP Z_t over time, relative to its baseline value Z_0 . We take the path of global mean temperature T_t relative to a reference level T_0 as given, and denote by $\hat{T}_t \equiv T_t - T_0$ the path of excess temperature. Global mean

temperature affects TFP through the structural damage function $\{\zeta_s\}_{s \geq 0}$:

$$Z_t = Z_0 \exp \left(\int_0^t \zeta_s \hat{T}_{t-s} ds \right). \quad (6)$$

The structural damage function ζ_s governs the persistence of the effect of transitory global temperature shocks on TFP. When ζ_s is a Dirac mass point at $s = 0$, global temperature shocks have purely transitory level effects. When ζ_s is a positive function that asymptotes to zero, global temperature shocks have persistent level effects. When ζ_s is a positive function that asymptotes to a positive value, global temperature shocks have growth effects.

When temperature $T_t \equiv \bar{T}$ is constant, the economy converges to its steady-state with the corresponding value of TFP $\bar{Z} = Z_0 \exp \left((\bar{T} - T_0) \int_0^\infty \zeta_s ds \right)$. This expression highlights that the cumulative damage function $\int_0^\infty \zeta_s ds$ determines the long-run impact of global temperature changes. In that case, ζ_s needs to be integrable to obtain a well-defined steady-state. Hence, under growth effects, there is no well-defined steady-state and the economy asymptotes to zero for *any* amount of permanent warming. Yet, in Figure 3 we do not find evidence supporting growth effects and our estimates uncover persistent level effects instead.

We do not model the feedback between the economy and emissions, and associated externalities, because we focus on climate damages. Thus, the competitive equilibrium is efficient as is standard in the neoclassical growth model.

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.

We denote by $\{\hat{T}_t^{\text{SCC}}\}_{t \geq 0}$ the path of excess warming implied by a one-time pulse of one ton of CO2 emitted at time 0. The average response in Dietz et al. (2021) indicates that after a 1 gigaton pulse temperature rises steadily and eventually stabilizes at +0.002°C af-

ter 15 years. We remain conservative and use the lower end of available temperature responses: we define $\{\hat{T}_t^{\text{SCC}}\}_{t \geq 0}$ as three quarters of the the multi-model mean in Dietz et al. (2021), which stands within but at the lower end of the 95% interval that spans the range of climate model outcomes. Doing so ensures that historical emissions are consistent with historical warming data. When we use the multi-model mean, the SCC rises by a third. Our welfare numbers remain unchanged as they do not depend on the temperature response to a CO2 pulse, but instead on a particular warming scenario.

We then construct a productivity path $\{Z_t^{\text{SCC}}\}_{t \geq 0}$ according to equation (6) in which we use the temperature path $\{\hat{T}_t^{\text{SCC}}\}_{t \geq 0}$ 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 $K_0^{\text{SCC}} = K^{\text{ss}}$, 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 given implicitly by:

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

where *ss* superscripts denote initial 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 $\mathcal{C} = \int_0^\infty e^{-\rho t} u'(C^{\text{ss}})(C^{\text{ss}} - C_t^{\text{SCC}}) dt = \frac{1}{\rho} \frac{C^{\text{ss}} - \bar{C}^{\text{SCC}}}{C^{\text{ss}}}$, where $\frac{C^{\text{ss}} - \bar{C}^{\text{SCC}}}{C^{\text{ss}}}$ is the consumption-equivalent welfare loss from the warming implied by the CO2 pulse. These identities indicate that the SCC is equal to the present *stock* valuation of *flow* consumption-equivalent welfare losses from the 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.

4.2 Estimation Strategy

Our next step is to estimate the structural damage function ζ_s . To do so, we match the reduced-form impulse response functions of output to global temperature shocks from Figure 7. 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 function. We set risk-aversion to $\gamma = 1$. The capital share is $\alpha = 0.33$. The annual capital depreciation rate is $\delta = 0.08$. Our choice of annual

pure rate of time preference $\rho = 0.02$ follows Rennert et al. (2022) and is consistent with a 2% annual interest rate in steady-state.⁸ Of course, the equilibrium path of consumption in the model determines the effective consumption-based discount rate. We assess the robustness of our results with respect to the rate of time preference in Section 5.4 below.

In the second step, we invert our model to estimate the sequence of TFP that corresponds to a temperature shock. We leverage that the actual temperature shocks that arise during our sample are small as in Figure 6 and therefore imply output and capital fluctuations of the order of 1%. Therefore, we can use a first-order perturbation of the model around the initial steady-state. For any sequence of excess temperature $\hat{T}_t$, we denote by $\hat{z}_t$ the resulting log deviation in TFP, and by $\hat{y}_t$ the log deviation in output along the transition. We emphasize that we use log-linearization for *estimation* only, *not* for *counterfactuals*.

Proposition 1. (*Model inversion*)

There exists $\mathcal{K}_{t,s}$ given in Appendix B.3, that only depends on steady-state objects and is independent from $\{\zeta_s\}_{s \geq 0}$, such that, to a first order in $\{\hat{T}_t\}_{t \geq 0}$:

$$\hat{y}_t = \hat{z}_t + \alpha \int_0^\infty \mathcal{K}_{t,s} \hat{z}_s ds.$$

Proof. See Appendix B.3. □

Proposition 1 delivers an identification result. Given observed output response $\hat{y}_t$, we can recover the underlying sequence of productivity shocks $\hat{z}_t$. The first component in Proposition 1 corresponds to the direct effect of productivity on output. The second component corresponds to the equilibrium response of capital. It is an integral over all times because investment is forward-looking and capital accumulates slowly over time.

The main content of Proposition 1 lies in this second component. By log-linearizing equilibrium conditions and solving explicitly for the equilibrium sequence of capital, we relate capital deviations to the sequence of productivity shocks through the sequence-space Jacobian $\mathcal{K}_{t,s}$ (Auclert et al., 2021; Bilal and Goyal, 2023). In the context of the neoclassical growth model, this Jacobian admits a closed-form expression as a function of parameters and steady-state objects. When $\text{Id} - \alpha \mathcal{K}$ is invertible—where Id denotes the

⁸This framework immediately accommodates balanced productivity growth. Provided we adjust the rate of time preference and the baseline capital depreciation rate, standard rescaling arguments ensure that allocations and welfare would be identical in counterfactuals when the baseline economy is in steady-state or on a balanced growth path.

identity map, for instance when α is small enough—productivity shocks are identified. Proposition 1 allows us to obtain the sequence of TFP $\hat{z}_t$ that correspond to any sequence of temperature shocks $\hat{T}_t$.

We use Proposition 1 to estimate ζ_s . We consider the response of output to an *observed* temperature shock in Figure 3(b), that corresponds to the underlying temperature path $\hat{T}_t$ in Figure 3(a). Proposition 1 delivers the corresponding sequence of productivity shocks $\hat{z}_t$. We then identify ζ_t as the innovations to these sequences as per equation (6).

This approach is consistent with households having rational expectations about future temperature shocks: after a temperature shock, households expect temperature to remain persistently elevated as in Figure 3(a). One advantage of this approach is that we identify damage functions off of empirical impulse responses to a shock that is itself persistent. Thus, counterfactuals that focus on a permanent increase in temperature build on moments identified from responses to a persistent shock—though not a fully permanent shock—rather than a purely transitory shock.

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 to be of the form $A(e^{-Bs} - e^{-Cs})$.

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


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.

Hence, instead of exactly inverting the model, we estimate A , B and C for ζ_s using Non-Linear Least Squares to minimize the squared deviations from the equation in Proposition 1 for the first 10 years only.

4.3 Estimation Results

Figure 13 shows our estimation results. Panel (a) displays the underlying temperature path from Figure 3. Panel (b) reveals that the estimated model closely fits the empirical response of output given its limited degrees of freedom. Of course, the model fit relies on our constrained functional form: if we did not constrain the damage function, the fit would be one-to-one.

Figure 13: Output, Capital and Productivity Global Temperature Shocks



Notes: Estimation results from matching the model impulse response to the empirical response of output to global temperature shocks. Panel (a): underlying temperature path. Panel (b): output responses to this internally persistent temperature path. Panel (c): implied productivity shocks; confidence intervals based on the Delta-method. Panel (d): non-targeted capital responses to internally persistent temperature path. Dashed lines: data. Solid lines: model fit. 90% (dark area) and 95% (light area) confidence intervals.

Panel (c) depicts the estimated structural damage function ζ_s . It coincides with the productivity responses to a one-time transitory global temperature shock of 1°C. It implies a short-run productivity loss of 4% that takes place two years after the temperature shock. Despite the corresponding temperature shock being transitory, the impact on pro-

ductivity decays only slowly and persists for up to 10 years. The confidence bands reflect the confidence intervals around our empirical output response.

We test whether the estimated model also delivers empirically plausible implications for capital. In the estimation, we have not used any information about the empirical impulse response of capital to a global temperature shock. Panel (d) compares the prediction of our estimated model to the data from Figure 11. Despite being non-targeted, the response of capital in the model is close to its empirical counterpart.

How do the productivity effects of global temperature shocks compare to those associated with local temperature shocks? Given that the empirical responses are substantially smaller for local temperature shocks as shown in Figure 7, such shocks likely also imply smaller damages. To answer this question quantitatively, we repeat our estimation but targeting the impulse response of output to local temperature shocks.

Figure B.1 in Appendix B.4 displays the productivity effects of local temperature shocks. The cumulative productivity effect of local temperature shocks is more than seven times smaller than under global temperature shocks. We conclude that global temperature shocks have much larger effects on economic fundamentals.

