

NBER WORKING PAPER SERIES

HOW MUCH WILL GLOBAL WARMING COOL GLOBAL GROWTH?

Ishan B. Nath
Valerie A. Ramey
Peter J. Klenow

Working Paper 32761
<http://www.nber.org/papers/w32761>

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

We are grateful to David Hémous, Ben Jones, and James Stock for thoughtful paper discussions, and to Adrien Bilal, Marshall Burke, Tamma Carleton, Steve Cicala, Graham Elliott, James Hamilton, Solomon Hsiang, Oscar Jorda, Ezra Oberfield, Richard Rogerson, Barbara Rossi, Esteban Rossi-Hansberg, Christopher Sims, John Van Reenen, and participants at numerous conferences and seminars for helpful comments. Jean-Felix Brouillette, Valeska Kohan, and Walker Lewis provided excellent research assistance. Any views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve System, its staff, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Ishan B. Nath, Valerie A. Ramey, and Peter J. Klenow. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Much Will Global Warming Cool Global Growth?

Ishan B. Nath, Valerie A. Ramey, and Peter J. Klenow

NBER Working Paper No. 32761

July 2024

JEL No. E01,O44

ABSTRACT

Does a permanent rise in temperature decrease the level or growth rate of GDP in affected countries? Differing answers to this question lead prominent estimates of climate damages to diverge by an order of magnitude. This paper combines indirect evidence on economic growth with new empirical estimates of the dynamic effects of temperature on GDP to argue that warming has persistent, but not permanent, effects on growth. We start by presenting a range of evidence that technology flows tether country growth rates together, preventing temperature changes from causing growth rates to diverge permanently. We then use data from a panel of countries to show that temperature shocks have large and persistent effects on GDP, driven in part by persistence in temperature itself. These estimates imply projected future impacts that are three to five times larger than level effect estimates and two to four times smaller than permanent growth effect estimates, with larger discrepancies for initially hot and cold countries.

Ishan B. Nath
Federal Reserve Bank of San Francisco
inath33@gmail.com

Valerie A. Ramey
Hoover Institution
Stanford University
434 Galvez Mall
Stanford, CA 94305
and NBER
vramey@stanford.edu

Peter J. Klenow
Department of Economics
579 Jane Stanford Way
Stanford University
Stanford, CA 94305-6072
and NBER
Klenow@Stanford.edu

1 Introduction

Projections of the economic impact of global warming play a critical role in research on optimal climate policy (e.g. Golosov, Hassler, Krusell and Tsyvinski, 2014; Acemoglu, Akcigit, Hanley and Kerr, 2016; Barrage, 2020), cost-benefit analysis on emissions reductions proposals (e.g. Ricke, Drouet, Caldeira and Tavoni, 2018; Burke, Davis and Diffenbaugh, 2018), and analysis of climate change adaptation (e.g. Desmet and Rossi-Hansberg, 2015). Despite the central importance of the mapping from global temperature change to global GDP in climate change economics, there is no consensus on the likely size of the effects. The most commonly used estimates in the literature differ by an order of magnitude, with first order implications for climate policy (Moore and Diaz, 2015). Estimates that follow from the seminal Nordhaus (1992) DICE model suggest that, in the no abatement emissions scenario, temperature changes will cost the global economy approximately two to three percent of GDP in 2099.¹ In contrast, a second strand of damage estimates that follow from the work of Burke, Hsiang and Miguel (2015) suggest that global warming will cost the world economy 20 to 30% of GDP in 2099.

The sharp divergence of estimates arises from disagreement over whether a permanent change in temperature will affect levels or growth rates of income in the long run. Damage estimates in DICE and related work are calibrated to evidence on sector-by-sector climate change impacts (see Nordhaus and Moffat (2017) for a summary) that allow warming in each future period to affect output only in that period.² Conversely, Burke, Hsiang and Miguel (2015) use historical data to estimate the effects of temperature on aggregate GDP. They find that temporary temperature shocks have lasting effects on income, leading them to conclude that permanent changes in temperature will affect the long run *growth rate* of income. Their paper projects that countries will experience permanent growth effects from warming and diverge perpetually in their income trajectories, with hot countries growing ever poorer and cold countries experiencing accelerating growth as they warm. However, their paper also cautions that it is difficult to discern level effects from growth effects precisely. Due to the limited number of available countries

¹Barrage and Nordhaus (2023) find a 1.6% global GDP loss from 3°C of warming, and a 3.1% loss when incorporating adjustments for possible tipping points and unmeasured non-market impacts.

²Other estimates using sector-level micro data, such as Cruz (2021) and Nath (2021), find similar magnitudes for the contemporaneous impact of temperature on GDP.

and years and the (mostly) small and transitory fluctuations in historical temperatures, it is inherently challenging to use a purely empirical approach with historical data to project the effects of large, permanent changes in future temperatures.

This paper combines a model-based interpretation of facts about economic growth with empirical estimates from a new econometric model of the dynamic effects of temperature on GDP to make a new set of long-run projections of the impact of global warming on global incomes. We start by presenting a simple, stylized model of endogenous growth across countries that we use to clarify the conditions under which changes in temperature could cause permanent changes in country-specific growth rates. In the model, technological progress can diffuse across borders such that each country's productivity is determined by a combination of domestic and international factors. Countries differ permanently in their levels of income, but they all grow at the same rate in the long-run. The speed of convergence toward these parallel growth paths — as well as the persistence of effects from a transitory shock — is determined by a parameter that governs the share of domestic growth that depends on foreign technologies. In the model, countries can follow permanently divergent growth paths only if this parameter is zero, such that domestic growth depends exclusively on domestic factors.

We present a range of evidence that global growth is tied together across countries, which suggests that country-specific shocks are unlikely to cause permanent changes in country-level growth rates. In the data, we show that countries at the frontier of global technology tend to grow at similar rates, with no discernible correlation between domestic growth and domestic innovation or investment rates. Relatedly, we find that differences in levels of income across countries persist strongly, while growth differences tend to be transitory. This is consistent with a model in which countries follow a common growth process but can vary dramatically in their levels of income. Finally, we show that the evolution of TFP in rich countries explains a meaningful, though modest, portion of TFP growth in non-OECD countries. Together, these facts point to a model in which international spillovers prevent countries from differing permanently in their growth rates as global temperatures change. This argument relates closely to that of Dell, Jones and Olken (2009), who show that growth effects of temperature can only be reconciled with the global cross-sectional gradient of income and temperature if convergence or adaptation forces prevent any growth effects from being permanent. This does not rule

out, however, that permanent temperature changes can have *persistent* effects on growth for years, or even decades, as countries transition toward a new steady-state.

To investigate whether future changes in temperature are likely to have persistent or purely temporary effects on growth, we use 1960-2019 data from a panel of countries to estimate the dynamic effects of temperature on GDP. We highlight several difficulties in specifying the econometric relationship between temperature and GDP that can account for part of the divergence in existing estimates. These concerns include omitted lags that can create the appearance of growth effects, the approach to modeling nonlinearities in temperature, and the complications introduced by the serial correlation of temperature. To confront these issues, we estimate a nonlinear state-dependent model in which temperature shocks (rather than temperature itself) are the treatment and their effects on both temperature and GDP vary with a country's average temperature. To allow even greater flexibility, we also estimate a linear model country-by-country and then assess how the estimated effects vary nonlinearly with country mean temperature as well as with other factors, such as the level of development.

We use Jordà's (2005) local projections method to estimate the dynamic response of both temperature and GDP to the temperature shock. The impulse response functions show that shocks to temperature have remarkably persistent effects on GDP, with the direction of the effect depending on a country's initial temperature. In hot countries (25°C), an unexpected 1°C increase in temperature reduces GDP by approximately one percentage point in the year of a shock. GDP remains depressed for years after the shock, with a slightly larger effect on output five years after the shock is realized. Cold countries show the opposite pattern, with unexpected increases in temperature boosting output persistently for several years, though the effects are smaller than in hot countries. Consistent with previous work, our estimates imply a bliss point of approximately 13°C, which is that experienced by the 10th percentile of the present day global population. Most of the countries that are innovation leaders have average temperatures near the bliss point, suggesting that warming is less likely to affect the pace of frontier TFP growth.

An important finding of our empirical estimates is that temperature shocks themselves show substantial persistence, which varies with a country's average temperature. When hot countries experience a shock to temperature in a given year, approximately 40% of the effect of the shock persists in the following year, and 20% remains five years later.

Thus, shocks to temperature in the historical record consist of a mixture of transitory and permanent components, even when we control for year fixed effects. The persistence is greater when we do not control for year fixed effects. This leads us to interpret the persistent effects of temperature shocks on GDP as arising from a combination of a lagged effect of the initial shock and the persistence of temperature itself. We use our simple growth model to interpret the joint persistence of GDP and temperature in the data, and show that it implies a convergence parameter consistent with permanent changes in temperature having medium-term growth effects that last for over a decade.

Finally, we use the empirical impulse response functions to project the impact of global warming on individual country economic growth through 2099. In particular, we use the ratio of the cumulative response of income to the cumulative response of temperature over a 10-year horizon to represent the long-run effect of a given increment of temperature on GDP. Importantly, we allow the effects to depend on a country's initial temperature according to the nonlinear estimates, which imply that hot countries will be harmed by warming and cold countries helped.

Our projections suggest that 3.7°C of warming could reduce global GDP by 7-12% in 2099 relative to a scenario with no warming. These damage estimates are three to five times larger than estimates that assume level effects from temperature change with no medium-run growth effects, and two to four times smaller than projections in which the growth effects are permanent. Deviations from previous projections are especially sharp in initially hot and cold countries. For Sub-Saharan Africa, our projections imply that warming reduces output by 21%. In contrast, estimates that assume only a level effect of temperature would suggest a 5% decline in this region, and those that assume a permanent growth effect would suggest an 88% reduction by 2099. Conversely, in colder Europe, our estimates suggest warming will increase GDP by about 0.6%, whereas a permanent growth-effect projection would imply a near doubling of income.

Our effort builds on several related papers that project the effects of global warming on the global economy. Dell, Jones and Olken (2012) pioneered the empirical approach of using historical data to explore the temperature-GDP relationship, and showed that temperature has persistent negative effects on output in poor countries. Burke, Hsiang and Miguel (2015) followed by highlighting the nonlinear effects of temperature, with rising temperatures benefiting colder regions and harming hotter ones. Their paper was

also the first to couple estimates from historical data with climate model forecasts of future temperature change to project the effects of global warming on country-level GDP.

More recent work employs a variety of empirical methods aimed at discerning whether a permanent increase in temperature has level or growth effects on income. Burke and Tanutama (2019) use sub-national panel data from 37 countries to show that temperature shocks have persistent effects on output. Colacito, Hoffmann and Phan (2019) find persistent effects of summer temperatures on state-level output in the U.S. Kahn, Mohaddes, Ng, Pesaran, Raissi and Yang (2021) estimate an autoregressive distributed lag model on country-level data and find that persistent absolute deviations of temperature from historical norms have negative growth effects. Bastien-Olvera, Granella and Moore (2022) use country-specific time-series regressions that filter out high-frequency variation, and find that temperature has persistent effects on output in many countries. Most recently, Bilal and Känzig (2024) use time series regressions of global GDP on global temperature to make future projections. Each of these papers concludes that their estimates are consistent with permanent effects of temperature on growth. Conversely, Newell, Prest and Sexton (2021) conduct a cross-validation exercise comparing the out-of-sample predictive power of a variety of specifications and conclude that temperature has only a level effect on income. Their preferred specification implies that a business-as-usual warming trajectory will cost 1-3% of global GDP by end-of-century, similar to the magnitude of effects in existing DICE-style damage functions.

A small number of related papers also motivate their empirical work using growth models. Kalkuhl and Wenz (2020) and Casey, Fried and Goode (2023) use closed-economy Solow and Ramsey growth models, respectively, to show that capital stock dynamics govern recovery from a transitory shock, whereas TFP drives long-run growth.³ The empirical work in these papers shows that the level of temperature does not affect growth when controlling for year-to-year changes in temperature, which they interpret to be consistent with long-run changes in temperature having level effects. These papers distinguish between exogenous and endogenous growth in TFP as the key criteria determining the long-run effects of global warming. In contrast, this paper argues that economic growth is endogenous only at the global level. In our projections, a permanent change

³Kahn, Mohaddes, Ng, Pesaran, Raissi and Yang (2019) also present a model of economic growth in which countries differ permanently in their baseline growth rates. The growth rate in their model falls with deviations from historical temperatures in either direction, but is independent of the level of temperature.

in a country's temperature has persistent medium term growth effects that eventually recede into long run level effects. Thus, we project effects of warming that are much larger than those that assume exogenous growth at the country level, and much smaller than those that assume endogenous growth at the country level.