5 The Welfare Impact of Climate Change

5.1 Representing Climate Change

To evaluate the consequences of climate change, we specify a path for global mean temperature. The baseline year $t = 0$ corresponds to 2024. The world subsequently warms by 3°C above pre-industrial levels by 2100, after which temperature asymptotes to 3.3°C. This scenario is broadly consistent with IPCC business-as-usual scenarios that imply 3 to 4°C of warming by 2100 (Lee et al., 2023). Given that the world has warmed by approximately 1°C since pre-industrial times, this scenario implies 2°C of additional warming since $t = 0$ (2024) by year $t = 76$ (2100).

We construct two counterfactuals to highlight the role of global temperature. In the first counterfactual, we use the structural damage function estimated under *global* temperature shocks ζ_s^{global} in Figure 13(c) to construct productivity changes using equation (6) together with excess temperature $\hat{T}_t$. In the second counterfactual, we instead use the structural damage function estimated under *local* temperature shocks ζ_s^{local} in Figure B.1(c), Appendix B.4, using again equation (6) together with the same excess temperature

path $\hat{T}_t$.

Our counterfactuals compare allocations and welfare in an economy that warms according to $\hat{T}_t$, to allocations and welfare in an economy that remains in steady-state under $\hat{T}_t \equiv 0$. Welfare losses from climate change are defined as an equivalent percent decline in steady-state consumption. The SCC is defined in equation (7) and is independent from the global warming scenario because it relies on the temperature response to a given CO2 pulse $\{\hat{T}_t^{\text{SCC}}\}_{t \geq 0}$. Conversely, the welfare calculations are independent from $\{\hat{T}_t^{\text{SCC}}\}_{t \geq 0}$. To solve for counterfactuals, we use standard global numerical methods to obtain the global solution—we only use log-linearization for estimation.

5.2 Welfare and the Social Cost of Carbon


Figure 14 presents our main results. Panel (a) depicts the path of global mean temperature. Panel (b) reveals that output drops rapidly as global temperature rises, relative to a world that is not warming. In 2050, output declines by 19%. In 2100, output is 46% below what it would have been without climate change. This substantial decline reflects accumulated productivity losses that eventually reach 34%. These impacts are statistically significant at the 5% level.

Panel (c) highlights the adverse impact of lower productivity on capital accumulation. Initially, investment rises as households anticipate lower income going forward and therefore save, following standard permanent income logic. Capital starts decumulating rapidly thereafter under the pressure of lower output. By 2100, capital is 37% below what it would have been without climate change.

Panel (d) reveals that consumption drops, eventually reaching a 37% loss by 2100. This substantial decline in consumption translates into large welfare losses. Panel (e) shows that the 2024 welfare impact of climate change amounts to a 25% 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 43% loss. All these values are statistically significant at the 5% level.

Our results indicate that the impact of climate change is substantial. The welfare cost of climate change is 640 times the cost of business cycles, or 10 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

Figure 14: Transitional Dynamics Under Climate Change



Notes: Transitional dynamics of the estimated model under the scenario in panel (a). Solid blue lines: model estimated under global temperature. Shaded blue: 90% confidence intervals. Dashed red lines: model estimated under local temperature. Shaded red: 90% confidence intervals. Confidence intervals based on the Delta-method.

to the effect of the 1929 Great Depression in the United States. 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 living in the 1929 Great Depression, *forever*.

Panel (f) uses our structural damage function to construct the SCC. We obtain a SCC of \$1,347 per ton. This value is more than seven times larger than the \$185 per ton value in Rennert et al. (2022). The 95% confidence interval for the SCC ranges from \$469 per ton to \$2,264 per ton. Despite non-trivial uncertainty, even the lower bound of that confidence interval is several times larger than conventional SCC estimates.

We demonstrate that our focus on global temperature shocks is the main driver of our conclusions. Under local temperature damage functions, Figure 14 shows that climate

change implies a long-run output decline of 6%, a present value welfare cost of 3% and a SCC of \$178 per ton. None of these effects are statistically significant at the 5 or 10% level. These values and the associated uncertainty are consistent with results in Nordhaus (1992), Dell et al. (2012), Burke et al. (2015), Nath et al. (2022), and Rennert et al. (2022). We conclude that global temperature effects are both larger and more precisely estimated than local temperature effects.


5.3 Growth Accounting

If the economic effects of global temperature are so large, why were they not noticed after nearly 1°C of global warming since 1960? We answer this question by analyzing the historical impact of climate change. We start the economy in 1960 and feed in the realized path of warming until 2019, after which we impose constant temperature. We construct counterfactual changes in output relative to a baseline economy that remains in steady-state. We then add these changes directly to the data.

Figure 15 displays the results. Panel (a) reveals that climate change is responsible for moderate but persistent reductions in the world’s annual growth rate. In the 1960s, there is little warming and so few effects on economic growth. By 2019, potential growth without climate change deviates more systematically from realized growth with climate change. Panel (a) highlights that historical warming occurs in small increments. Warming shocks thus have moderate economic year-to-year effects in comparison to other economic shocks. The analysis in Section 2 detects these effects that are otherwise hidden behind background economic variation.

Panels (b) and (c) show that the annual growth effects of climate change eventually accumulate because climate change is a permanent shift, despite having an initially moderate effect on growth. Panel (b) indicates that climate change reduces the world growth rate by as much as a third of baseline growth in the 21st century. Panel (c) shows that this growth slowdown implies that world GDP per capita would be 19% higher today had no warming occurred between 1960 and 2019. Even though in this counterfactual we hold temperature constant at its 2019 level in all subsequent years, economic losses continue to accumulate after 2019. These delayed impacts are due to the lagged productivity effects embedded in our estimated damage functions $\{\zeta_s\}_s$ and to the internal transitional dynamics of the neoclassical growth model. By 2040, output is 25% below its potential due to climate change: one quarter of the economic losses caused by past warming are yet to

Figure 15: Growth Accounting With Climate Change



Notes: Impact of past climate change on world GDP. Panel (a): world output growth rate with (solid blue) and without (dashed red) climate change. Horizontal lines: sample averages. Panel (b): fraction of growth rate lost to climate change (annual growth loss out of 1960-2019 mean). Horizontal line: sample average. Dashed line: linear regression fit. Panel (c): world output with (solid blue) and without (dashed orange) climate change, normalized to one in 1960.


materialize.

5.4 Sensitivity

Given the sizable magnitude of our results, we investigate which parameters may be particularly important for them. Figure 16 displays how our results depend on four key choices: the rate of time preference ρ , our treatment of expectations, 2100 global mean temperature, and the climate sensitivity.

Panel (a) shows 2024 welfare losses as a function of the rate of time preference ρ , and panel (b) shows the corresponding SCC. As expected, a higher rate of time preference lowers welfare losses and the SCC: households then discount more damages that are far in the future. Our baseline rate of time preference $\rho = 0.02$ is consistent with Rennert et al. (2022) and with the secular decline in interest rates. However, even at rates of time preference above 0.04 we still obtain sizable welfare losses in excess of 15%. The corresponding SCC remains two to three times as large as the high end of previous estimates, and still seven times larger than the SCC based on local temperature under the same rate of time preference. By contrast, as we approach very low discount rates consistent with Stern (2006), welfare losses exceed 30% and the SCC rises above \$3,000 per ton. Welfare losses are less sensitive to the discount rate than the SCC because welfare losses represent an annualized flow of losses, while the SCC is a discounted stock valuation.

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



Notes: Sensitivity of welfare costs and Social Cost of Carbon in 2024 with respect to the rate of time preference (ρ), 2100 global mean temperature, the climate sensitivity and treatment of expectations. Solid blue lines: model estimated using global temperature shocks under baseline expectations. Dotted blue lines: model estimated using global temperature shocks with temperature shock surprises. Dashed red lines: model estimated using local temperature shocks under baseline expectations with productivity shocks only.

Panels (c) and (d) show welfare losses and the SCC when we vary 2100 temperature relative to pre-industrial levels. Welfare losses under 10% 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 3°C to 4°C above pre-industrial levels under business as usual: global mean temperatures already largely exceed 1°C since pre-industrial levels, and 2023 reached 1.45°C since pre-industrial levels. By contrast, pessimistic scenarios under which global mean temperatures reach 6°C since pre-industrial levels in 2100 lead to present value welfare losses of 50%. Of course, in Panel (d), the 2024 SCC is independent from the warming scenario because it only depends on the temperature response to a CO₂ pulse.

Panels (e) and (f) display how the climate sensitivity affect our conclusions. The climate sensitivity governs how carbon emissions map into current and future warming.

Consequently, welfare losses to a given warming scenario in as in panel (e) are independent from the climate sensitivity. However, as shown in panel (f), the SCC is not. Our main analysis uses a strongly conservative climate sensitivity: three quarters of the mean climate sensitivity in Dietz et al. (2021). This choice allows our analysis to remain more closely consistent with the historical link between emissions and warming, but is the lower end of climate sensitivities produced by leading climate models. When we use the mean climate sensitivity, the SCC exceeds \$1,700 per ton. With a larger climate sensitivity, the SCC exceeds \$3,000 per ton.

Figure 16 also shows how our conclusions change when we treat household expectations differently. In our main estimation, we assume that households have rational expectations about the temperature path following a temperature shock. An alternative is to assume that households are surprised every period by persistently elevated temperatures following a temperature shock. Under this assumption, we linearly combine our estimated impulse response functions to obtain the output responses to a one-time transitory temperature shock. We then target these responses to a transitory shock to estimate structural damage functions, instead of estimating damage functions first as in our baseline. We provide more details in Appendix B.4. The dotted lines in Figure 16 displays our results under this alternative treatment of expectations. The results are very close to our baseline, highlighting that our baseline treatment of expectations is not driving our results.

This 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 \$1,367 per ton and a 25% welfare loss from a moderate warming scenario. These effects are comparable to experiencing the 1929 Great Depression, 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. Most decarbonization interventions cost \$80 per ton of CO₂ abated (Bistline et al., 2023). A conventional SCC value of \$185 per ton 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. 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 \$36 per ton, making unilateral emissions reduction prohibitively expensive. Under our new estimates however, the DCC of the United States becomes \$273 per ton and thus largely exceeds policy costs. In that case, unilateral decarbonization policy is cost-effective for large economies such as 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”. *National Bureau of Economic Research 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”. *National Bureau of Economic Research Working Paper Series* 31327.
- Bilal, Adrien and Shlok Goyal** (2023). “Some Pleasant Sequence-Space Arithmetic in Continuous Time”. *Social Science Research Network Working Paper* 4634993.
- 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 and Kyle Emerick** (2016). “Adaptation to Climate Change: Evidence from US Agriculture”. *American Economic Journal: Economic Policy* 8.3, 106–40.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel** (2015). “Global non-linear effect of temperature on economic production”. *Nature* 527.7577, pp. 235–239.
- Burke, Marshall, Mustafa Zahid, Noah Diffenbaugh, and Solomon M Hsiang** (2023). “Quantifying Climate Change Loss and Damage Consistent with a Social Cost of Greenhouse Gases”. *National Bureau of Economic Research Working Paper Series* 31658.
- Callahan, Christopher W. and Justin S. Mankin** (2023). “Persistent effect of El Niño on global economic growth”. *Science* 380.6649, pp. 1064–1069.
- 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”. *National Bureau of Economic Research Working Paper Series* 30678.
- Cruz, José-Luis and Esteban Rossi-Hansberg** (2023). “The Economic Geography of Global Warming”. *The Review of Economic Studies* 91.2, pp. 899–939.

- 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". *Social Science Research Network Working Paper* 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.
- Dingel, Jonathan I., Kyle C. Meng, and Solomon M. Hsiang** (2023). "Spatial correlation, trade, and inequality: Evidence from the global climate". *National Bureau of Economic Research Working Paper Series* 25447.
- Domeisen, Daniela I. V., Elfatih A. B. Eltahir, Erich M. Fischer, Reto Knutti, Sarah E. Perkins-Kirkpatrick, Christoph Schär, Sonia I. Seneviratne, Weisheimer Antje, and Heini Wernli** (2023). "Prediction and Projection of Heatwaves". *Nature Reviews Earth and Environment* 4, 36–50.
- 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*.
- 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 Series* 20352.

- Hsiang, Solomon M., Kyle C. Meng, and Mark A. Cane** (2011). "Civil conflicts are associated with the global climate". *Nature* 476, 438–441.
- Joos, F., R. Roth, J. S. Fuglestedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. 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.
- Kahn, Matthew E., Kamiar Mohaddes, Ryan N.C. Ng, M. Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang** (2021). "Long-term macroeconomic effects of climate change: A cross-country analysis". *Energy Economics* 104, p. 105624.
- 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".
- Kotz, Maximilian, Anders Levermann, and Leonie Wenz** (2024). "The Economic Commitment of Climate Change". *Nature* 628, pp. 551–557.
- Krusell, Per and Anthony A. Smith** (2022). "Climate change around the world". *National Bureau of Economic Research Working Paper Series* 30338.
- Lee, Hoesung, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barrett, 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-Møller** (2021). "Local projection inference is simpler and more robust than you think". *Econometrica* 89.4, pp. 1789–1823.
- Montiel Olea, José Luis, Mikkel Plagborg-Møller, Eric Qian, and Christian K. Wolf** (2024). "Double Robustness of Local Projections and Some Unpleasant VARithmetic". *National Bureau of Economic Research Working Paper Series* 32495.
- Moore, Frances C., Moritz A. Drupp, James Rising, Simon Dietz, Ivan Rudik, and Gernot Wagner** (2024). "Synthesis of Evidence Yields High Social Cost of Carbon Due to

- Structural Model Variation and Uncertainties". *National Bureau of Economic Research Working Paper Series* 32544.
- Moscona, Jacob and Karthik A. Sastry** (Oct. 2022). "Does Directed Innovation Mitigate Climate Damage? Evidence from U.S. Agriculture*". *The Quarterly Journal of Economics* 138.2, pp. 637–701.
- 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** (2022). "How Much Will Global Warming Cool Global Growth?" *National Bureau of Economic Research Working Paper Series* 32761.
- National Oceanic and Atmospheric Administration** (2005). *Volcanos and Climate*. [Online Resource](#).
- (2009). *Climate Change: Incoming Sunlight*. [Online Resource](#).
- (2023). *What are El Niño and La Niña?* [Online Resource](#).
- Neal, Timothy** (2023). "The Importance of External Weather Effects in Projecting the Macroeconomic Impacts of Climate Change". *UNSW Economics Working Paper* 2023-09.
- 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** (2024). "Climate defaults and financial adaptation". *European Economic Review* 170, p. 104866.
- 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". *Working Paper*.
- Seneviratne, S. I., X. Zhang, M. Adnan, W. Badi, C. Derczynski, A. Di Luca, S. Ghosh, I. Iskandar, J. Kossin, S. Lewis, F. Otto, I. Pinto, M. Satoh, S.M. Vicente-Serrano, M.**

- Wehner, and B. Zhou** (2021). “Chapter 11: Weather and Climate Extreme Events in a Changing Climate”. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*; Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.). Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 1513–1766.
- Seneviratne, Sonia I., Markus G. Donat, Andy J. Pitman, Reto Knutti, and Robert L. Wilby** (2016). “Allowable CO₂ Emissions Based on Regional and Impact-Related Climate Targets”. *Nature* 529, 477–483.
- Sims, Christopher A.** (1986). “Are forecasting models usable for policy analysis?” *Quarterly Review* 10, 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 Working Paper*.
- Wartenburger, R., M. Hirschi, M. G. Donat, P. Greve, A. J. Pitman, and S. I. Seneviratne** (2017). “Changes in Regional Climate Extremes as a Function of Global Mean Temperature: an Interactive Plotting Framework”. *Geoscientific Model Development* 10, 3609–3634.
- Zappalà, Guglielmo** (2023). “Estimating Sectoral Climate Impacts in a Global Production Network”. *IMF Working Paper* 2023/053.

Online Appendix

The Macroeconomic Impact of Climate Change

Adrien Bilal[†]

Diego R. Känzig[‡]

Contents

A. Empirics	52
A.1. Data	52
A.1.1. Economic Data	52
A.1.2. Climate Data	52
A.2. Statistical Properties of Global Temperature Shocks	56
A.3. Accounting for Estimation Uncertainty in Temperature Shocks	59
A.4. Accounting for the Persistence in the Temperature Response	60
A.5. Alternative Estimation Models	61
A.6. Searching for Influential Observations	64
A.7. Reverse Causality	65
A.8. Nonlinearities in the Impact of Temperature Shocks	71
A.9. Jointly Estimating Local and Global Shocks	73
A.10. The Role of Time Fixed Effects	74
A.11. External Temperature Shocks	75
A.12. Impacts of Extreme Events	76
A.13. Additional Empirical Results	78
A.14. Additional Robustness Checks	79
A.14.1. Results Based on One-step Forecast Error Temperature Shocks	83
A.15. Regional Impacts	89
B. Model	91
B.1. Equilibrium	91
B.2. Linearization	91
B.3. Model Inversion: Proof of Proposition 1	93
B.4. Estimation	95

[†]Stanford University, CEPR and NBER. E-mail: adrienbilal@stanford.edu.

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

A Empirics

A.1 Data

A.1.1 Economic Data

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

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

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. (2022). 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 2000 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 1901-2000 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). Temperature 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

Figure A.1: Global Average Temperature Since 1950



Notes: 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 $\approx 2\%$.

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.

Country-level temperatures. We use the Berkeley Earth gridded temperature data to construct population- and area-weighted country-level mean temperatures. In our analyses, 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 1950 and 1980, 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). We define these threshold percentiles such that the extreme heat, drought, extreme precipitation and extreme wind indices have a baseline probability of 0.05, 0.25, 0.01 and 0.01, respectively. 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 1950 and 1980. As a robustness exercise, we used alternative thresholds computed based on data from 1900 to 1930, yielding very similar results. 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).

Descriptive statistics. Our main data set spans the period from 1960 to 2019. We drop countries for which we have fewer than 20 non-missing observations of temperature and real GDP per capita. This leaves us with 173 countries. Our results are robust to restricting the selection of countries further. In Appendix [A.14](#), we replicate our results based on the original panel datasets used in Dell et al. (2012) and Burke et al. (2015).

In Table [A.1](#), we present some descriptive statistics on the main variables of interest. In Panel (a), we report statistics on our global time-series variables. In Panel (b), we show statistics for the country-level variables. Specifically, we report the number of non-

missing observations, the mean, median, standard deviation as well as the minimum and the maximum observation.

Table A.1: Descriptive Statistics

	Obs	Mean	SD	Median	Min	Max
<i>Panel (a): Global variables</i>						
Global temperature anomaly	60	0.36	0.30	0.34	-0.15	1.03
Global temperature shock	60	0.00	0.12	-0.01	-0.24	0.28
World real GDP per capita growth	59	2.06	1.47	2.13	-1.74	6.36
Oil price change	59	8.55	30.32	1.70	-47.79	167.83
US Treasury yield	60	5.04	3.31	5.00	0.12	14.78
<i>Panel (b): Country-level variables</i>						
Local temperature anomaly	10379	0.40	0.57	0.35	-1.89	3.33
Local temperature shock	10379	0.01	0.46	0.00	-2.59	2.89
Real GDP per capita growth	9090	2.07	6.31	2.23	-67.01	94.17
Investment per capita growth	8938	6.58	23.61	4.68	-98.36	499.01
TFP growth	5716	0.33	4.90	0.47	-65.22	83.10
Labor productivity growth	8353	1.75	6.63	1.79	-67.31	142.17
Extreme heat days	10379	0.10	0.08	0.08	0.00	0.87
Drought days	10033	0.29	0.10	0.27	0.05	0.91
Extreme precipitation days	10033	0.01	0.01	0.01	0.00	0.08
Extreme wind days	10033	0.01	0.01	0.01	0.00	0.06


Notes: Descriptive statistics for our global and country-level variables. We report the number of non-missing observations, the mean, standard deviation, median, and min and max for the main variables used in our analysis over the period 1960-2019.