Before proceeding, it is worth acknowledging that a number of important topics in the economics of climate change go beyond the scope of this paper. These include, but are not limited to, valuation of non-market damages (e.g. Hsiang et al., 2013; Carleton et al., 2022), non-temperature effects such as hurricanes and coastal flooding (e.g. Balboni, 2019; Desmet et al., 2021; Fried, 2022), climate tipping points (e.g. Lemoine and Traeger, 2016; Dietz, Rising, Stoerk and Wagner, 2021), valuation of uncertainty and risk aversion (e.g. Weitzman, 2009; Traeger, 2014; Barnett, Brock and Hansen, 2020), and analysis of partial and general equilibrium mechanisms for adaptation (e.g. Burke and Emerick, 2016; Cruz and Rossi-Hansberg, 2021; Moscona and Sastry, 2021; Nath, 2021; Rudik, Lyn, Tan and Ortiz-Bobea, 2021; Conte, 2022; Bilal and Rossi-Hansberg, 2023).

This paper focuses specifically on the expected impact of rising temperatures on GDP, an exercise which is itself subject to a range of important caveats. To start with, our estimates are based on the effects of changes in temperature over the range experienced by the different countries in our sample. When we make projections for the already-hot countries going forward, we are relying on temperature environments beyond anything experienced in our data. Thus, projections for the already-hot countries are tenuous. In addition, we make projections for each country in isolation based on their current and future temperatures, and the simple model we use for interpretation accounts for global interconnectedness only through the diffusion of technology. As mentioned, climate change could unleash other global effects, such as tipping points, international migration, supply chain disruptions, and shifting trade flows. To the extent that this has not yet happened in our sample, our estimates exclude those effects.

The rest of the paper proceeds as follows. Section 2 lays out a simple model of global growth along with related patterns in the data. Section 3 highlights some econometric challenges in the analysis of temperature-GDP relationships, and discusses how best to avoid them. Section 4 presents our econometric framework and estimates. Section 5 offers long run projections of the effect of global warming on GDP by country and for the world as a whole. Section 6 concludes.

2 Background on Globally-Interconnected Growth

2.1 A Stylized Model of Global Growth

As discussed in the introduction, projecting climate damages depends critically on whether a permanent change in temperature has permanent *level* effects or *growth* effects. Given the consensus that long run growth is driven by technological improvements, the key question becomes whether a permanently higher temperature will affect the level or the growth rate of technology in the long run.

To clarify the conditions for level versus growth effects of rising temperatures, we present a stylized model of country technology growth rates. We provide the full model in Appendix A, and present only key equations and intuition here. As we proceed, we have in mind that temperature could have lasting effects via the efficiency or profitability of investments in technological improvements.

In this model, country i 's income per capita can be expressed as:

$$Y_{it}/L_{it} \propto M_{it}^{\frac{1}{\sigma-1}} \cdot Q_{it}.$$

A country is richer the higher its mass of intermediate good varieties M and the higher its process efficiency Q . The number of varieties is linked to the size of the local market. The parameter $\sigma > 1$ is the elasticity of substitution between; the lower this elasticity, the greater the “love of variety” and therefore the gains from having more variety.

A country's process efficiency, in turn, evolves according to

$$Q_{it} \propto \mu_{it} \cdot (Q_{it-1})^{1-\omega} (Q_{t-1}^*)^\omega. \quad (1)$$

Here μ denotes the efficiency of technology adoption and innovation efforts in a country-year. Q^* is the process efficiency of countries at the technological frontier. The parameter $0 \leq \omega \leq 1$ governs the degree to which a country builds on technology in the frontier countries versus its own previous technology level. The process efficiency in frontier countries, meanwhile, follows

$$Q_{t+1}^* \propto \mu_t^* \cdot Q_t^*.$$

One can think of shocks to adoption efficiency μ_{it} in a follower (non-frontier) country as generating transition dynamics in process efficiency. To convey the role of ω in such transition dynamics, suppose a country is on its balanced growth path with constant μ_i , and then is hit by a temporary negative shock to its technology adoption efficiency, μ_{it} , say due to rising temperatures. Figure 1(a) illustrates that a temporary shock has a purely temporary impact on a country's technology if $\omega = 1$. That is, if a country builds solely on technology in the frontier countries, then it will quickly recover. At the other extreme, when $\omega = 0$ and a country builds only on its own technology, the level effect is permanent. In intermediate cases ($0 < \omega < 1$) temporary shocks have persistent but not permanent effects on the level of a country's technology.

Figure 1(b) displays the effect of a *permanent* negative shock to a country's adoption efficiency. When $\omega = 1$, this has a permanent level effect. When $\omega = 0$, however, there is a permanent growth effect since a country is developing its own technology in isolation and will forever make less progress. When $0 < \omega < 1$ there is a persistent growth effect that builds to a larger permanent level effect. This is because future innovators build on inferior domestic technology. But there is no long run growth effect so long as $\omega > 0$, as in the presence of cross-country knowledge spillovers countries eventually converge back to the global growth path determined by growth in frontier countries. For a non-frontier country, its own adoption efforts ultimately have level effects but not growth effects.


2.2 Evidence Consistent with Globally-Interconnected Growth

When contemplating the effect of a country's population size, human capital, climate, or other characteristic, a key question is whether one should think of its long run growth rate as connected to global knowledge spillovers ($\omega > 0$) or entirely independent ($\omega = 0$). In this section, we provide three pieces of evidence that point to interconnected growth in a general context. In the following section, we present some suggestive cross-country evidence that this interconnectedness also applies to temperature.


First, rich countries have grown at similar rates in recent decades despite very different rates of domestic innovation. The left panel of Figure 2 plots patents filed in the U.S. by origin country domestic employment in 2019 across OECD economies. Not surprisingly, OECD countries with higher employment patent proportionately more. The right panel, however, shows that larger OECD countries exhibit no faster TFP growth.

Figure 1: Impact of ω on Speed of Convergence

(a) Recovery from a Transitory Shock in Year 0




(b) Growth Path Following a Permanent Shock in Year 0



Notes: Graphs display model simulations of how the effects of shocks to a country's efficiency of technology adoption, μ , vary with the degree of international knowledge spillovers, ω . Panel (a) shows the effects of a temporary shock, and panel (b) shows the effects of a permanent shock relative to the baseline balanced growth path (gray line). $\omega = 1$ represents the case in which countries build only on global frontier technologies, and $\omega = 0$ represents the case in which each country has access to only its own technologies.

These patterns are consistent with ideas flowing across OECD economies – countries that innovate more do not seem to be growing faster.

Figure 2: Employment and Patents / Employment and Growth



Notes: Patents are from the U.S. PTO in 2019. Employment and TFP are from Penn World Table 10.0.

In addition to patenting, countries have persistently different investment rates in and levels of human and physical capital. Such differences are strongly correlated with country income *levels*, but not country income *growth rates*. Easterly, Kremer, Pritchett and Summers (1993) and Klenow and Rodriguez-Clare (2005) document the weak connection of investment rates with income growth rates, and Klenow and Rodriguez-Clare (1997), Hall and Jones (1999), Caselli (2005), and Jones (2016) document the strong connection between investment rates and income levels.

For the second piece of evidence, we test whether differences across countries in TFP levels and growth rates persist over time, using data from 1960 to 2019 from the Penn World Tables (PWT). Table 1 presents regressions of levels and growth rates of TFP on year effects and a single country fixed effect. We run these regressions one country at a time such that the null hypothesis is that each country's TFP level or growth rate is the same as the global average over the sample period. One can reject common TFP *levels* for 55 to 70% of countries, depending on the specification, but one can reject common TFP *growth rates* for only 2 to 9% of countries. Thus, the data shows that level differences across countries persist, but growth differences tend to be transitory rather than permanent. We also show this pattern visually in the top left panel of Appendix Figure A-8, which is adapted from Müller, Stock and Watson (2022). The table also shows that the regression results hold with and without PWT outlier countries, and with and without adjusting for possible variety effects linked to the scale of a country's employment.

Table 1: Tests of Country Differences in TFP Levels and Growth Rates

	(1)	(2)	(3)
<i>Dependent Variable: Log Level of TFP</i>			
Average p-value on Country FE	0.179	0.180	0.118
Percent of Countries with p-value < 0.05	54.9%	52.8%	69.7%
<i>Dependent Variable: Growth Rate of TFP</i>			
Average p-value on Country FE	0.773	0.475	0.514
Percent of Countries with p-value < 0.05	2.0%	9.0%	7.9%
Year FE	✓	✓	✓
Without Penn World Table Data Flag Countries		✓	✓
No Variety Adjustment			✓
Observations	3978	3471	3471
Countries	102	89	89

Notes: Data is over 1960 to 2019 from Penn World Table 10.0. For each country and year, we multiply the variable $rtfpna$ by the 2005 ratio of $ctfp/rtfpna$ for that country. We exponentiate the result by the inverse of the “labsh” variable to obtain TFP in labor-augmenting form. We then net out the potential contribution of variety by dividing by employment raised to the power $1/\sigma - 1$ using $\sigma = 4$. For the middle column we exclude countries the PWT flags as being outliers. For the third column we also drop the variety adjustment.

Finally, the third piece of evidence on interconnected growth comes from estimating ω using equation (1) from our simple model. Recall that ω governs the degree to which a country builds on the world frontier technology. $\omega = 1$ implies that a country builds solely on the world frontier technology whereas $\omega = 0$ implies that the country builds only on its own technology. We take logs and allow adoption efficiency μ_{it} to follow a country-specific AR(1) process. See Appendix A.3 for details on the estimation procedure.

Table 2 presents the resulting estimates. The estimates imply a high weight (above 90%) on a country’s own previous technology level, but at the same time a statistically significant weight on frontier technologies (5% to 13%). When we constrain the coefficients to sum to 1, the weights are about 92.5% and 7.5% on domestic versus foreign technologies, respectively. A value of $\omega = 0.075$ ($1 - \omega = 0.925$) implies persistent effects

on TFP levels from transitory shocks to country adoption efficiency, μ , and substantial medium run growth effects from a permanent change in μ . This is important to keep in mind when we later consider the possibility that higher temperatures in a country hinder its technology adoption and thereby have persistent effects on country TFP. The final column shows there may be some upward bias in OLS estimates of ω , as the true ω which generates an OLS ω of 0.075 in simulated data is closer to 0.07.

Table 2: Regressions of Q_{it} on Q_{it-1} and Q_{it-1}^*

	<u>Unconstrained</u>		<u>Constrained</u>		<u>Bias-Corrected ω</u> Consistent with the constraint
	Coeff. on $\ln Q_{it-1}$	Coeff. on $\ln Q_{it-1}^*$	Coeff. on $\ln Q_{it-1}$	Coeff. on $\ln Q_{it-1}^*$	
Baseline	0.931 (0.006)	0.100 (0.012)	0.925 (0.005)	0.075 (0.005)	0.071
OECD Q*	0.935 (0.007)	0.133 (0.022)	0.928 (0.006)	0.072 (0.006)	0.063
No Employment Weighting	0.923 (0.006)	0.047 (0.018)	0.926 (0.005)	0.074 (0.005)	0.061
No Variety Adjustment	0.926 (0.006)	0.081 (0.009)	0.924 (0.006)	0.076 (0.006)	0.069
With Outlier Countries	0.890 (0.007)	0.103 (0.021)	0.890 (0.007)	0.110 (0.007)	0.073

Notes: The underlying data is from Penn World Table version 10.0. The baseline row uses U.S. TFP net of a variety adjustment as a proxy for Q^* , weights countries by their employment, and excludes PWT outlier countries from the sample. The regression specification is equation (1), taking logs and allowing for μ_{it} to follow an AR(1) process with country-specific intercept, serial correlation, and innovation variance. The bias-corrected ω is the one that generates the constrained empirical OLS $\hat{\omega}$ when OLS estimation is carried out on simulated data. See Appendix A.3 for details.

In sum, we have presented three pieces of evidence in support of globally interconnected growth (i.e., $\omega > 0$). Frontier countries tend to grow together despite large differences in domestic innovation, level differences in incomes across countries tend to persist while growth differences do not, and lagged frontier country growth explains a significant, though modest, share of growth in non-frontier countries. This evidence adds to a large body of existing research finding that country growth rates are tethered together in the long run. This work includes the conditional convergence literature, which finds

that per capita incomes tend to converge to parallel growth paths but with different levels determined by fundamentals such as investment in physical capital, human capital, and technology. The development accounting literature sheds further light by decomposing the extent to which country-specific factors contribute to differences in long-run levels of development. And the technology diffusion literature presents evidence of technology flowing across countries through patents, foreign direct investment, and trade. We summarize each of these rich bodies of evidence further in Appendix C.