A.2 Statistical Properties of Global Temperature Shocks

In this appendix, we discuss some of the statistical properties of global temperature shocks in more detail.

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.14, our results are robust with respect to the number of lags for the temperature shock.

Figure A.2: Autocorrelation of Global Temperature Shock



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

Table A.2: Granger-causality Tests

Variable	p-value
Real GDP	0.669
Population	0.885
Brent price	0.937
Commodity price index	0.842
Treasury 1Y	0.952
Overall	0.939

Notes: 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.


Forecastability. A desirable feature of “shocks” is that they should not be forecastable by past information (Ramey, 2016). In our context, if global temperature 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.2 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.

The role of El Niño and other temperature variability. Are our results driven by specific sources of temperature variability such as El Niño events? To answer this question, we net out variation coming from El Niño by controlling for ENSO indices in our main specification. The results are shown in Figure A.3.

The responses are similar to our baseline estimates, suggesting that our main results capture a common effect of global temperature on economic activity that does not depend heavily on being driven by El Niño or other sources of climate variability. A related concern is that major volcanic eruptions may affect world real GDP through other channels than temperature, for instance by limiting air travel. Controlling for volcanic eruptions also yields virtually unchanged results.

Figure A.3: The Role of El Niño and Other Temperature Variability



Notes: Impulse responses of real GDP per capita to a global temperature shock estimated based on (2), controlling for El Niño and volcanic eruptions. Dark and light shaded areas: 90 and 95% confidence bands for our baseline estimates.


²We 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.

A.3 Accounting for Estimation Uncertainty in Temperature Shocks

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 construct the confidence bands using bootstrapping techniques. We resample the shock and controls using a Wild bootstrap and then compute bootstrapped series of our outcome variables based on our autoregressive model. We repeat this procedure a 1,000 times and re-estimate our local projection specification for each iteration of the bootstrap. Based on the bootstrapped distribution, we can then compute confidence bands for all our objects of interest.

Figure A.4 compares the confidence bands based on our baseline lag-augmentation approach with the bootstrapped confidence bands. The coverage is similar, suggesting that taking estimation uncertainty in the global temperature shock into account turns out to be inconsequential in the context of our application.

Figure A.4: The Role of Estimation Uncertainty in Temperature Shocks



Notes: Impulse responses of real GDP per capita to a global temperature shock, estimated based on (2). Solid black line: point estimate. Dark and light shaded areas are 90 and 95% confidence bands based on our simple lag-augmentation approach. Red dotted and dashed lines: 90 and 95% confidence bands based on our bootstrap, taking estimation uncertainty in the temperature shock into account.

A.4 Accounting for the Persistence in the Temperature Response

As we have seen in the main text, global temperature shocks lead to a relatively persistent increase in temperature. We estimate this increase based on the following model:

$$T_{t+h} - T_{t-1} = \alpha_h + \phi_h^T T_t^{\text{shock}} + \mathbf{x}_t' \beta_h + \varepsilon_{t+h}, \quad (\text{A.1})$$

where $\{\phi_h^T\}_{h=0,\dots,H}$ is the global temperature response to a global temperature shock.

Approach. To account for the persistence in the temperature response, we construct the response to a counterfactual scenario where the global temperature increase is purely transitory, i.e. temperature increases by 1°C on impact and zero after. Following Sims (1986), we achieve this by introducing a series of shocks T_h^{shock} at each horizon h to impose the desired temperature response $\tilde{\phi}^T$. The series of shocks $\mathbf{T}^{\text{shock}}$ can then be obtained through

$$\underbrace{\begin{pmatrix} T_0^{\text{shock}} \\ T_1^{\text{shock}} \\ \vdots \\ T_H^{\text{shock}} \end{pmatrix}}_{\mathbf{T}^{\text{shock}}} = \underbrace{\begin{pmatrix} 1 & 0 & \cdots & 0 \\ \phi_1^T & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \phi_H^T & \phi_{H-1}^T & \cdots & 1 \end{pmatrix}^{-1}}_{(\Phi^T)^{-1}} \underbrace{\begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}}_{\tilde{\phi}^T}.$$

With the shock series $\mathbf{T}^{\text{shock}}$ implying a purely transitory temperature response at hand, the impulse responses of GDP, $\tilde{\theta}_h$, can be obtained through

$$\underbrace{\begin{pmatrix} \tilde{\theta}_0 \\ \tilde{\theta}_1 \\ \vdots \\ \tilde{\theta}_H \end{pmatrix}}_{\tilde{\theta}} = \underbrace{\begin{pmatrix} T_0^{\text{shock}} & 0 & \cdots & 0 \\ T_1^{\text{shock}} & T_0^{\text{shock}} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ T_H^{\text{shock}} & T_{H-1}^{\text{shock}} & \cdots & T_0^{\text{shock}} \end{pmatrix}}_{\mathcal{T}^{\text{shock}}} \underbrace{\begin{pmatrix} \theta_0 \\ \theta_1 \\ \vdots \\ \theta_H \end{pmatrix}}_{\theta}.$$


The obtained impulse responses $\tilde{\theta}$ correspond to the effects on GDP following global temperature shock that leads to a one-time, purely transitory increase in global temperature. Based on these responses it is then straightforward to compute the responses to global temperature shocks of arbitrary persistence. For inference, we rely on bootstrap-

ping techniques.

It is important to note that this method is not robust to the Lucas critique. The underlying assumption is that the effects of a series of unanticipated temperature shocks are equivalent to an anticipated path announced at time zero. However, given that economic agents have historically paid little attention to temperature shocks, this assumption may be less restrictive than in other contexts.

Results. Figure A.5 shows the results. Once accounting for the internal persistence of temperature, the effects on GDP become much less pronounced and persistent.

Figure A.5: The Effect of a Transitory Global Temperature Shock



Notes: Impulse responses of global mean temperature and world real GDP per capita to a global temperature shock, estimated based on (2) on the period 1960-2019, transformed to a counterfactual responses to a completely transitory temperature shock computed using Sims (1986) method. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

A.5 Alternative Estimation Models

In this appendix, we explore the sensitivity of our results with respect to alternative estimation models.

One-step local projection. Recall, our baseline empirical model consists of a two step approach: (i) estimate temperature shocks using the Hamilton (2018) filter and (ii) estimate the impulse responses using local projections.


Here, we present results if we instead estimate the responses directly, projecting cumulative changes in GDP directly on temperature levels. This is closer to the distributed

lag models commonly used in the literature (Dell et al., 2012; Burke et al., 2015).

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h T_t + \mathbf{x}_t' \boldsymbol{\beta}_h + \varepsilon_{t+h}, \quad (\text{A.2})$$

where y_t is the outcome variable of interest, T_t^{shock} is the temperature shock and θ_h is the dynamic causal effect of interest at horizon h . Importantly, $\mathbf{x}_t$ contains sufficient lags of GDP growth and temperature to account for the serial correlations in both variables.

Figure A.6: Impulse Responses based on One-Step Estimation



Notes: Impulse responses of global mean temperature and world real GDP per capita to a global temperature shock, estimated based on (A.2) on the period 1960-2019. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

The responses are very similar to our baseline. This should not come as a surprise. In fact, if we rely on one-step ahead forecast errors as the relevant temperature shock measure and include the same set of controls in the shock regression (1) and the local projection (2), the two approaches yield the exact same results by the Frisch–Waugh–Lovell theorem. This discussion also highlights that previous literature implicitly relied on “temperature shocks”, i.e. deviations from the temperature trend, implicitly or explicitly.

Vector autoregression model. Our main empirical specification relies on local projection techniques. In this appendix, we alternatively estimate the responses based on VAR techniques. Starting point is the following structural vector moving-average representation

$$\mathbf{Y}_t = \mathbf{B}(L)\mathbf{S}\varepsilon_t, \quad (\text{A.3})$$

where $\mathbf{Y}_t$ is a $k \times 1$ vector of annual time series, $\boldsymbol{\varepsilon}_t$ is a vector of structural shocks driving the economy with $\mathbb{E}[\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'] = \mathbf{I}$, $\mathbf{B}(L) \equiv \mathbf{I} + \mathbf{B}_1 L + \mathbf{B}_2 L^2 + \dots$ is a matrix lag polynomial, and $\mathbf{S}$ is the structural impact matrix.

Assuming that the vector-moving average process (A.3) is invertible, it admits the following VAR representation:

$$\mathbf{A}(L)\mathbf{Y}_t = \mathbf{S}\boldsymbol{\varepsilon}_t = \mathbf{u}_t, \quad (\text{A.4})$$


where $\mathbf{u}_t$ is a $k \times 1$ vector of reduced-form innovations with variance-covariance matrix $\mathbb{E}[\mathbf{u}_t \mathbf{u}_t'] = \boldsymbol{\Sigma}_u$ and $\mathbf{A}(L) \equiv \mathbf{I} - \mathbf{A}_1 L - \dots$ is a matrix lag polynomial. Truncating the VAR to order p , we can estimate the model using standard techniques and recover an estimate of $\mathbf{B}(L)$.