2.3 GDP - Temperature Correlations Consistent with Our Model

If temperature affects a country through its technology adoption efficiency — μ_{it} in our model — then temperature can affect the level of GDP per capita, but not its long-run growth. To explore whether the predictions of our theory are consistent with the data, we estimate two cross-country regressions. In the first, we regress the log level of GDP per capita in 2019 on country average temperature. In the second, we regress the average growth rate of GDP per capita from 1970 to 2019 on country average temperature.⁴

Figure 3 shows scatter plots of the data across countries along with the fitted values of the regressions. The left panel shows a clear negative relationship between the log level of per capita GDP and average temperature across countries. In the global sample of countries in 2019, each 1°C increase in long-run average temperature is associated with 8.2% lower GDP per capita with a t -statistic of 7.7.⁵ In contrast, the right panel shows little association between per capita GDP *growth* and average temperature across countries. The regression estimate indicates that each 1°C is associated with only 0.025% lower average annual growth over the 50-year period, and is not statistically different from zero.⁶ Thus, these temperature results are consistent with the more general evidence in favor of globally-interconnected growth presented in the last section.

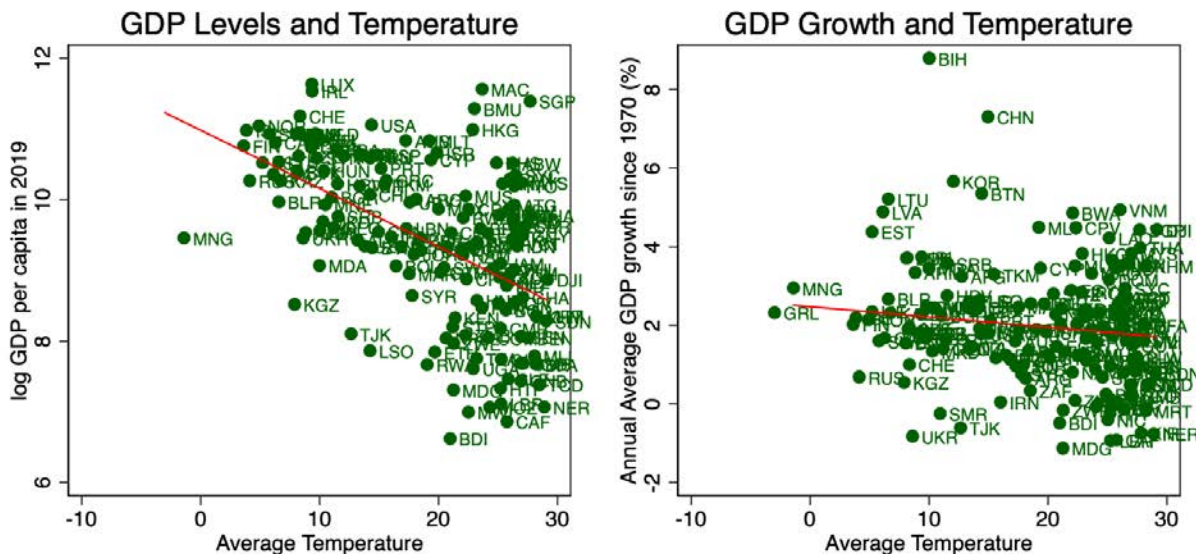
The cross-country evidence has advantages and disadvantages. As argued by Dell et al. (2009), the cross-sectional comparison between temperature and GDP levels has the

⁴The GDP levels variables are PPP variables from the Penn World Tables (PWT) data and the growth rate variables are constructed from the World Bank's World Development Indicators (WDI) data in constant local currency units. The WDI and temperature data are discussed below in Section 4.1.

⁵We exclude a small number of oil-intensive economies for this analysis that are disproportionately hot and rich. When including these countries, each 1°C is associated with 6.8% lower GDP per capita.

⁶Adding an additional control for the initial log level of GDP in 1970 approximately doubles the coefficient on temperature to 0.05% annual growth per °C, which is still small relative to the level differences and connotes a long run level effect in the conditional convergence literature *a la* Barro (1991).

Figure 3: Global Correlation With Temperature in Levels and Growth



Notes: Temperature data is from the Global Meteorological Forcing Dataset. GDP levels are from the Penn World Tables and GDP growth rates are from World Development Indicators.

advantage of capturing the long-run effects of exposure to heat on the macroeconomy, inclusive of both the accumulation of effects over time and whatever efforts have been made to adapt to them. However, cross-sectional regressions have several disadvantages. First, they are subject to omitted variable bias. For example, hotter countries often have weaker institutions and it is difficult to ascertain whether temperature caused the weaker institutions and if so whether that causal link is still relevant in the modern world. Second, the cross-country estimates do not reveal anything about the intermediate-run effects of temperature. We cannot tell whether temperature shocks have purely transitory effects on economic output or persistent impacts that create growth effects in the medium term before countries converge back to the common growth path determined by global factors. For these reasons, we turn to country panel data analysis in the next two sections.

3 Modeling of Temperature Effects on GDP in Panel Data

Using panel variation in historical data provides a potential path to improving causal identification and studying the effects of temperature on GDP at different time horizons. Estimating these effects accurately and using them to make out-of-sample projections of the effects of future global warming on future GDP, however, is also an inherently challenging exercise. As Newell, Prest and Sexton (2021) demonstrate, point estimates

can vary widely depending on the specification, and the imprecision of the estimates often makes them statistically indistinguishable. Using estimates based on variation within the historical range of temperatures to project the effects of steadily rising future temperatures adds another layer of complications.

This section outlines some of the econometric challenges associated with modeling the relationship between temperature and GDP, and demonstrates that the wide range of estimates in the literature owes partly to specifications that impose constraints that are not consistent with the data. The last part of the section introduces our econometric framework, which is designed to avoid these potential issues.

3.1 Modeling Challenges

We highlight three challenges to estimating the effect of temperature on GDP. First, we illustrate how restricting lagged effects to be zero can lead to faulty inferences about growth versus levels effects. Second, we show that the most widely-used nonlinear model of temperature effects does not identify coefficients based on within-country variation alone and does not capture the data as well as a state-dependent alternative. Third, we argue that, because temperature is serially correlated, temperature shocks rather than temperature itself should be used to estimate dynamic causal effects. Projecting the GDP effects of future climate change from historical data requires taking into account the dynamic response of both temperature and GDP following a shock.

3.1.1 Level vs. Growth Effects and the Importance of Including Lags

Previous papers such as Dell, Jones and Olken (2012), Newell, Prest and Sexton (2021), and Casey, Fried and Goode (2023) point out that models such as the Burke, Hsiang and Miguel (2015) baseline specification force temperature to have growth effects because they regress the *first difference* of log GDP on the *level* of a polynomial in temperature. Dell et al. and Casey et al. argue that one should instead include both a level and a first difference of temperature to determine whether temperature has level or growth effects. This requires no serial correlation in either GDP growth or temperature. Both GDP growth and temperature display significant serial correlation, however, so sufficient lags of *both* variables must be included to generate unbiased causal estimates of temperature

on GDP. The literature often excludes lags of GDP growth, temperature, or both.⁷

We illustrate the importance of including lags of both GDP growth and temperature using a stylized linear time series model that relates GDP growth to temperature:

$$\Delta y_{it} = \rho \Delta y_{it-1} + \beta T_{it} + \theta_1 T_{i,t-1} + \theta_2 T_{i,t-2} + \mu_i + \mu_t + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, \sigma_\eta^2)$$

$$T_{it} = \gamma T_{it-1} + \lambda_i + \lambda_t + \zeta_{it}, \quad \zeta_{it} \sim \mathcal{N}(0, \sigma_\zeta^2).$$

Δy_{it} is GDP growth in country i in year t , T_{it} is temperature, and the μ 's and λ 's represent country and year fixed effects.

If temperature has only a transitory, one-period effect on the log level of GDP it must be the case that $\theta_1 = -\beta(1 + \rho)$ and $\theta_2 = \beta\rho$. That is, the coefficients on the lagged values of temperature must reverse the previous effect on GDP growth. This is what Newell et al. (2021, p.4-5) mean by *sign reversal*. With no serial correlation of GDP growth ($\rho = 0$), temperature must enter as a first difference, i.e. $\theta_1 = -\beta$ and $\theta_2 = 0$, as argued by Dell et al. (2012) and Casey et al. (2023). However, GDP growth is serially correlated in the data, so the more general restrictions immediately above apply.

What happens if we estimate the model with the lagged temperature terms omitted, as in the baseline model on which Burke et al. (2015) (hereafter BHM) base their GDP projections? To answer this question, in Appendix B we conduct some simple Monte Carlo simulations from a model in which temperature has a temporary, contemporaneous negative effect on the level of GDP and both temperature and GDP are serially correlated. The regressions on simulated data in Appendix Table A-1 have several key takeaways. First, a specification that omits the lags of both temperature and GDP will imply that temperature has a growth effect on GDP, and not capture that this effect is temporary rather than permanent. Second, when temperature is serially correlated, omitting lagged temperature causes the coefficient on contemporaneous temperature to be biased away from its true value. Third, while controlling for lagged temperature is sufficient to eliminate the bias on the contemporaneous coefficient on temperature, it is also necessary to control for lagged GDP growth to recover unbiased estimates for lagged temperature.

⁷Notable exceptions are Kahn, Mohaddes, Ng, Pesaran, Raissi and Yang (2019), Acevedo, Mrkaic, Novta, Pugacheva and Topalova (2020), and Berg, Curtis and Mark (2023), who include lags of both GDP growth and temperature in their baseline specifications.

The algebraic example and the Monte Carlo experiment illustrate two main points. First, even without serial correlation of GDP or temperature, the BHM baseline constrains temperature to have a growth effect because it rules out reversals that turn the effect into a temporary effect on the level of GDP. Lagged values of temperature must be included in order to detect the reversal. Second, serial correlation in temperature and GDP growth imply that enough lags of both must be included to obtain unbiased estimates. How many lags one must include depends on the serial correlation properties of GDP growth and temperature and whether there are lagged effects of temperature.


Does this issue matter in actual data? Further investigation reveals that it does. To show the importance in practice, we estimate the BHM baseline model using our new data set described in the next section. The model follows BHM in regressing GDP growth on a quadratic in temperature, a quadratic in precipitation, and country and year fixed effects. It omits BHM's country-specific quadratic trends since Newell, Prest and Sexton (2021) show that they over-saturate the data.

Figure 4 shows the estimated cumulative marginal effects of temperature on GDP by temperature level when zero, one, and two lags of the polynomial in temperature are included. The specifications with temperature lags also include one lag of GDP growth. The version with no temperature lags implies that the effects of temperature on GDP growth vary with temperature itself, with positive effects for colder countries and negative effects for warmer countries. The slope of the line is statistically different from zero at many points. However, when one or two lags of temperature are included, the relationship flattens and eliminates the negative effects. One might be tempted to conclude from these estimates that the effects of temperature on GDP are transitory. However, the next sections describe additional problems with this specification.

3.1.2 Modeling Nonlinear Temperature Effects

One of the most important contributions of Burke, Hsiang and Miguel (2015) is their consideration of nonlinear effects of temperature. Citing evidence such as agricultural studies of inverse U-shape relationships between crop yields and temperature, they hypothesized that the effects of temperature on aggregate GDP are likely to be nonlinear. Their baseline model specifies a quadratic in temperature, which has also been used in many subsequent papers.

Figure 4: Estimated Cumulative Marginal Effects in a Simplified BHM Model
Effects of Adding Lags of Temperature



Notes: Estimates from regressions of GDP growth on a quadratic in temperature and precipitation, as well as country- and year-fixed effects in our new dataset. One lag of GDP growth is included in the specifications that have temperature lags. The estimates shown are for the marginal effects and are summed over current and lagged temperature. The solid dots denote estimates that are statistically different from zero at the 90 % level. GDP data are from the World Development Indicators, and temperature data is from the Global Meteorological Forcing Dataset. We describe the sample in more detail in Section 4.1.

Consider a model with BHM’s type of nonlinearity:

$$\Delta y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + X_{it} + \eta_{it}, \quad (2)$$

where Δy_{it} is the growth rate of per capita GDP in country i in year t , T_{it} is temperature in country i in year t , X_{it} is a set of control variables that include country and time fixed effects and possibly lags of variables, and η_{it} is the error term. This type of nonlinearity in a fixed effects model results in the demeaned squared variable itself being a function of the group mean. Thus, the source of identification is not strictly from “within group” variation (McIntosh and Schlenker, 2006).

We propose an alternative specification to capture nonlinearities in the way temperature affects GDP. In particular, we argue that the nonlinearity is more likely to be *state dependence*, i.e., a shock to temperature will have different effects on GDP depending on the country’s mean historical temperature.

Consider the simple case in which temperature is not serially correlated, so the shock is equivalent to the deviation from mean. We can decompose temperature in country i in year t into a country effect \bar{T}_i , a common year effect \bar{T}_t , and the shock τ_{it} . That is, $T_{it} = \bar{T}_i + \bar{T}_t + \tau_{it}$. Substituting this expression into the quadratic in temperature in (2) and combining terms that vary only by country or time with the fixed effect terms in the X_{it} 's yields the following:

$$\Delta y_{it} = \beta_1 \cdot \tau_{it} + 2\beta_2 \cdot \bar{T}_i \cdot \bar{T}_t + 2\beta_2 \cdot \bar{T}_i \cdot \tau_{it} + 2\beta_2 \cdot \bar{T}_t \cdot \tau_{it} + \beta_2 \tau_{it}^2 + X_{it} + \eta_{it}. \quad (3)$$

This decomposition shows that including temperature as a quadratic implies that the temperature shock τ_{it} enters nonlinearly in several terms: a quadratic term, interaction terms between the temperature deviation and both country and year effects, as well as an interaction between country and year effects. Moreover, there are implied parameter constraints across the various terms. Critically, the second and third terms in (3) show that the variation used to identify β_2 comes not only from within-country variation in temperature, but also from variation in the average temperature across countries, \bar{T}_i . Thus time-invariant country variables may bias the estimation of the temperature quadratic.

Our proposed alternative model contains one nonlinear term that appears in the BHM quadratic specification — the interaction of the temperature deviation τ_{it} with the country mean temperature \bar{T}_i — but it omits the other three nonlinear terms. To be specific, our state-dependent model is:

$$\Delta y_{it} = (\theta_1 + \theta_2 \cdot \bar{T}_i) \cdot \tau_{it} + X_{it} + \eta_{it} = \theta_1 \cdot \tau_{it} + \theta_2 \cdot \bar{T}_i \cdot \tau_{it} + X_{it} + \eta_{it},$$

A non-zero θ_2 allows the effect of a temperature shock on GDP to depend on a country's average temperature. To determine which model better fits the data, we estimate a model that contains both BHM's quadratic in temperature and our state-dependent term. For the reasons given in the last section, we include three lags of temperature and GDP growth as controls along with the fixed effects.

Table 3 shows the estimates using our data. Column 1 shows the BHM baseline quadratic-in-temperature specification. Both the linear and quadratic term coefficients are statistically significant and the magnitudes imply marginal effects on contemporaneous GDP growth that change from positive to negative at temperatures above 11 degrees.

Table 3: Testing the Quadratic in Temperature vs. State-Dependent Model

	Dependent Variable: GDP Growth in year t		
	(1)	(2)	(3)
Temperature $_{it}$	0.479* (0.276)	0.881** (0.282)	0.876** (0.286)
Temperature $_{i,t}^2$	-0.022** (0.007)	0.002 (0.012)	
$\tau_{it} * \overline{\text{temp}}_i$		-0.074** (0.029)	-0.069** (0.017)

Notes: τ_{it} is the temperature shock. All regressions contain country- and year- fixed effects and three lags of temperature and GDP growth. ** indicates p-value < 0.05, * indicates p-value < 0.1

Column 2 shows the estimates with both the quadratic term and our state-dependent term. The coefficient on the quadratic term falls to zero, while the state-dependent term is negative and statistically different from zero. Thus, the quadratic term is no longer informative once the state-dependent term is included. Column 3 shows the estimates for the model with just the linear term and our state-dependent term. Both are statistically different from zero. The estimates imply that the effects of temperature on current GDP switch from positive to negative for country mean temperatures of 13 and above. In sum, the data favor the state-dependent model over the quadratic-in-temperature model.⁸

3.1.3 Dynamic Causal Effects of Temperature on GDP

Finally, we discuss two issues related to dynamic treatment effects. The first issue is that most of the literature has used temperature itself as the implicit exogenous treatment. However, because temperature in each country is serially correlated, temperature itself cannot be used as the treatment. Estimation of causal effects in a dynamic context requires not only the usual conditions of instrument relevance and exogeneity, but also a third condition — lead/lag exogeneity — which requires that an instrument not be correlated with any future or past structural shock including its own leads or lags (Stock and Watson,

⁸Kahn et al. (2021) consider the absolute value of the deviation of temperature from a long-run moving average. When we add their nonlinear term to our model in Column 3, the resulting coefficient is not statistically different from zero, whereas the estimated coefficients and standard errors on the linear and state-dependent terms are similar to those in Column 3 of our table.

2018). When temperature is serially correlated, a regression of GDP growth on current temperature confounds the effects of a current shock to temperature with the effects of past shocks to temperature. This is why macroeconomic analyses routinely use *shocks* to estimate causal effects, as discussed in Ramey (2016). For this reason, we use identified temperature shocks as our treatment, where the shock is identified as the innovation to temperature in a nonlinear time series model.

The second issue is how to translate the estimated coefficients on the temperature shocks to project the effects of sustained increases in temperature on GDP. If a shock leads to a persistent change in temperature, that persistence must be accounted for when mapping estimated effects of temperature on GDP to make projections. Acevedo, Mrkaic, Novta, Pugacheva and Topalova (2020) and Newell et al. (2021) are two recent papers noting that temperatures are serially correlated. However, it is not clear whether these papers accounted for the estimated persistence of temperature when they constructed their GDP projections.

An additional complication is that temperature can have both transitory (“weather”) and permanent components (“climate change”). Even the transitory component can lead to changes in temperature that last several years, such as El Niño events. Thus, a shock to temperature can impact future GDP through both a delayed direct effect of temperature on GDP and through persistence in the temperature response itself. Decomposing the temperature shocks into transitory and permanent components is difficult in samples with a few decades of data.⁹

We use a procedure that puts more weight on the permanent component and accounts for the persistence of temperature when making projections. Specifically, we use our state-dependent local projections model to estimate impulse response functions of both temperature and log GDP to identified temperature shocks. As the horizon increases, the effects of the transitory component of temperature should die out, so that the effects of the permanent (or very persistent) component are dominant.¹⁰

To scale our estimates of the GDP effects by the temperature treatment in the historical sample, we compute the *cumulative response ratio* (CRR), defined as the integral under the

⁹Several papers have attempted to isolate the lower frequency component of temperature. Dell, Jones and Olken (2012) study the effects of changes in 15-year average temperatures. Bastien-Olvera, Granella and Moore (2022) employ time series filters to extract low frequency variation in temperature.

¹⁰See Hamilton (2018) for a discussion of this idea as a way to detrend data.

impulse response of log GDP divided by the integral of the impulse response of temperature, both up to horizon H . This cumulative response ratio is analogous to the cumulative fiscal multipliers introduced by Mountford and Uhlig (2009). In the temperature context, the CRR allows us to scale the estimated effect of a temperature shock on GDP by the cumulative change in temperature that drove the effect.¹¹ The CRR at short horizons will be dominated by the transitory component of shocks to temperature, whereas the CRR at longer horizons will be dominated by the permanent component.

3.2 Our Econometric Framework

In the previous section, we established four key points about specifying GDP-temperature models: (i) including lags of both temperature and GDP in the model helps avoid biases and faulty inference about growth effects; (ii) the state-dependent model dominates the quadratic-in-temperature for modeling nonlinear effects; (iii) coefficients on temperature shocks rather than temperature itself should be used to estimate dynamic causal effects; and (iv) cumulative response ratios (CRRs) should be used to scale the GDP effects by the cumulative changes in temperature. In this section, we incorporate these lessons into our econometric model for estimating the effects of temperature on GDP.

To make causal statements, we need to identify an exogenous shock to temperature. Our temperature shock is the innovation to temperature from a nonlinear autoregressive model of country temperature. For this shock to be valid, we require that country GDP not affect country temperature. While it is reasonable to assume that idiosyncratic changes in a country's GDP do not affect a country's temperature, there is the worry that country GDP is correlated with global GDP and that global GDP affects contemporaneous global temperatures. In one of our main specifications, we control for year fixed effects, which should lessen this concern. In specifications without fixed effects, we include current and lagged USA TFP growth (as a measure of frontier technology growth), as well as lags of global GDP. Another factor that lessens the magnitude of the potential bias is that only around 10 to 20% of the effect of a pulse of CO_2 on temperature occurs in the first year (Carleton et al., 2022).

¹¹In the applied time series literature, *cumulative* often refers to cumulative growth rates, e.g. $y_{t+H} - y_{t-1}$ as in Acevedo et al. (2020). In contrast, our measure uses cumulative *level* effects, measured as the integral under the impulse response function of levels, i.e., $\sum_{h=0}^H (y_{t+h} - y_{t-1}) = \sum_{h=0}^H y_{t+h} - (H + 1) \cdot y_{t-1}$.

To estimate the temperature shock, we project temperature in each country on its own lags, which we interact with country mean temperatures to allow the dynamics to vary across colder and hotter countries, as well as country fixed effects and either year fixed effects or global controls (global GDP growth and U.S. TFP growth). We then use the innovation in this nonlinear regression in the state-dependent regressions for temperature and GDP. In particular, we estimate the temperature shock τ_{it} as the innovation to temperature in the following:

$$T_{it} = \sum_{j=1}^p \gamma_j T_{i,t-j} + \sum_{j=1}^p \theta_j T_{i,t-j} \cdot \bar{T}_i + \mu_i + \mu_t + \tau_{it} \quad (4)$$

where T_{it} is temperature in country i in year t , \bar{T}_i is country mean temperature, μ_i are country fixed effects, μ_t are either year fixed effects or global controls, and p refers to the number of lags included. The second summation term allows the coefficients on lagged temperature to vary with country mean temperature. We include these lag interactions because we find that the dynamic response of temperature to a temperature shock is different in hot versus cold countries.¹²

We next estimate the impulse responses of temperature and GDP per capita to the estimated temperature shock from (4). To do this, we use Jordà's (2005) local projection method. This simple, intuitive method estimates the effect of a treatment in period 0 on the variable h periods after the treatment by regressing the variable at horizon $t + h$ on the shock at t , as well as lagged control variables. The coefficient on the shock at t is the estimate of the impulse response function at h . The local projection method is particularly useful in the case of nonlinear models, since obtaining impulse response functions from a nonlinear structural vector autoregression is challenging.¹³

The two sets of local projections (for temperature and GDP per capita) are as follows:

$$T_{i,t+h} = \alpha_0^h \tau_{it} + \alpha_1^h \tau_{it} \cdot \bar{T}_i + X_{it} + \zeta_{it}, \quad h = 1, \dots, H. \quad (5)$$

¹²In theory, our use of the sample average of temperature as the state variable is problematic because climate change should make temperature nonstationary. In our sample, however, the rise in temperature is small. And we obtain very similar results if we instead use average temperatures before 1980, before the temperature increases became perceptible.

¹³Most of the literature that studies state-dependence in fiscal multipliers uses local projections (e.g. Auerbach and Gorodnichenko (2013) and Owyang, Ramey and Zubairy (2013)).

where $X_{it} = \{T_{i,t-j}, T_{i,t-j} \cdot \bar{T}_i\}_{j=1}^p, \mu_i, \mu_t$.

$$y_{i,t+h} - y_{i,t-1} = \beta_0^h \tau_{it} + \beta_1^h \tau_{it} \cdot \bar{T}_i + Z_{it} + \epsilon_{it}, \quad h = 0, \dots, H. \quad (6)$$

where $Z_{it} = \{T_{i,t-j}, T_{i,t-j} \cdot \bar{T}_i, \Delta y_{i,t-j}\}_{j=1}^p, \mu_i, \mu_t$.

In the set of H regressions in (5), temperature in each year $t + h$ is regressed on the estimated temperature shock in year t , as well as controls X_{it} . The estimate of $\alpha_0^h + \alpha_1^h \cdot \bar{T}_i$ represents the impulse response function at horizon h . The second part of this expression allows the effects of the shock to vary with country mean temperature. The set of regressions starts at horizon $h=1$ because of the unit normalization, i.e., the impact effect at $h=0$ is normalized to unity in equation (4) that identifies the shock.

In the $H+1$ regressions described in (6), we regress the difference between log GDP per capita (y) at time $t + h$ and time $t - 1$ (before the shock hits) on the temperature shock in period t as well as controls Z_{it} . The impulse response of log GDP per capita at horizon h is $\beta_0^h + \beta_1^h \cdot \bar{T}_i$. Both sets of controls X_{it} and Z_{it} contain lags of temperature, lags of temperature interacted with country mean temperature, country fixed effects, and either year fixed effects or global controls (world GDP growth and U.S. TFP growth). Z_{it} additionally contains lags of country GDP growth.¹⁴

While there may be efficiency gains to estimating the regressions jointly using Seemingly Unrelated Regressions (SUR), we estimate them separately to preserve as many observations as possible. Our temperature data extend from 1950 to 2015, but our GDP data extend at most from 1960 to 2019 (since some countries enter the sample later). Each time we increase the horizon h , we lose another year of observations. Joint estimation of the regressions requires a fixed sample, so many observations would be lost.