The main identification problem is then to find the structural impact matrix $\mathbf{S}$. From the linear relation between the structural shocks and the reduced-form innovations, we obtain the following covariance restrictions $\mathbf{S}\mathbf{S}' = \boldsymbol{\Sigma}_u$. We assume that temperature shocks can impact on all variables in the VAR contemporaneously, while other shocks only affect temperature with a lag. This is motivated by the fact that emissions increases usually translate into temperature with a substantial lag. The identifying restriction can be implemented via the Cholesky decomposition of $\boldsymbol{\Sigma}_u$, denoted by $\tilde{\mathbf{S}}$.

In terms of model specification, $\mathbf{Y}_t$ includes global temperature and real GDP growth. To mitigate concerns about non-invertibility, we also include oil price growth and the U.S. treasury yield. The lag order is set to 4 and we also include our recession dummies as an exogenous variable.

Figure A.7 shows the results. The estimated impacts turn out to be consistent with our local projection evidence. As expected, the shape of the impulse response is not exactly as in Figure 3 because the VAR extrapolates from the first four autocovariances between GDP and temperature to obtain impacts at higher horizons, while the local projection in Figure 3 directly estimates these impacts at higher horizons.

Figure A.7: VAR Responses



Notes: Impulse responses of global temperature and real GDP per capita to a global temperature shock, estimated based on our VAR model (A.4). Solid lines: point estimates. Dark and light shaded areas: 90 and 95% confidence bands.


A.6 Searching for Influential Observations

Section 2.2 displays the identifying variation in a scatter plot. The negative relationship between temperature and GDP turns out to be a robust one and does not appear to be driven by a particular set of extreme observations.

Nevertheless, there were two potentially influential temperature shocks: a strong negative temperature shock in 1964 that was followed by a significant economic upswing and a large positive temperature shock in 1977. This latter observation precedes the second oil shock and the following Volker disinflation, even though we already control for these events through our set of recession dummies.

To formally assess the role of influential observations, we perform a jackknife exercise. Specifically, we censor one shock value at a time to zero, and re-run our local projection. To account for the differential impact on our controls, we also include a dummy variable for the year we censor.

Figure A.8: Sensitivity of the Response to Global Temperature Shocks



Notes: Baseline response of real GDP per capita to a global temperature shock (in black), together with the responses obtained from the jackknife, censoring one shock value at a time (in gray). Dark and light shaded areas: 90 and 95% confidence bands for our baseline response.

Figure A.8 shows our baseline response in black, together with the responses from the jackknife exercise in gray. The estimated impact of temperature on GDP is not driven any single extreme shock. When censoring certain shocks we can get even bigger impacts, while when dropping others the effects can be somewhat attenuated. In all cases, the peak effect is always larger than 7-8% and well within the confidence bands. Excluding the 1977 shock value corresponds to one of the more attenuated responses in the jackknife. However, even in this case we still find a sizeable effect.

A.7 Reverse Causality

In this appendix, we describe how we account for reverse causality. We assume that we start from detrended, stationary variables. We specify for GDP:

$$y_t = \sum_{s=-\infty}^t T_s \theta_{t-s} + \varepsilon_t,$$

where T_s is the temperature deviation, and ε_t is a possibly autocorrelated shock. We are interested in estimating the vector θ . For temperature, we specify:

$$T_t = \sum_{s=-\infty}^t y_s \gamma_{t-s} + \tau_t,$$

where τ_t is a possibly autocorrelated shock, and we know γ . Without loss of generality, We normalize the variance of T_t and y_t to 1. The local projection estimates:

$$P_h^{YT} \equiv \text{Cov}[y_{t+h} - y_{t-1}, T_t | \mathbf{x}_{t-1}] \equiv \text{Cov}_{t-1}[y_{t+h} - y_{t-1}, T_t],$$

where $\mathbf{x}_{t-1}$ is our vector of controls, and we denote $\text{Cov}_{t-1}[\bullet, \bullet] \equiv \text{Cov}[\bullet, \bullet | \mathbf{x}_{t-1}]$. We have:

$$\begin{aligned} P_h^{YT} &= \text{Cov}_{t-1}[y_{t+h} - y_{t-1}, T_t] \\ &= \text{Cov}_{t-1} \left[\sum_{s=-\infty}^{t+h} T_s \theta_{t+h-s} + \varepsilon_{t+h} - y_{t-1}, T_t \right] \\ &= \text{Cov}_{t-1} \left[\sum_{s=-\infty}^{t+h} T_s \theta_{t+h-s} + \varepsilon_{t+h}, T_t \right] \\ &= \text{Cov}_{t-1} \left[\sum_{s=-\infty}^{t+h} T_s \theta_{t+h-s}, T_t \right] + \text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] \\ &= \text{Cov}_{t-1} \left[\sum_{s=-\infty}^{t-1} T_s \theta_{t+h-s}, T_t \right] + \text{Cov}_{t-1} \left[\sum_{s=0}^h T_{t+s} \theta_{h-s}, T_t \right] + \text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] \end{aligned}$$

The third equality holds because we include lagged GDP y_{t-1} in our vector of controls $\mathbf{x}_{t-1}$. From the last line, as long as we include sufficiently many lags of temperature and GDP in our vector of controls to cover the moving average structure, the first term is zero. We proceed under the assumption that we include sufficiently many lags. Hence,

$$\begin{aligned} P_h^{YT} &= \text{Cov}_{t-1} \left[\sum_{s=0}^h T_{t+s} \theta_{h-s}, T_t \right] + \text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] \\ &= \sum_{s=0}^h \theta_{h-s} \text{Cov}_{t-1} [T_{t+s}, T_t] + \text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] \end{aligned}$$

There are two sources of reverse causality: internal persistence (first term), and residual shocks to GDP (second term). Of course, residual shocks can also affect the covariance in

the first term, but one can condition this channel out with the deconvolution procedure that conditions on the realized temperature path after a shock.

We start with the second term:

$$\begin{aligned}
\text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] &= \text{Cov}_{t-1} \left[\varepsilon_{t+h}, \sum_{s=-\infty}^t y_s \gamma_{t-s} + \tau_t \right] \\
&= \sum_{s=-\infty}^t \gamma_{t-s} \text{Cov}_{t-1} [\varepsilon_{t+h}, y_s] + \text{Cov}_{t-1} [\varepsilon_{t+h}, \tau_t] \\
&= \gamma_0 \text{Cov}_{t-1} [\varepsilon_{t+h}, y_t] + \text{Cov}_{t-1} [\varepsilon_{t+h}, \tau_t] \\
&= \gamma_0 \text{Cov}_{t-1} [\varepsilon_{t+h}, y_t] \\
&= \gamma_0 \text{Cov}_{t-1} [\varepsilon_{t+h}, \theta_0 T_t + \varepsilon_t] \\
&= \gamma_0 \text{Cov}_{t-1} [\varepsilon_t, \varepsilon_{t+h}] + \gamma_0 \theta_0 \text{Cov}_{t-1} [\varepsilon_{t+h}, T_t]
\end{aligned}$$

The third equality obtains because we again assume that control for enough lags of GDP. The fourth equality obtains because we assume that structural shocks are orthogonal. The fifth equality follows from substituting the equation for output and noting that we include enough lags of temperature as controls. Re-arranging, we obtain:

$$\text{Cov}_{t-1} [\varepsilon_{t+h}, T_t] = \frac{\gamma_0}{1 - \gamma_0 \theta_0} \text{Cov}_{t-1} [\varepsilon_t, \varepsilon_{t+h}]$$

We now denote by $P_s^{TT} = \text{Cov}_{t-1} [T_{t+s}, T_t]$ the (observed) autocovariance function of the temperature process. We also denote by $E_s = \text{Cov}_{t-1} [\varepsilon_t, \varepsilon_{t+s}]$ the (unobserved) autocovariance function of the GDP residuals. We have shown that our local projection estimator is:

$$P_h^{YT} = \sum_{s=0}^h \theta_{h-s} P_s^{TT} + \frac{\gamma_0}{1 - \gamma_0 \theta_0} E_h.$$

The first term represents the how internal persistence to the temperature process affects our estimator. We start our discussion by abstracting from the bias in the second term.

If we are only interested in the response to a purely transitory temperature shock—i.e. the θ 's—then we can directly correct our estimator for this internal persistence using the observed autocovariance P_s^{TT} . However, if we want to reconstruct the unbiased GDP

response to a temperature shock with the same amount of persistence as in the data, i.e.

$$\tilde{\theta}_h = \sum_{s=0}^h \theta_{h-s} \text{Cov}_{t-1}[\tau_t, \tau_{t+s}],$$

we need to construct the autocovariance function of the structural temperature shocks $\mathcal{V}_s \equiv \text{Cov}_{t-1}[\tau_t, \tau_{t+s}]$. Since we assumed that we know the γ 's, we can simply residualize the temperature process using lagged GDP and the known γ 's and obtain the τ 's.

Now we turn to the second term. This term is the classic reverse causality bias. However, since we assume that we know γ_0 , we can construct E_h as a function of θ and known covariances, and then solve for θ . Indeed, we have:

$$\begin{aligned} E_h &= \text{Cov}_{t-1}[\varepsilon_t, \varepsilon_{t+h}] \\ &= \text{Cov}_{t-1} \left[y_t - \sum_{s=-\infty}^t \theta_{t-s} T_s, y_{t+h} - \sum_{s=-\infty}^{t+h} \theta_{t+h-s} T_s \right] \\ &= \text{Cov}_{t-1} \left[y_t - \theta_0 T_t, y_{t+h} - \sum_{s=0}^h \theta_{h-s} T_{t+s} \right] \\ &= \text{Cov}_{t-1}[y_t, y_{t+h}] - \theta_0 \text{Cov}_{t-1}[y_{t+h}, T_t] - \sum_{s=0}^h \theta_{h-s} \text{Cov}_{t-1}[y_t, T_{t+s}] + \theta_0 \sum_{s=0}^h \theta_{h-s} \text{Cov}_{t-1}[T_t, T_{t+s}] \\ &= P_h^{YY} - \theta_0 P_h^{YT} - \sum_{s=0}^h \theta_{h-s} P_s^{TY} + \theta_0 \sum_{s=0}^h \theta_{h-s} P_s^{TT}. \end{aligned}$$

The third equality obtains because we controls for enough lags. In the fourth equality we defined $P_h^{YY} = \text{Cov}_{t-1}[y_t, y_{t+h}]$ the known autocovariance function of output. We recognized the local projection P_h^{YT} . We defined as $P_s^{TY} = \text{Cov}_{t-1}[y_t, T_{t+s}]$ the local projection of temperature on output (the “reverse” of our baseline local projection).