To assess robustness, we also estimate a linear local projections model for each country separately, controlling for the global variables. This model is more flexible because it allows each country to react differently to the global controls and to have its own lag

¹⁴Including lagged GDP growth is tantamount to assuming a unit root in log GDP. In robustness checks, we specify lags in log levels and obtain similar results. We excluded precipitation variables because they were not significant and their presence did not change the estimated impulse responses.

coefficients. However, year effects are not identified, so global variable controls must be used. We then create a cross-country data set of estimated impulse responses for each horizon and regress them on country mean temperature.

4 Empirical Estimates

4.1 Data


In our main specifications, we use data from the World Bank's World Development Indicators on GDP per capita in constant local currency units. The data covers at most 1960-2019 (omitting COVID years), with early years missing for some countries. We combine these GDP data with temperature data from the Global Meteorological Forcing Dataset (GMFD) version 3, produced by researchers at Princeton (Sheffield, Goteti and Wood, 2006). GMFD is a reanalysis dataset that combines observational data with local climate models to reconstruct estimates of historical temperature at the $0.25^\circ \times 0.25^\circ$ resolution throughout the world. We calculate country-level average temperature in each year as the population-weighted average of temperature across pixels. These data are available from 1950 to 2015. Despite the temperature data extending only to 2015, we are able to use the GDP data through 2019 for estimating the response of GDP at forward horizons.

In the specifications without year fixed effects, we control for world GDP growth and a measure of frontier TFP growth. World GDP growth is in constant U.S. dollars and is from the World Development Indicators. Our measure of frontier TFP growth is based on annual utilization-adjusted U.S. TFP (Fernald, 2014).

4.2 Empirical Estimates

We estimate three versions of the model presented in Section 3.2. The first uses year fixed effects in all equations. The disadvantage of this specification is that it eliminates global warming from the data. Thus, we also estimate a specification without year effects and instead control for global economic variables in (6). The global control variables are contemporaneous U.S. TFP growth and three of its lags (as an indicator of frontier technology) and three lags of world GDP growth. The third version is the estimation of the linear model country-by-country with the second-step regression of the country-specific impulse responses on country temperature. We discuss the results of this third model as part of the robustness checks.

Figure 5: Contemporaneous Impact of a 1°C Temperature Shock on GDP Per Capita



Notes: Graph shows the initial impact of a 1°C temperature shock on log GDP estimated using the local projections specification in Equation 6. The effect is allowed to vary with long-run average historical country temperature, which is shown on the x-axis. Left panel shows estimates for the specification with year fixed effects, and right panel shows the corresponding estimates for the specification with U.S. TFP controls instead. Temperature data are from GMFD, and GDP data are from the World Development Indicators. 95% confidence interval is shown in blue. This figure shows contemporaneous effects at horizon $h = 0$, whereas Figure 6 documents the persistence of the effects.

Figure 5 displays estimates via (6) of the contemporaneous impact of an unanticipated 1°C shock to temperature in year t on log GDP per capita in year t . Temperature shocks are estimated as the residual in (4). The effects of temperature shocks on log GDP per capita are allowed to vary with a country's long-run average temperature across the historical sample, such that the effects of an unusually hot year can differ across hot and cold countries. The left panel shows the specification that controls for year fixed effects and the right panel shows the specification with global controls.

The estimates in Figure 5 reveal that temperature has meaningful contemporaneous effects on GDP per capita. The estimates using year fixed effects imply that, in the hottest countries in the world (about 28°C in the historical sample), a 1°C temperature shock reduces GDP in the same year by about 1.3%. The effects of temperature shocks are smaller in places that are less hot, and positive in very cold countries. In cold countries such as Norway, which has an historical long-run average temperature of approximately 5°C , a 1°C temperature shock increases annual output by about 0.75%. The bliss point for temperature implied by these estimates is 13.2°C , which is similar to previous estimates in the literature such as Burke, Hsiang and Miguel (2015). Our estimates are similar when year fixed effects are omitted in the right panel of Figure 5.

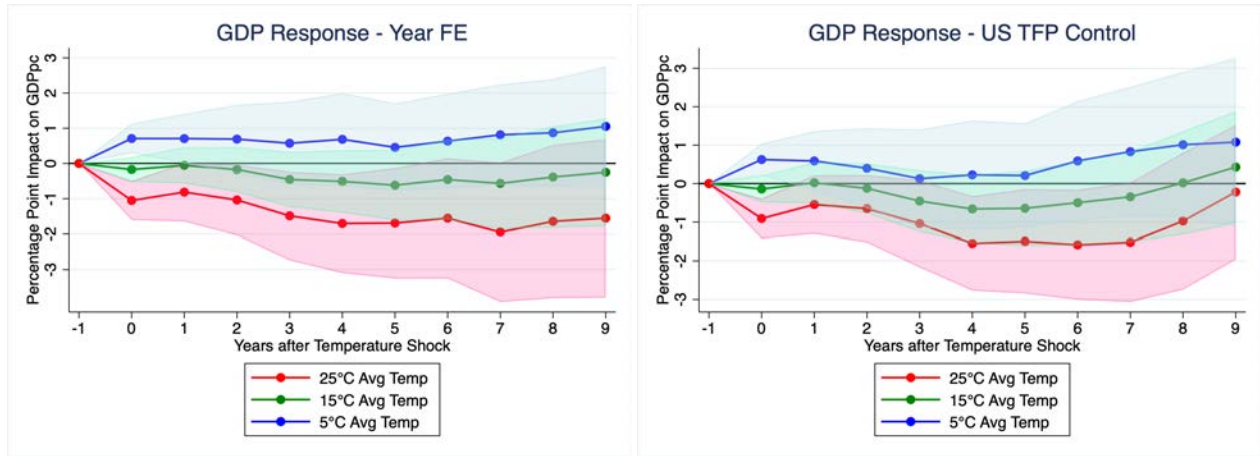
Next, we turn our attention to the *persistence* of the effects of a temperature shock on GDP per capita. Figure 6a displays the estimates from (6) of the impact of a shock in year t on GDP per capita in year $t + h$ for a 10-year horizon. The left panel shows results with year fixed effects, and the right panel with controls for U.S. TFP instead. In each graph, the three separate impulse response functions represent the effects for a range of long-run average country temperatures. The red line shows the effects on a hot country with average temperature of 25°C, such as India or Indonesia. Similarly, the green line shows the effects for a country with a moderate long-run temperature of 15°C, such as Greece or Portugal, and the blue line for a cold country with a long-run temperature of 5°C, such as Norway or Sweden.

Figure 6a demonstrates that the effects of temperature on GDP are remarkably persistent in both hot and cold countries. In the specification with year FE shown in the left panel, point estimates show no evidence that GDP per capita recovers back to trend over the 10-year horizon after the shock hits in year t . In hot countries (25°C), the initial 1°C shock in year t reduces GDP per capita by about 1.1% on impact, and remains depressed by approximately 1.7% in the fifth year following the shock. Conversely, in cold countries (5°C), GDP per capita rises by 0.7% in the year of the shock and remains 0.7% above expectations five years later. Effects continue to persist in the years that follow, though the confidence intervals unsurprisingly grow larger as more lags enter the estimate. The estimates in the right panel that control for U.S. TFP instead of year FE show similar levels of persistence in hot countries over the first seven years after a temperature shock, though with imprecise evidence that GDP per capita recovers to trend afterward.

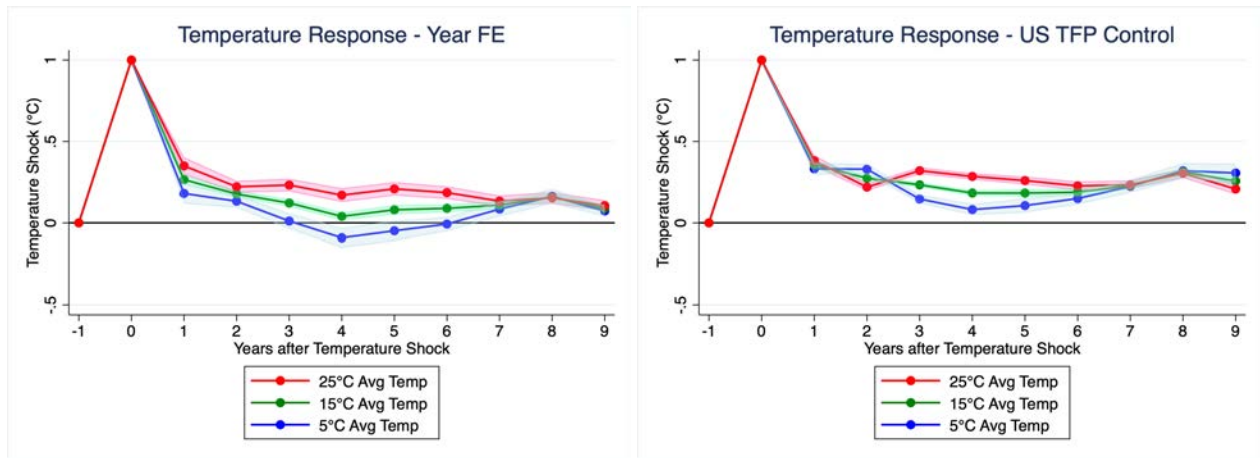
In order to interpret the persistence of the GDP impacts of the temperature shock, it is useful to measure the persistence of temperature itself. Figure 6b shows the impulse response function of temperature in the years following an initial shock estimated using the specification in (5). In the graph, we set the year 0 shock equal to 1°C by construction. Recall that the temperature shock is defined in (4) as the residual in year t from a regression of temperature on country fixed effects and its own lags. The values in each proceeding year, $t + h$, represent the proportion of the initial shock that persists in the years that follow. Like the GDP IRF, we allow the persistence of temperature to differ by a country's long-run average temperature. The red, green, and blue lines represent hot (25°C), moderate (15°C), and cold (5°C) climates, respectively.

Figure 6: Dynamic Empirical Response of Temperature and GDP per capita


(a) GDP per capita Response



(b) Temperature Response



(c) Cumulative Response Ratio



Notes: Graphs show local projections estimates of the persistent effects of an unanticipated 1°C temperature shock in year 0. Panel (a) and panel (b) show estimates for the path of GDP per capita and temperature following the shock over a 10-year horizon, estimated using Equations 6 and 5, respectively. Panel (c) shows the cumulative response ratio of the integrals of the GDP and temperature effects up to each horizon. The left graph in each panel contains the specification with year fixed effects, and the right graph contains the specification with a control for contemporaneous U.S. TFP instead. Blue, green, and red lines represent cold (5°C), moderate (15°C), and hot (25°C) countries, respectively, and the 95% confidence intervals are shown with corresponding color shading. Confidence intervals in Panel (c) are bootstrapped. Temperature data are from GMFD, and GDP data are from the World Development Indicators.

The results in Figure 6b demonstrate that temperature shocks display substantial persistence. In hot countries, the specification on the left with year fixed effects suggests that a temperature shock of 1°C in year t is followed, on average, by a temperature shock of 0.35°C in year $t + 1$, where the shock is defined in (5) as relative to the predicted value based on the information available in year t . Approximately 20% of the shock persists even in the 5th year after it is realized in hot countries, and 10% remains even in the 10th year thereafter. The results in the right panel with U.S. TFP controls show even greater persistence in temperature shocks, with over 20% of the initial shock persisting even a decade later. Thus, we conclude that temperature shocks in hotter places consist of a combination of transitory and permanent components. The specification without year fixed effects captures more of the permanent component of temperature shocks, but remarkably it appears even in the specification that removes the aggregate global trend in temperature from the estimating variation. Temperature shocks also persist temporarily in countries with moderate and cold climates, though they dissipate by the fourth or fifth year after they begin in the year FE specification.

Figure 6c brings together the dynamic GDP per capita and temperature estimates to calculate the *cumulative response ratio*, which we defined in Section 3.1.3 as the ratio of the integrals of the GDP and temperature impulse response functions at each horizon. We interpret this value as the total GDP effect of a given pulse of temperature change, accounting for both the lasting impact of the initial shock and continuing impacts caused by the persistence of the shock itself. Thus, we use the cumulative response ratio at longer time horizons as a measure of the long-run level effect of a given increment of temperature change. The results in the left panel of Figure 6c with year fixed effects suggest that in hotter (25°C) countries, each 1°C increase in temperature reduces GDP by about five percentage points in the long-run, while in colder (5°C) countries the same change would raise GDP by about five percentage points. The effects in the right panel with a U.S. TFP control are broadly similar, though somewhat smaller due to point estimates that suggest slightly more recovery of GDP and slightly more persistence of temperature. In that specification, 25°C countries lose about three percentage points of GDP from each 1°C , and 5°C countries gain about two percentage points.¹⁵

¹⁵We do not yet show confidence bands for the multipliers. Using the Ramey and Zubairy (2018) 1-step method to estimate standard errors, we obtain standard errors at the 10-year horizon of anywhere between 2 and 3.8 depending on the mean country temperature. However, the multipliers from the 1-step procedure

Appendix D presents a number of robustness checks. Appendix Figure A-1 shows the contemporaneous and persistent effects of the temperature shock on GDP using several alternative approaches. Panel (a) shows results for a specification that controls for log levels of lagged GDP rather than first differences as control variables, and panels (b) and (c) show estimates using alternative sources of temperature data from the Berkeley Earth Surface Temperature dataset and the University of Delaware climate dataset. The results are very similar throughout.