Hence, we obtain the collection of equations (some nonlinear) indexed by h :

$$P_h^{YT} = \sum_{s=0}^h \theta_{h-s} P_s^{TT} + \frac{\gamma_0}{1 - \gamma_0 \theta_0} \left\{ P_h^{YY} - \theta_0 P_h^{YT} - \sum_{s=0}^h \theta_{h-s} P_s^{TY} + \theta_0 \sum_{s=0}^h \theta_{h-s} P_s^{TT} \right\}$$

The θ 's are the unknowns. Everything else is known or observable. The only nonlinearity comes from θ_0 . Conditional on θ_0 , these are linear equations. Hence, we examine the

equation for θ_0 separately. We obtain:

$$P_0^{YT} = \theta_0 + \frac{\gamma_0}{1 - \gamma_0\theta_0} \left\{ 1 - \theta_0 P_0^{YT} - \theta_0 P_0^{TY} + \theta_0^2 \right\},$$

where recall that we normalized $V_0 = 1$ and $Y_0 = 1$. We also note that $P_0^{YT} = P_0^{TY}$ by definition (but not at higher lags). Multiplying by $1 - \gamma_0\theta_0$,

$$P_0^{YT}(1 - \gamma_0\theta_0) = \theta_0(1 - \gamma_0\theta_0) + \gamma_0 \left\{ 1 - 2\theta_0 P_0^{YT} + \theta_0^2 \right\}.$$

Re-arranging, we observe that the quadratic terms cancel out. Hence the equation for θ_0 is actually also linear. We obtain:

$$\theta_0 = \frac{P_0^{YT} - \gamma_0}{1 - \gamma_0 P_0^{YT}}.$$

We have thus constructed an unbiased estimator of θ_0 .

Then given θ_h , we construct θ_{h+1} by induction. We have (using some changes of indices):

$$\begin{aligned} P_{h+1}^{YT} &= \theta_{h+1} + \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \\ &+ \frac{\gamma_0}{1 - \gamma_0\theta_0} \left\{ P_{h+1}^{YY} - \theta_0 P_{h+1}^{YT} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TY} - P_0^{YT} \theta_{h+1} + \theta_0 \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} + \theta_0 \theta_{h+1} \right\} \end{aligned}$$

Re-arranging:

$$\begin{aligned} \left(1 + \frac{\gamma_0}{1 - \gamma_0\theta_0} \left\{ -P_0^{YT} + \theta_0 \right\} \right) \theta_{h+1} &= P_{h+1}^{YT} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \\ &- \frac{\gamma_0}{1 - \gamma_0\theta_0} \left\{ P_{h+1}^{YY} - \theta_0 P_{h+1}^{YT} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TY} + \theta_0 \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \right\}. \end{aligned}$$

Further re-arranging:

$$\frac{1 - \gamma_0 P_0^{YT}}{1 - \gamma_0\theta_0} \theta_{h+1} = \frac{1}{1 - \gamma_0\theta_0} P_{h+1}^{YT} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} - \frac{\gamma_0}{1 - \gamma_0\theta_0} \left\{ P_{h+1}^{YY} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TY} + \theta_0 \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \right\}.$$

Therefore:

$$\begin{aligned}\theta_{h+1} = & \frac{1}{1 - \gamma_0 P_0^{YT}} P_{h+1}^{YT} - \frac{1 - \gamma_0 \theta_0}{1 - \gamma_0 P_0^{YT}} \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \\ & - \frac{\gamma_0}{1 - \gamma_0 P_0^{YT}} \left\{ P_{h+1}^{YY} - \sum_{s=0}^h \theta_s P_{h+1-s}^{TY} + \theta_0 \sum_{s=0}^h \theta_s P_{h+1-s}^{TT} \right\}.\end{aligned}$$

This correction delivers the θ 's after adjusting for reverse causality. We observe that the “classic” reverse causality adjustment scales with γ_0 .

We can then construct $\tilde{\theta}_h$, the response to a persistent temperature shock τ_t . We start from the unbiased θ 's. Then, we construct the autocovariance function of the τ 's, i.e. $\mathcal{V}$. We have:

$$\begin{aligned}\mathcal{V}_h &= \text{Cov}_{t-1}[\tau_{t+h}, \tau_t] \\ &= \text{Cov}_{t-1} \left[T_{t+h} - \sum_{s=-\infty}^{t+h} y_s \gamma_{t+h-s}, T_t - \sum_{s=-\infty}^t y_s \gamma_{t-s} \right] \\ &= \text{Cov}_{t-1} \left[T_{t+h} - \sum_{s=0}^h y_{t+s} \gamma_{h-s}, T_t - y_t \gamma_0 \right] \\ &= P_h^{TT} - \gamma_0 P_h^{TY} - \sum_{s=0}^h \gamma_{h-s} (P_s^{YT} - \gamma_0 P_s^{YY})\end{aligned}$$

The third equality follows from including enough controls. Then we construct:

$$\tilde{\theta}_h = \sum_{s=0}^h \theta_{h-s} \mathcal{V}_s.$$

Implementation. In practice, we need a sequence γ . We construct a central case, and some alternatives for robustness.

The central case uses the following parameters. We use $\eta = 1$ for CO2, CH4 and SO2: emissions move one-for-one with output, which is consistent with a Cobb-Douglas production function.

Average world CO2 emissions during our 1960-2019 sample are $\bar{E}^{\text{CO2}} = 22.5 \text{ Gt/y}$.³ The temperature response in Celsius to a 100 Gt pulse in Dietz et al. (2021) is well-

³See <https://ourworldindata.org/co2-emissions>.

approximated by:

$$100 \times \phi_h^{\text{CO}_2} = a_{100}^{\text{CO}_2} \times (e^{-b^{\text{CO}_2} \times h} - e^{-c^{\text{CO}_2} \times h}) + d_{100}^{\text{CO}_2} \times (1 - e^{-f^{\text{CO}_2} \times h})$$

$$a_{100}^{\text{CO}_2} = 0.1878, b^{\text{CO}_2} = 0.083, c^{\text{CO}_2} = 0.2113, d_{100}^{\text{CO}_2} = 0.1708, f^{\text{CO}_2} = 0.2113.$$

Then we define: $\gamma_h^{\text{CO}_2} = \eta \times \bar{E}^{\text{CO}_2} \times \phi_h^{\text{CO}_2}$.

We use CH₄ emissions of $\bar{E}^{\text{CH}_4} = 125 \text{ Mt/y}$.⁴ The temperature response in Celsius to a 1 Mt pulse in Azar et al. (2023) is well-approximated by:

$$\phi_h^{\text{CH}_4} = a^{\text{CH}_4} \times (e^{-b^{\text{CH}_4} \times h} - e^{-c^{\text{CH}_4} \times h}) + d^{\text{CH}_4} \times (1 - e^{-f^{\text{CH}_4} \times h})$$

$$a^{\text{CH}_4} = 4.9970, b^{\text{CH}_4} = 0.1230, c^{\text{CH}_4} = 0.1376, d^{\text{CH}_4} = 0.0109, f^{\text{CH}_4} = 0.0019.$$

Then we define: $\gamma_h^{\text{CH}_4} = \eta \times \bar{E}^{\text{CH}_4} \times \phi_h^{\text{CH}_4}$.

We use SO₂ emissions of $\bar{E}^{\text{SO}_2} = 100 \text{ Mt/y}$.⁵ The temperature response in Celsius to a 1 Mt pulse in Albright et al. (2021) is well-approximated by:

$$\phi_h^{\text{SO}_2} = F^{\text{SO}_2} \times (A_1^{\text{SO}_2} e^{-h/\tau_1^{\text{SO}_2}} + A_2^{\text{SO}_2} e^{-h/\tau_2^{\text{SO}_2}} + A_3^{\text{SO}_2} e^{-h/\tau_3^{\text{SO}_2}})$$

$$F^{\text{SO}_2} = -0.0051$$

$$A_1^{\text{SO}_2} = 0.2537, \tau_1^{\text{SO}_2} = 0.6700, A_2^{\text{SO}_2} = 0.0269, \tau_2^{\text{SO}_2} = 12, A_3^{\text{SO}_2} = 0.0010, \tau_3^{\text{SO}_2} = 352$$

Then we define: $\gamma_h^{\text{SO}_2} = \eta \times \bar{E}^{\text{SO}_2} \times \phi_h^{\text{SO}_2}$.

Finally, we define $\gamma = \gamma^{\text{CO}_2} + \gamma^{\text{CH}_4} + \gamma^{\text{SO}_2}$. Alternative, plausible choices of emissions-to-GDP elasticities or temperature sensitivity do not affect the reverse causality correction materially because it is small to begin with.

A.8 Nonlinearities in the Impact of Temperature Shocks

In this appendix, we investigate the role of potential non-linearities for the impact of temperature shocks. Our baseline model relies on the assumption of linearity. We motivate this feature by the fact that we only observe relatively small global temperature shocks in our sample. To further support this assumption, we estimate a local projection model, allowing for differential impacts of small and larger temperature shocks. Specifically, we

⁴See <https://www.iea.org/reports/global-methane-tracker-2023/overview>.