To further probe the robustness of the results, we estimate a linear local projections model of GDP per capita and temperature separately for each country, extending a method introduced by Berg et al. (2023). This procedure allows arbitrary heterogeneity of all the parameters across countries, which avoids the dynamic panel biases discussed in Pesaran and Smith (1995). In the second stage, we regress the country-specific estimates on mean country temperature, which produces estimates of impulse responses and cumulative response ratios by country mean temperature. These estimates are analogous to the results from our state-dependent panel model, only with more flexibility.¹⁶ The top panel of Appendix Figure A-2 shows the contemporaneous effects of temperature shocks across countries, and the second panel shows the average effects across the world's temperature distribution (for 5°C, 15°C, and 25°C countries). The results are extremely similar to the analogous panel specification without year fixed effects.

Appendix Figure A-3 uses the country-by-country local projections approach to shed light on two additional dimensions of robustness. The top panel shows how the average estimates differ when dropping two countries (Gabon and Kiribati) with outlier effects of greater than 10% GDP losses from temperature shocks. The effects on hot countries are slightly muted compared with the version that includes all countries. In contrast, the second panel of Figure A-3 shows that allowing for nonlinearity in the heterogeneous effects by country average temperature strongly increases the implied vulnerability of the hottest countries. Specifically, here we regress the coefficient on temperature shocks across countries on a fourth degree polynomial in long-run country average temperature. The effects on the hottest countries are nearly twice as large when allowing for

are not identical to those implied by the IRFs because 10 years of the sample must be dropped in the 1-step procedure. Nevertheless, these 1-step results are indicative of the imprecision of the estimates.

¹⁶Berg, Curtis and Mark (2023) follow the same procedure, but do not uncover a systematic relationship between temperature effects and country characteristics.

this nonlinearity. Finally, the bottom panel of Figure A-3 shows the net effects of both dropping outliers and allowing for nonlinear heterogeneity. At the 10% effect threshold for dropping outliers, the nonlinear heterogeneity effect is stronger for hot countries, but at lower outlier thresholds it is possible that these two forces would offset in magnitude.

In Figure A-4, we separately estimate the effects of temperature shocks on agricultural and non-agricultural GDP to gain additional information about mechanisms. We gather data on agricultural and non-agricultural GDP from Herrendorf, Rogerson and Valentinyi (2014), the UN National Accounts database, and the University of Groningen 10-Sector Database (Timmer, de Vries and De Vries, 2015), though the set of country-years in the resulting sample is substantially smaller than in the main analysis. Perhaps not surprisingly, we find that the effects on agricultural GDP are several times larger than on overall GDP per capita. In a country with long-run average temperature of 25°C, a 1°C temperature shock reduces agricultural GDP by about 4%, compared to about 1% for overall GDP. Somewhat more surprisingly, we find null effects of temperature on non-agricultural GDP, though the standard errors cannot rule out moderate impacts that would be consistent with the magnitudes in studies using micro-data such as Zhang, Deschenes, Meng and Zhang (2018) and Somanathan, Somanathan, Sudarshan and Tewari (2021). The large divergence between effects on agriculture and non-agriculture is also consistent with the micro-estimates in Nath (2021), though that paper also finds moderate effects on non-agricultural production.

Given the findings on agricultural and non-agricultural GDP, Table A-2 investigates the issue of state-dependence with regard to the level of development. The table shows results from regressing the coefficients for contemporaneous effects of a temperature shock from the country-by-country local projections approach on several variables. Column 1 shows the primary heterogeneity by long-run average temperature shown in Figure 6 — positive temperature shocks have more negative effects on hotter countries. In Columns 2, 3, and 4, we find little evidence that less developed or more agricultural economies are more susceptible to temperature shocks once we condition on long-run average temperature, though the estimates are imprecise. Due to the imprecision of these results, we interpret them with caution in light of both the regressions on sectoral GDP in this paper and other evidence using micro-data (Nath, 2021) that suggest that more developed and less agricultural countries are less susceptible to temperature.

4.3 Model-Based Interpretation of Empirical Results

We now bring together the estimates in Section 4.2 with the model presented in Section 2 to make progress on the key challenge in this paper, namely using historical fluctuations in temperature to draw inferences about large and permanent future warming. The projections to follow in Section 5 draw straight from the empirically estimated cumulative response ratios estimated in Figure 6c and do not rely on the model at all. But the model provides an additional lens through which to interpret the magnitude of the persistence in the estimates from historical data. Recall that, in the model, the convergence parameter ω governs both the persistence of level effects from a transitory shock and the persistence of growth effects from a permanent shock. For values of ω closer to 0, which imply weaker forces of global convergence, the level effects from a transitory shock persist longer before the economy recovers to trend, and the growth effects from a permanent shock last longer before the economy returns to the steady-state growth rate.

In order to interpret what the empirical estimates imply about long-run permanent changes in temperature, we estimate the value of ω consistent with the persistent GDP effects from the temperature shock process estimated in Section 4.2. While the historical record does not contain the ideal experiment of randomly assigned large and permanent changes in temperature, the temperature shocks we identify do contain a mixture of transitory and permanent components. The degree to which the corresponding GDP effects from these shocks persist is informative about the value of ω , which also governs the persistence of growth effects from hypothetical permanent changes in temperature when viewed through the lens of the model.

We estimate the value of ω implied by the empirical estimates as follows. We start by constructing a model simulation of a temperature shock with persistence that matches the empirical temperature IRF. In the simulation, each year's temperature shock affects that year's value of μ_{it} , which we assume remains constant in the absence of temperature shocks. Following Appendix equation (8), each year's shock to μ_{it} affects productivity and output both contemporaneously and in future periods, with the degree of persistence governed inversely by the value of ω .

We calibrate the magnitude of the temperature effect on μ_{it} to match the estimated contemporaneous impact of temperature on GDP in year 0 shown in Figure 6a for a 25°C


country. We then simulate a shock process in which we calibrate the magnitude of each period's shock to match the value from the temperature impulse response function, again for a 25°C country. The combination of the simulated temperature shock process and the calibrated magnitude of each year's temperature effect provides us with a sequence of values for μ_{it} , beginning with the initial shock in year 0. When combined with a chosen value for ω and σ , (8) implies a sequence of values for Q_{it} that result from the sequence of shocks to μ_{it} . The simulated path of Q_{it} implies a corresponding impulse response function for GDP. We set $\sigma = 4$ and search for the value of ω that minimizes the sum of squared errors between the simulated and empirical impulse response functions over the 10-year horizon.

Figure 7a displays an overlay of the empirical GDP impulse response function (in red) and its simulated counterpart (in black) for $\omega = 0.08$, the estimated value that most closely replicates the empirically estimated GDP persistence in the specification with year fixed effects. To connect the estimated ω from the impulse response function to the critical question of long-run permanent changes, Figure 7b shows the implied long-run growth path following a permanent shock to μ_{it} starting in year 0 in a simulated economy with $\omega = 0.08$. The orange line shows that the growth effects of the hypothetical permanent shock to productivity (e.g. from temperature) would persist for well over a decade, and that the eventual long run level effect would be many times larger than a level effect of the same magnitude with no persistence (the green line with $\omega = 0$). Appendix Figure A-5 shows the corresponding results for the specification with a U.S. TFP control instead of year FE, which implies an ω value of 0.21 and somewhat less persistence of growth effects. Overall, the empirical estimates imply a level of persistence that suggests that hypothetical permanent changes in temperature are likely to have growth effects that last substantially in the medium to long term, though not indefinitely.


It is worth noting that this estimate of $\omega = 0.08$ is remarkably close to the $\omega = 0.07$ estimate from Section 2.2 using indirect inference on historical growth patterns across countries. While we caution that it is possible for the persistence process of temperature shocks to differ from that of the more general drivers of growth, we take the striking similarity of these two very different methods of backing out ω as further support for a growth process in which country-specific growth effects can endure substantially (though not forever).

Figure 7: Model-Based Interpretation of Empirical Results

(a) Empirical vs. Simulated GDP Impulse Response Function
Year Fixed Effects Specification



(b) Transition Dynamics with $\omega = 0.08$



Notes: The red line in panel (a) shows the empirical impulse response function of the path of GDP following an unanticipated 1°C shock to temperature in year 0, estimated using Equation 6 with year fixed effects, with the 95% confidence interval shaded in pink. The black line shows a model simulation with $\omega = 0.08$ of the impulse response function following a shock with magnitude calibrated to match the contemporaneous effect in year 0, and persistence calibrated to match the impulse response function of temperature shown in Figure 6b. Panel (b) shows a model simulation of the medium-term growth trajectory following a permanent shock starting in year 0 with $\omega = 0.08$ in orange, and $\omega = 1$ in green for comparison.

A final point worth noting about the simulated IRFs is the critical importance of measuring the persistence in temperature itself. While the impulse response function for GDP shows no recovery during the 10-year window, it would not be correct to conclude from this that a transitory shock to temperature causes a permanent level effect on GDP since the shock to temperature is in fact not purely transitory. Thus, attributing the full path of the GDP effects to only the initial shock to temperature would overestimate the persistence of the effects. Instead, what we find through the model-based interpretation of the results is that the persistence in the GDP effects results from a combination of the lasting effects of the initial shock as well as the persistence of the temperature shock itself.

Overall, the empirically estimated IRFs are most consistent with a level of persistence in the effects of temperature that implies that global warming will have long-term, though impermanent, effects on economic growth. It is worth noting, however, that while $\omega = 0.08$ represents the best estimate to match the empirical impulse response functions, the standard errors on the 10-year horizon are large enough that these estimates are sufficient to rule out neither a substantially larger value of ω nor the edge case of $\omega = 0$. The projections in the next section demonstrate that even the seemingly small distinction between medium-run growth effects with $\omega = 0.08$ and permanent growth effects with $\omega = 0$ constitutes an enormous difference over the time scales relevant to global warming. This underscores the importance of combining the empirical estimates with the indirect inference presented in Section 2.2 to more convincingly rule out the case of permanent growth divergence across countries. Historically, growth rates did not diverge despite persistent differences in temperature.

5 Climate Change Impact Projections

5.1 Projection Approach

In this section, we use the empirical results from Section 4.2 to project the effects of global warming on the trajectory of GDP for 163 countries through the end of the 21st century. We take scientific projections of country-level population-weighted average temperature change directly from Burke, Hsiang and Miguel (2015), who use mean projected warming in RCP 8.5 across all global climate models included in the World Climate Research Programme's Coupled Model Intercomparison Project Phase 5 (CMIP5) (Tayler et al.,

2012). Following Burke, Hsiang and Miguel (2015), we use country-level projections for end-of-century warming and assume a linear increase in temperature from 2010 to 2099. The temperature projections from RCP 8.5 represent the median scenario from a warming trajectory with little emissions abatement and high fossil fuel use.

The population-weighted mean temperature increase across countries in Burke, Hsiang and Miguel (2015) is 4.3°C . This corresponds to about 3.7°C of global temperature change using the more widely cited metric of global mean surface temperature, since land warms faster than the oceans. The projections imply that warming will be spatially heterogeneous, ranging from 2.7°C to 5.8°C across countries. Thus, a country's initial temperature is not a sufficient statistic for its vulnerability to global warming, as climate models imply that some parts of the world will heat up more than others. The hottest countries in the world had population-weighted average annual temperatures of about 28.6°C (Mauritania and Niger) in the historical period from 1980-2010. By 2099, that number rises to about 33.4°C for the hottest country in the projection. Approximately 35% of the current global population lives in a country that will heat up to a level beyond the historical range of country-level temperatures in the given scenario, underscoring the importance of the caveat that projecting the effects of global warming necessarily requires out-of-sample extrapolation that is difficult to validate.

To quantify the economic effects of projected warming on country-level GDP in each future year across the century, we use the cumulative response ratios (CRRs) from Figure 6c that take the ratio of the integrals of the GDP response and temperature response to a given shock over the 10-year horizon following the initial impulse. The IRFs, shown in Figures 6a and 6b, are estimated using historical data on GDP and temperature from the 1960-2015 period as explained in Section 4. The CRRs represent the cumulative impact on GDP from the full dynamic path of a pulse to temperature, which we interpret as the long-run level effect of a given marginal change in long-run temperature. We apply the CRR separately to the change in temperature in each future year relative to the 2010 baseline, which is equivalent to treating that realization as the new permanent level of temperature. Note that the change in temperature in the BHM projection we consider is deterministic and linearly increasing throughout the century, which allows us to abstract from additional complications caused by the potential for differing temporal sequences for a given total change in future temperature.

Using the CRR to make climate change projections also requires incorporating the nonlinearity of the estimated effects. Recall from Section 4.2 that we allow the effects of temperature shocks to differ by long-run average country temperature. Specifically, the multipliers we estimate range from about a 6.0% loss per °C in the hottest parts of the world (28°C) to about a 4.8% gain per °C in the coldest parts of the world (5°C) in the specification with year fixed effects. We account for the nonlinear effects by applying the corresponding temperature-specific multiplier for each 0.1°C increment of warming that occurs in the projection. For instance, if a country warms from 25°C to 26°C, we apply the multiplier for a 25°C country to the first 0.1° of warming, the multiplier for a 25.1°C country to the next 0.1°, and so on.¹⁷ For countries that warm outside the range of historical observation, the multipliers rely on extrapolating the gradient of the GDP effects with respect to long-run average temperature beyond the range of the historical sample. For instance, at a country-level temperature of 32°C that is realized in the hottest places later in the century, our estimates imply that the long-run negative level effect on GDP per capita of an additional degree of warming is about 7.5%.

Before discussing the results, it is worth noting that these projections provide a conservative estimate of climate change impacts along one critical dimension: restricting the temperature effects to the 10-year horizon. Recall from Figure 7b that the model simulation with the estimated persistence parameter of $\omega = 0.08$ from the year FE specification implies that the growth effects from a permanent temperature change are likely to persist for well *over* a decade. However, we use lagged estimates of temperature effects from only the first decade following a shock, as estimates become excessively imprecise for longer time horizons given that there are only about five decades of available historical data. Encouragingly, the simulation with $\omega = 0.08$ does imply that the majority of the long-run level effect of a permanent change in temperature occurs within the first decade. In addition, the corresponding simulation in Appendix Figure A-5b shows that, with the $\omega = 0.21$ parameter implied by the specification with U.S. TFP controls, nearly the entire long-run effect of a permanent shock occurs within the first decade. Still, to the extent that the medium-term growth effects of temperature change last beyond a decade, our estimates will not account for the full effect.¹⁸

¹⁷Note that this requires dividing the multiplier at each temperature by 10 to convert from the effects of a 1°C change to the effects of a 0.1°C change.

¹⁸To clarify, since projected temperature is *trending* over the course of the 21st century rather than rising

5.2 Projection Results

Figure 8 displays estimates of the impact of projected warming on country level GDP by 2099, relative to a scenario with no warming. Panel (a) shows the estimates using the cumulative response ratio from the specification with year FE shown in Figure 6c. This projection allows for persistent, but not permanent, growth effects of temperature change. Panels (b) and (c) show contrasting estimates that assume level effects and permanent growth effects, respectively. The level effect projections in panel (b) use only the contemporaneous effect of temperature on GDP shown in Figure 5, rather than the full effects of the temperature pulse that accumulate over the 10-year horizon. This projection assumes that a permanent temperature change has no growth effects on GDP for any length of time, and that only contemporaneous temperature affects contemporaneous output. In contrast, the permanent growth effect projections in panel (c) use the estimates from Burke, Hsiang and Miguel (2015), and follow that paper in assuming that elevated temperature in future years permanently alters the growth rate.

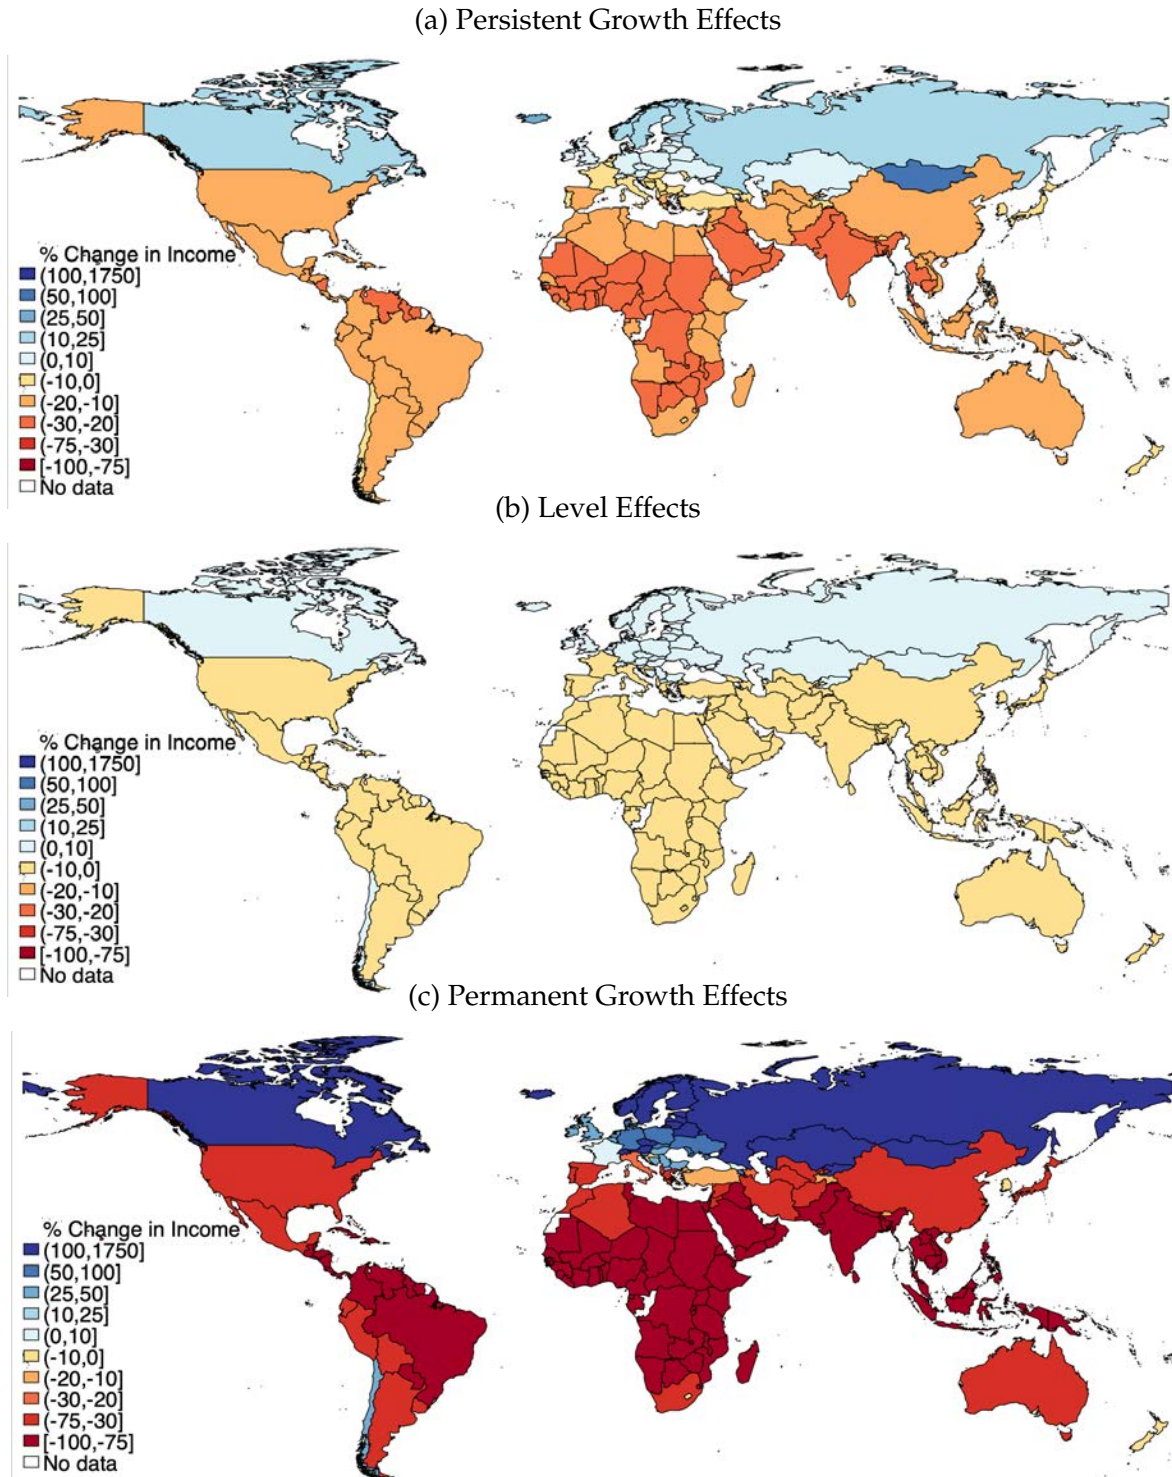
The results in Figure 8a show that the projections with persistent, but not permanent, growth effects from global warming imply large effects in absolute terms. The medium-term growth effect projections in panel (a) suggest that the hardest hit countries in the world will lose nearly 30% of their GDP to global warming on an annual basis by 2099. Warming reduces future income by at least 20% in 42 countries covering 33% of the present day global population, and by at least 15% in 93 countries covering 55% of the current global population. In total, 137 of the 163 countries representing about 92% of the existing global population expect to lose income from warming, while just under 8% of the population expects to gain. The median person in today's population distribution expects to lose about 16% of their income to warming by end-of-century.¹⁹

Comparing our estimates in Figure 8a to those in Figures 8b and 8c highlights that our estimated global warming impacts are markedly larger than the level effect estimates and smaller than the permanent growth effect estimates from previous work. The level

by a fixed permanent amount, the projected impact of global warming escalates each year and our estimates do suggest that global warming will permanently reduce global growth. The text paragraph applies to the long-run level effect on GDP that would result from a hypothetical permanent level change in temperature.

¹⁹This paragraph describes results from projections that use the empirical specification with year fixed effects. Appendix Figure A-7 and Table 4 below show the corresponding results for the specification with U.S. TFP controls instead of year fixed effects.

Figure 8: Projected Impacts of Unabated Global Warming on Country-Level GDP



Notes: Maps show the projected effects of unabated global warming on end-of-century country level GDP under different projection methods. “Persistent growth effects” estimates in panel (a) use the 10-year cumulative response ratio shown in Figure 6c, from the specification with year fixed effects, to calibrate the long-run level effect of each degree of projected warming. “Level effects” projections in panel (b) use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates.


effect projections shown in Figure 8b suggest that the hardest hit countries lose 7.3% of GDP from warming, approximately four times smaller than our persistent growth effect projections. The median person in today's global population loses only 3.1% from warming under this assumption, about five times less than when we allow persistent effects of temperature to accumulate over the 10-year horizon. The level effect projections using our estimates are very similar to the projections from Casey, Fried and Goode (2023), who project 8% losses in countries like India and a 3.4% decline in global GDP in the same emissions scenario. Conversely, the permanent growth effect projections from Burke, Hsiang and Miguel (2015) in Figure 8c suggest that the hardest hit countries are about 94% poorer than they would be in the absence of warming, as economies in the hottest places shrink dramatically. The median person in today's global population can expect to lose about 77% of their income to warming by 2099 under that set of projections.

To illustrate more concretely why the results from Figure 8 differ so sharply from each other, Figure 9 shows the projected path of income over the 21st century under each set of projections in two example countries, India and Sweden. The blue line represents the baseline trajectory of income in the absence of climate change for the given country.²⁰ The green line represents the modified trajectory using the level effect estimate in which only contemporaneous temperature affects contemporaneous GDP with no persistent effects of temperature. This estimate suggests that warming will have modest effects in both hot and cold countries. The red line represents the permanent growth effect projections in which hot and cold countries diverge permanently as the earth warms. Given that temperature is trending over the century, these projections imply accelerating growth in cold countries and decelerating growth in hot countries, which accumulates to large effects by 2099. In contrast, the orange lines show the projections that use the long-run level effects from our cumulative response ratios over the 10-year horizon. Our projections are consistent with persistent, but not permanent, growth effects from a given permanent change in temperature. Note that this actually implies permanent growth effects in our projections since the anticipated change in temperature is ever increasing rather than constant.


²⁰The figure uses baseline estimates from Scenario Two of the Shared Socioeconomic Pathway economic growth projections (Dellink, Chateau, Lanzi and Magné, 2017) commonly used in climate change economics research. Müller, Stock and Watson (2022) also provide a more comprehensive probabilistic set of projections of future baseline economic growth. Note, however, that the results in this paper are all presented in percentage changes so the baseline trajectory in the figure is used only for illustration.

Figure 9: Projected Impacts of Unabated Global Warming in Example Countries

(a) India



(b) Sweden



Notes: Graphs show the projected effects of unabated global warming on the trajectory of GDP under different projection methods for two example countries, India and Sweden. “Persistent growth effects” projections in orange use the 10-year cumulative response ratio shown in Figure 6c, from the specification with year fixed effects, to calibrate the effect of each degree of projected warming. “Level effects” projections in green use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” projections in red use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates. Corresponding projections for the specification with U.S. TFP controls instead of year fixed effects are shown in Appendix Figure A-6.