⁵See <https://ourworldindata.org/grapher/so-emissions-by-world-region-in-million-tonnes>.

consider temperature shocks below and above the 90th percentile of the global temperature shock distribution, which is 0.18°C , in absolute terms. The model we estimate is the following:

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h^{\text{small}} T_t^{\text{shock}} I(|T_t^{\text{shock}}| \leq z_{0.9}) + \theta_h^{\text{big}} T_t^{\text{shock}} I(|T_t^{\text{shock}}| > z_{0.9}) + \mathbf{x}_t' \boldsymbol{\beta}_h + \varepsilon_{t+h}, \quad (\text{A.5})$$

where θ_h^{small} are the dynamic causal effects of small and θ_h^{big} are the dynamic causal effects of big temperature shocks.

The results are shown in Figure A.9(a). The impact of small and large global temperature shocks is comparable. However, the responses are less precisely estimated. These observations motivate the use of both small and large shocks in our baseline specification.


We also investigate potential asymmetries in the impact of temperature shocks. Arguably, positive temperature shocks are more informative of future climate change than negative temperature shocks. How do the effects of positive and negative shocks compare? We consider a specification that allows for differential impacts of positive and negative shocks:

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h^{\text{pos}} T_t^{\text{shock}} I(T_t^{\text{shock}} > 0) + \theta_h^{\text{neg}} T_t^{\text{shock}} I(T_t^{\text{shock}} \leq 0) + \mathbf{x}_t' \boldsymbol{\beta}_h + \varepsilon_{t+h}, \quad (\text{A.6})$$

where θ_h^{pos} are the dynamic causal effects of positive and θ_h^{neg} are the dynamic causal effects of negative global temperature shocks.

Figure A.9(b) shows the results. Positive effects have comparable impacts to our baseline responses. If at all, the effects of positive shocks are somewhat more pronounced, even though they are less precisely estimated than when we include all shocks. The increased precision motivates our baseline specification, which includes both positive and negative shocks.

Figure A.9: The Role of Nonlinearities




Notes: Nonlinearities in the response of world real GDP per capita to a global temperature shock. Left panel: allowing for differential impacts of small shocks (in black) and large shocks (in red) based on (A.5). We use the 90th percentile of the temperature shock distribution as the threshold. Right panel: allowing for differential impacts of positive shocks (in red) against baseline (in black) based on (A.6). Dark and light shaded areas: 90 and 95% confidence bands.

A.9 Jointly Estimating Local and Global Shocks

As our baseline, we estimate the impacts of global and local temperature shocks separately, in individual models. In this appendix, we report the results when we estimate the impacts jointly in the same local projection specification.

Figure A.10 displays the results. The jointly estimated responses are very similar to our baseline responses. This is especially true for the impulse responses to the global temperature shock. For the local temperature shock, the effects are slightly attenuated, and lie closer to the responses from the univariate model with time fixed effects than to the responses from the univariate model with global controls. This is intuitive as both, global temperature shocks and time fixed effects net out common variation in local temperature shocks.

Figure A.10: Joint Responses of Local and Global Temperature Shocks



Notes: Impulse responses of GDP per capita to global and local temperature shocks, estimated based on (4a). Lines: point estimates. Dark and light shaded areas: 90 and 95% confidence bands.


A.10 The Role of Time Fixed Effects

In this appendix, we shed further light on the role of time fixed effects. Figure A.11(a) compares the impulse responses of GDP to local temperature shocks with and without time fixed effects. The responses from the local temperature shock specification with time fixed effects are strikingly close to the baseline with global controls. The coverage of the confidence bands is also comparable. Overall, these results suggest that our controls successfully account for common economic shocks.

While we cannot include time fixed effects in our baseline specification with global temperature, we alternatively consider a specification based on an intermediate level of aggregation. Specifically, we construct distance-weighted external temperature shocks. To do so, we first construct an external temperature measure that weights the shocks in the surrounding countries by their physical distance: $T_{i,t}^{\text{dist, ex}} = \sum_{j \neq i} d_{ij} T_{j,t}$, where d_{ij} is proportional to the inverse geodesic distance between countries i and j and are normalized to sum to one for each country i . We source these distances from the TRADHIST database. In a next step, we create a distance-weighted external temperature shock measure by applying the Hamilton (2018) filter.

These shocks are close in spirit to our land-based global temperature shocks, but because the weights differ by country, the shock measure will too. Therefore, we are able to control for time fixed effects in this setting.

Figure A.11: The Role of Time Fixed Effects



Notes: Impulse responses of real GDP per capita. Responses to a local temperature shock (Panel (a)) and responses to a distance-weighted external temperature shock (Panel (b)), estimated based on a specification with global controls (3) or with time FE (4b). Solid line: point estimate. Dark and light shaded areas/dashed and dotted lines: 90 and 95% confidence bands. Sample of countries differs from main analysis due to availability of trade data at the beginning of the sample.

Figure A.11(b) shows how the inclusion of time fixed effects affects the response of distance-weighted temperature shocks. In both cases, real GDP falls significantly with a peak effect in excess of 5%. The shape of the responses is a bit different, with the model without time fixed effects generating more front-loaded impacts. However, the cumulative effect is very similar in both cases, suggesting that our larger impacts of global temperature shocks are not the consequence of an omitted control variable due to the omission of time fixed effects.


A.11 External Temperature Shocks

As discussed in Section 3.3, we also construct external trade-weighted temperature shocks, in addition to the external distance-weighted temperature shocks discussed in Appendix A.10 above. Trade weights are based on total trade flows (imports plus exports) in 1960 from [TRADHIST](#). We then normalize total trade flows such that all weights sum to one for any given country.

Figure A.12 shows the impulse responses of world real GDP to the two external temperature shocks, estimated in separate specifications. Temperature shocks to a country's trading partners (trade-weighted external temperature shocks) lead to a notable fall in output, comparable to the impacts of local, idiosyncratic temperature shocks at first. The response, however, turns out to be less persistent and reverses after about 5 years.

Temperature shocks to a country's neighbors (distance-weighted external temperature shocks) have more pronounced effects on GDP. The peak effect materializes after 4 years and stands around 5%. The response features also a considerable degree of persistence. Overall, the estimated impacts are comparable to the effects estimated for common land temperature shocks, and substantially larger than for idiosyncratic local temperature shocks.

Figure A.12: The Impact of External Temperature Shocks




Notes: Impulse responses of world real GDP per capita to external temperature shocks, estimated based on (3). Left panel: response to a trade-weighted shock. Right panel: responses to a distance-weighted shock. Solid lines: point estimates. Dark and light shaded areas: 90 and 95% confidence bands. Sample of countries differs from main analysis due to availability of trade data at the beginning of the sample.

A.12 Impacts of Extreme Events

Figure 8 shows that global temperature shocks strongly correlate with the exposure to extreme weather events: extreme temperature, drought, extreme precipitation, and extreme wind. Here we project world real GDP on extreme event exposure directly, so that we can aggregate up the impact of global temperature on GDP through extreme events. We use

the panel local projection specification (3), except that we replace the global temperature shock on the right-hand-side with extreme event exposure. We denote by ψ_h^X the impact of an increase in exposure for extreme event X on country-level GDP at horizon h .

Figure A.13: The Impact of Extreme Events on GDP



Notes: Impulse responses of world real GDP per capita to extreme events, estimated based on (3) with our expanded set of controls. Extreme weather variables record the share of cell-days in a given year and country where temperature, precipitation, or wind speed are above/below a threshold. We define threshold using the daily weather distribution in 1950-1980. Temperature: above 95th percentile. Drought: below 25th percentile. Extreme precipitation: above 99th percentile. Wind: above 99th percentile. Though not necessary for our results, we smooth the precipitation and wind measures with a backward-looking (current and previous two years) moving average to remove their inherent noise. Responses are normalized to the peak increase in frequency from Figure 8: graphical responses report $\psi_h^X / (\max_t \theta_t^X)$. Solid lines: point estimates. Dark and light shaded areas: 90 and 95% confidence bands.

Figure A.13 displays our results. Graphically, we normalize the estimated impact normalized by the peak frequency rise in exposure from Figure 8 to ease interpretation: we report $\frac{\phi_h^X}{\max_t \theta_t^X}$. Extreme weather events lead to a significant and persistent fall in GDP. The response is particularly pronounced for extreme heat and extreme precipitation and


droughts but also extreme wind has substantial adverse effects, even though somewhat less precisely estimated.

To construct the aggregate impact of global temperature on GDP through extreme events, we further need to adjust the estimates ϕ_h^X for internal persistence. Underlying the estimates in Figure A.13, extreme event exposures turn out to have very low internal persistence. Thus, the estimates in Figure A.13 largely represent the GDP impact of a one-time increase in extreme events. Nevertheless, we convert the estimates ϕ_h^X in response to a realized rise in extreme event frequency to estimates in response to a one-time fully transitory surge in extreme event frequency using the method in Sims (1986). We denote those adjusted estimates by ψ_h^X . In practice, the ϕ_h^X and ψ_h^X are close. We then aggregate these estimates using the definition of Θ_h in Section 3.3.

A.13 Additional Empirical Results

In this appendix, we present some additional empirical results.


Effect of global temperature shocks on ocean and land temperature. Figure A.14 shows the impulse response of (area-weighted) ocean surface and land temperature. For our land temperature measure, we exclude Antarctica. As expected given that land warms more than oceans, a global temperature shock of 1°C leads to an increase in ocean surface temperature by around 0.8°C and an increase in land temperature by about 1.5°C



Notes: Impulse responses of ocean surface and land temperature to a global temperature shock, estimated based on (2). Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands.

Additional local temperature shock responses. Figure A.15 shows the responses of investment, capital and productivity to a local temperature shock. As for the output responses, these effects are by an order of magnitude smaller than for global temperature shocks. Controlling for time fixed effects makes again little difference.

Figure A.15: Transmission of Local Temperature Shocks



Notes: Impulse responses of investment per capita, the capital stock per capita, total factor productivity and labor productivity to a global temperature shock, estimated based on panel local projections (3). Labor productivity: output over employment. Total factor productivity: Penn World Tables. Solid line: point estimate from specification with global controls. Dark and light shaded areas: 90 and 95% confidence bands. Dashed line: point estimate based on model with time fixed effects.

A.14 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. We start by examining


the role of data choices, the construction of the temperature shock, and the number of lags included.

Figure A.16 collects the results. Panels (a)-(b) assess the sensitivity with respect to the GDP and temperature data we use. 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 Panel (c), we replicate our results with the datasets from Burke et al. (2015) and Dell et al. (2012). We obtain their datasets from the respective replication packages, merge our global temperature shock, and compute the impulse responses to the shock. We obtain similar results with their datasets.

Panel (d) assesses additional ways of constructing of the temperature shocks. 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.


Panels (e)-(f) evaluate 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 baseline value and vice versa. Our results turn out to be robust with respect to the lag order. In fact, in the main text, we show that our results even survive when we control up to 10 lags of real GDP.

Figure A.16: Sensitivity of the Average Effect of Global Temperature Shocks



Notes: Sensitivity of the effects of global temperature shocks on real GDP per capita to a global temperature shock, with respect to data choices, the construction of the temperature shock, and the number of lags included. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands, respectively.

Figure A.17: Sensitivity of the Average Effect of Global Temperature Shocks II



Notes: Sensitivity of the effects of global temperature shocks on real GDP per capita to a global temperature shock, with respect to the sample period. Solid line: point estimate. Dark and light shaded areas: 90 and 95% confidence bands, respectively.

Figure 5 in the main text shows that our point estimates are similar in a longer sample (1900-2019, based on a smaller selection of countries), a shorter sample (1985-2019), and stopping the sample prior to the Great Recession (1960-2007).⁶ However, are the responses also statistically significant? Figure A.17 shows the impulse responses on the three periods that we consider, together with the associated confidence bands. The esti-

⁶The longer sample includes 18 advanced economies, specifically Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, UK, and USA. We control for the world wars using a dummy for the years 1914-1918 and 1939-1945. We further control for the Great Depression and the Great Recession using a recession dummy (1929-1939 and 2007-2009). In the shorter sample, we add the European debt crisis to our recession dummy (2011-2012). This period was marked by low growth globally and is relatively more important in the shorter sample, which is why we control for it.

ated impact is significant in all the alternative sample periods.


Overall, these results further illustrate the robustness of our finding that global temperature shocks 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.

A.14.1 Results Based on One-step Forecast Error Temperature Shocks


As our baseline, we measure the temperature shocks as two-step ahead forecast errors, motivated by the period of the climatic variation we aim to capture. A more common choice in the literature is to construct temperature shocks as one-step ahead forecast errors, as in Bansal and Ochoa (2011) and Nath et al. (2022). We have already showed that the GDP response is virtually identical when using the one- or two-step ahead forecast error as the relevant shock measure. For completeness, we present all our main results based on the one-step ahead temperature forecast error. The results are shown in Figures A.18-A.23. Our results turn out to be virtually identical using this alternative shock measure.

Figure A.18: Alternative Global and Local Temperature Shocks

(a) Global temperature shock




(b) Global vs. local temperature shock




Notes: Panel (a): Global temperature shocks, computed as in Hamilton (2018) with $(h = 1, p = 2)$, over the post-World War II era. Panel (b): 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 shocks computed based on the Hamilton (2018) approach with $(h = 1, p = 2)$, over our sample from 1960. Local shocks computed based on population-weighted country-level temperature data.

Figure A.19: Time-series Results Based on One-Step Ahead Forecast Error




Notes: Reproduces Figures 3 and 4 for a global temperature shock measured as a one-step ahead forecast error.

Figure A.20: Panel Results Based on One-Step Ahead Forecast Error I




Notes: Reproduces Figures 5, 7 and 10 for global, local and external temperature shocks measured as a one-step ahead forecast errors.

Figure A.21: Panel Results Based on One-Step Ahead Forecast Error II




Notes: Reproduces Figures 8 and 9(a) for global, local, and external temperature shocks measured as a one-step ahead forecast errors.

Figure A.22: Panel Results Based on One-Step Ahead Forecast Error III



Notes: Reproduces Figure 11 for a global temperature shock measured as a one-step ahead forecast error.

Figure A.23: Panel Results Based on One-Step Ahead Forecast Error IV




Notes: Reproduces Figure 12 for a global temperature shock measured as a one-step ahead forecast error.

A.15 Regional Impacts

We study how the impact of global temperature varies 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 A.24 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.

Figure A.24: Heterogeneous Effects of Global Temperature Shocks



Notes: Impulse responses of real GDP per capita to a global temperature shock for different groups of countries. We estimate these responses based on (3), with our expanded set of controls, conditioning on the different groups. 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. Solid line: point estimate. Dark and light shaded areas: 95% confidence bands.

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: they are more uniformly detrimental than local temperature shocks. Temperate countries also display a response that is economically large. Only colder countries display a somewhat smaller effect that is also not statistically significant.

Figure A.24(b) shows the responses by income per capita. 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 8,000 USD), and high income countries (real GDP per capita above 8,000 USD). Relative to the evidence based on local temperature, we find again more uniformly detrimental impacts: real GDP per capita falls across all income groups. Poor countries experience the most significant and persistent decline. Middle-income countries also see a considerable decrease in output. Only high-income countries are relatively more insulated, with a somewhat smaller and less enduring impact.

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:

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

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

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

In steady-state,

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

B.2 Linearization

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

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

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

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

We define:

$$X_t = \begin{pmatrix} \hat{K}_t \\ \hat{C}_t \end{pmatrix}.$$

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

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

where:

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

We have an initial condition $\hat{K}_0$, and a terminal condition $\hat{C}_t \rightarrow 0$. 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 $\mathcal{X}_t$:

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

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 small log deviations:

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

Since (in small log deviations):

$$\widehat{y}_t = \widehat{z}_t + \alpha \widehat{k}_t,$$

we are left with calculating $\widehat{k}_t$ as a function of the sequence $\widehat{z}$. We express:

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

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

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

where we defined:

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

Thus,

$$\widehat{k}_t = \frac{e^{-\kappa t}}{K} \int_0^\infty \left\{ B_1 \mathbb{1}[s \leq t] e^{\kappa s} - (\varepsilon_S B) J_{t,s} \right\} \widehat{z}_s ds \equiv \int_0^\infty \mathcal{K}_{t,s} \widehat{z}_s ds,$$

where we defined:

$$\mathcal{K}_{t,s} = \frac{e^{-\kappa t}}{K} \left\{ B_1 \mathbb{1}[s \leq t] e^{\kappa s} - (\varepsilon_S B) J_{t,s} \right\}.$$

Hence, we have obtained:


$$\widehat{y}_t = \widehat{z}_t + \alpha \int_0^\infty \mathcal{K}_{t,s} \widehat{z}_s ds.$$

These equations conclude the proof of Proposition 1.

B.4 Estimation

Figure B.1 displays the productivity and capital depreciation effects of local temperature shocks, discussed in the main text.

Figure B.1: Output, Capital and Productivity After Local Temperature Shocks



Notes: Estimation results from matching the model impulse response to the empirical response of output to local temperature shocks. Panel (a): underlying temperature path. Panel (b): output responses to this internally persistent temperature path. Panel (c): implied productivity shocks. Panel (d): non-targeted capital responses to internally persistent temperature path. Dashed lines: data. Solid lines: model fit. 90% (dark area) and 95% (light area) confidence intervals based on the Delta-method.

An alternative estimation strategy is to 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 *before* matching the model to the data. The interpretation of this approach is that households are surprised by elevated temperature each period after a global temperature shock.

We follow Sims (1986) to obtain the impulse response to one-time transitory temperature shocks. It is equivalent to using a recursive approach. Indeed, denote by $\tilde{y}_t$ the un-

known impulse response function of output to a transitory temperature shock. In discrete data and under linearity: $\hat{y}_t = \sum_{s=0}^t \hat{T}_{t-s} \tilde{y}_s$. We then obtain $\tilde{y}_t = (\hat{y}_t - \sum_{s=0}^{t-1} \hat{T}_{t-s} \tilde{y}_s) / \hat{T}_0$ 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 $\hat{z}_t, \hat{\Delta}_t$. We then identify $\zeta_s = \hat{z}_s$ and $\delta_s = \hat{\Delta}_s / \Delta_0$.

References Appendix

- Albright, Anna Lea, Cristian Proistosescu, and Peter Huybers** (2021). “Origins of a Relatively Tight Lower Bound on Anthropogenic Aerosol Radiative Forcing from Bayesian Analysis of Historical Observations”. *Journal of Climate* 34, 8777–8792.
- Azar, Christian, Jorge García Martín, Daniel JA. Johansson, and Thomas Sterner** (2023). “The Social Cost of Methane”. *Climatic Change* 176.
- Bansal, Ravi and Marcelo Ochoa** (2011). “Temperature, Aggregate Risk, and Expected Returns”. *National Bureau of Economic Research Working Paper Series* 17575.
- 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** (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** (2022). “How Much Will Global Warming Cool Global Growth?” *National Bureau of Economic Research Working Paper Series* 32761.
- NOAA National Centers for Environmental Information** (2023a). *Climate at a Glance: Global Time Series*. [Online Resource](#). (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, pp. 2–16.