Table 4: Projected Effects of Unabated Global Warming on 2099 Income

Region	Persistent Growth Effects	Level Effects	Permanent Growth Effects
<i>Panel A - Year Fixed Effect Specification</i>			
Global GDP	-11.5	-2.2	-26.6
Global Population Average	-16.4	-3.6	-58.7
Sub-Saharan Africa	-20.6	-4.8	-86.1
Middle East & North Africa	-20.1	-4.3	-82.5
Asia	-18.0	-4.0	-73.3
South & Central America	-16.1	-3.3	-74.6
North America	-9.6	-1.4	-20.0
Europe	0.6	0.4	96.6
<i>Panel B - U.S. TFP Control Specification</i>			
Global GDP	-6.8	-1.9	-26.6
Global Population Average	-10.0	-3.1	-58.7
Sub-Saharan Africa	-13.0	-4.2	-86.1
Middle East & North Africa	-12.1	-3.7	-82.5
Asia	-11.0	-3.4	-73.3
South & Central America	-9.5	-2.8	-74.6
North America	-4.8	-1.2	-20.0
Europe	0.2	0.4	96.6

Notes: Table show the projected effects, in percent changes, of unabated global warming on end-of-century GDP under different projection methods. Country level temperature projections come from Burke, Hsiang and Miguel (2015). “Persistent growth effects” projections use the 10-year cumulative response ratio shown in Figure 6c to calibrate the effect of each degree of projected warming. “Level effects” projections use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” projections use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates. Panel (a) shows results for the local projections specification with year fixed effects, and panel (b) shows results for the specification with a contemporaneous control for U.S. TFP instead.

Table 4 summarizes the projection results from all three methods at the global and regional scale. Note that, while the country-level estimates are all expressed in percentage terms that do not depend on assumptions about baseline growth in the absence of climate change, summarizing the results at an aggregate level requires weighting countries by the size of their economies or populations. Rather than assuming that the current distribution of global GDP and population stays constant in the future, we aggregate to the global level using the average country-level baseline GDP and population projections from the five Shared Socioeconomic Pathway scenarios (Dellink et al., 2017) that forecast expected future trends under a range of assumptions about the speed of global growth and the rate of convergence in the absence of warming.

Using the weights from future baseline growth projections, Table 4 shows that our estimates imply a decline in global GDP of 11.5% in the specification with year FE, which is over five times larger than the level effect estimates and less than half as large as the permanent growth effect estimates. In the specification with U.S. TFP controls, our estimates imply a decline in global GDP of 6.8%, which is over three times larger than the level effect estimates, and nearly four times smaller than the permanent growth effect estimates. As shown in Figure 6, the implied long-run level effect of temperature is smaller in this specification, which features modestly more persistence of temperature shocks and somewhat more rapid recovery of GDP to trend.

The population-weighted average decline in income is substantially larger than the impact on world GDP, as global warming disproportionately hurts poorer regions with a larger share of the global population than of global GDP. The persistent growth effect projection implies that the average population-weighted global income loss is over 16%, which is again about four times larger than implied by a level effect projection and four times smaller than implied by a permanent growth effect projection.

Appendix Table A-3 shows the results for a moderate emissions scenario with substantial global reductions in greenhouse gas emissions and global temperature change of 2°C. In this scenario, the decline in global GDP and population-weighted incomes are 5.6% and 8.6%, respectively, in the year fixed effects specification, and 3.2% and 5.1% in the U.S. TFP control specification. The comparison with the level effect projection is similar to the high emissions scenario in Table 4, and the distinction relative to the permanent growth effect projection is even larger.

Finally, Appendix Table A-4 shows the estimated effects on the global economy of the warming that has already occurred from the 1950s to the 2010s. Depending on specification, our estimates suggest that the global economy is about \$400-\$800 billion (0.4-0.8%) smaller in the present day than it would be in the absence of the past half century of warming, with damages averaging over 5% of GDP in Africa and the Middle East. These damages are exacerbated by the fact that hotter countries have warmed at a faster rate than cooler ones on average in the historical record.

Regional comparisons of projected future climate damages reiterate that poorer and hotter regions suffer the greatest harm. Table 4 shows that the largest damages occur in Africa, the Middle East, and Asia, where lost income averages approximately 20% of GDP

in the specification with year fixed effects. Given the projected global population heavily concentrates in these regions, the persistent growth effect estimates with year fixed effects suggest that the median person in the 2099 global population suffers an 18% income loss from global warming. The corresponding loss for the median global agent in 2099 is 4% and 86% when assuming level effects and permanent growth effects, respectively.

Appendix Figure A-8 shows that our persistent growth effects projections are consistent with historical patterns of growth convergence. In contrast, the permanent growth effect projections imply that countries diverge in their income trajectories in a way that has not been previously observed, as we documented in Section 2.2. Overall, the results underscore the critical importance of allowing projected warming to have medium-term rather than permanent impacts on country-level economic growth.

In assessing the implications of these results, it is worth restating and acknowledging a number of important limitations and caveats. First, this paper focuses on the question of growth versus level effects and leaves aside a number of other relevant questions, such as the feasibility and efficacy of adaptation through channels such as trade, migration, or innovation, and the feedback between growth, emissions, and temperature change.

Second, we project effects on country levels assuming no change in the growth rate of the global frontier technology. Global warming could in principle alter the growth rate of frontier countries and hence all countries. However, the projected average effect in the most technologically innovative countries is approximately zero, depending on the exact selection of countries and weights that define the “frontier” (see, for instance, the modest positive effects of warming in most European countries). Still, we cannot rule out that the geographic distribution of global innovation will shift in such a way that temperature does slow down future frontier growth, and it is worth noting that Bilal and Känzig (2024) project large adverse effects on global growth.

Third, as acknowledged above, our projections likely somewhat understate the effects by limiting the medium growth effects of a given increment of temperature change to the 10-year horizon when the results suggest that they could persist for longer. The best estimates from the simulation in Figure 7b suggest that adjusting for this limitation would increase the preferred estimates by about 30%, which would leave unchanged the qualitative conclusion that they are both several times larger and several times smaller than the most prominent other estimates in the literature.

6 Conclusion

A critical question for assessing the potential damage from rising global temperatures is whether the result will be lower GDP per capita than otherwise or instead a lower long run growth rate of GDP per capita. Estimates in the literature vary widely on this point, from the contemporaneous level effects of Barrage and Nordhaus (2023) to the permanent growth effects of Burke, Hsiang and Miguel (2015). Their estimated losses in GDP per capita in 2099 differ by an order of magnitude as a result.

In this paper we estimate the dynamic effects of temperature on GDP and find that they build and persist, but eventually level off. Thus permanently higher temperatures appear to hurt the level of long run GDP per capita but not its long run growth rate. Compared to the literature that estimates contemporaneous level effects only, we find it is crucial to allow lagged temperature to affect future GDP per capita in a given country. In contrast to the literature that estimates permanent growth effects, we incorporate the persistence of changes in temperature and project that temperature has an effect on GDP growth for years but eventually fades.

We emphasize that level (but not growth) effects are consistent with a large literature finding that country growth rates are tethered together by technology diffusion. The estimates we obtain for the strength of knowledge spillovers based on GDP per capita alone are remarkably close to what is needed to rationalize the dynamic effects of temperature on GDP in a given country. Levels can diverge, but growth rates converge back to the rate dictated by a common technological driver. We argue that the pace of technological progress for the world is not likely to be disrupted directly by rising temperatures because most of frontier research is conducted in initially colder OECD countries.

Our estimates imply impacts in 2099 that are three to five times larger than contemporaneous level effect estimates, but one-half to one-fourth as large as estimates based on permanent growth effects, with especially stark differences for initially hot and cold countries. We leave it to future work to assess the implications of these climate damage estimates for cost-benefit analysis on climate change policy.

References

- Acemoglu, Daron, *Introduction to modern economic growth*, Princeton university press, 2008.
- and Fabrizio Zilibotti, “Productivity differences,” *The Quarterly Journal of Economics*, 2001, 116 (2), 563–606.
- , Philippe Aghion, and Fabrizio Zilibotti, “Distance to frontier, selection, and economic growth,” *Journal of the European Economic association*, 2006, 4 (1), 37–74.
- , Ufuk Akcigit, Douglas Hanley, and William Kerr, “Transition to clean technology,” *Journal of political economy*, 2016, 124 (1), 52–104.
- Acevedo, Sebastian, Mico Mrkaic, Natalija Novta, Evgenia Pugacheva, and Petia Topalova, “The effects of weather shocks on economic activity: what are the channels of impact?,” *Journal of Macroeconomics*, 2020, 65, 103207.
- Akcigit, Ufuk, Sina T Ates, and Giammario Impullitti, “Innovation and trade policy in a globalized world,” Technical Report, National Bureau of Economic Research 2018.
- Alfonso-Cristancho, Rafael, Tatiana Andia, Tatiana Barbosa, and Jonathan H Watanabe, “Definition and classification of generic drugs across the world,” *Applied health economics and health policy*, 2015, 13, 5–11.
- Alvarez, Fernando E, Francisco J Buera, and Robert E Lucas, “Idea flows, economic growth, and trade,” Technical Report, National Bureau of Economic Research 2013.
- Alviarez, Vanessa I, Javier Cravino, and Natalia Ramondo, “Firm-embedded productivity and cross-country income differences,” Technical Report, National Bureau of Economic Research 2020.
- Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple, “Innovation and production in the global economy,” *American Economic Review*, 2018, 108 (8), 2128–2173.
- Auerbach, Alan J and Yuriy Gorodnichenko, “Fiscal Multipliers in Recession and Expansion,” *Fiscal Policy after the Financial Crisis*, 2013, p. 63.

- Balboni, Clare Alexandra, "In harm's way? infrastructure investments and the persistence of coastal cities." PhD dissertation, London School of Economics and Political Science 2019.
- Barnett, Michael, William Brock, and Lars Peter Hansen, "Pricing uncertainty induced by climate change," *The Review of Financial Studies*, 2020, 33 (3), 1024–1066.
- Barrage, Lint, "Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy," *The Review of Economic Studies*, 2020, 87 (1), 1–39.
- and William D Nordhaus, "Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model," Technical Report, National Bureau of Economic Research 2023.
- Barro, Robert J, "Economic growth in a cross section of countries," *The quarterly journal of economics*, 1991, 106 (2), 407–443.
- , "J., and Xavier Sala-I-Martin," *Economic growth*, 1995, 2.
- , "Convergence and modernisation," *The economic journal*, 2015, 125 (585), 911–942.
- and Xavier Sala i Martin, "Convergence," *Journal of political Economy*, 1992, 100 (2), 223–251.
- Bastien-Olvera, BA, F Granella, and FC Moore, "Persistent effect of temperature on GDP identified from lower frequency temperature variability," *Environmental Research Letters*, 2022, 17 (8), 084038.
- Basu, Susanto and David N Weil, "Appropriate technology and growth," *The Quarterly Journal of Economics*, 1998, 113 (4), 1025–1054.
- Berg, Kimberly A, Chadwick C Curtis, and Nelson Mark, "Gdp and temperature: Evidence on cross-country response heterogeneity," Technical Report, National Bureau of Economic Research 2023.
- Bilal, Adrien and Diego R Känzig, "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature," Technical Report, National Bureau of Economic Research 2024.

- and Esteban Rossi-Hansberg, “Anticipating Climate Change Across the United States,” Technical Report, National Bureau of Economic Research 2023.
- Buera, Francisco J and Ezra Oberfield, “The global diffusion of ideas,” *Econometrica*, 2020, 88 (1), 83–114.
- Burke, Marshall and Kyle Emerick, “Adaptation to climate change: Evidence from US agriculture,” *American Economic Journal: Economic Policy*, 2016, 8 (3), 106–40.
- and Vincent Tanutama, “Climatic constraints on aggregate economic output,” Technical Report, National Bureau of Economic Research 2019.
- , Solomon M Hsiang, and Edward Miguel, “Global non-linear effect of temperature on economic production,” *Nature*, 2015, 527 (7577), 235–239.
- , W Matthew Davis, and Noah S Diffenbaugh, “Large potential reduction in economic damages under UN mitigation targets,” *Nature*, 2018, 557 (7706), 549–553.
- Cai, Jie, Nan Li, and Ana Maria Santacreu, “Knowledge diffusion, trade, and innovation across countries and sectors,” *American Economic Journal: Macroeconomics*, 2022, 14 (1), 104–45.
- Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang, “Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*,” *The Quarterly Journal of Economics*, 04 2022. qjac020.
- Caselli, Francesco, “Accounting for cross-country income differences,” *Handbook of economic growth*, 2005, 1, 679–741.
- Casey, Greg, Stephie Fried, and Ethan Goode, “Projecting the impact of rising temperatures: The role of macroeconomic dynamics,” Technical Report, Working Paper 2023.

Chaudhuri, Shubham, Pinelopi K Goldberg, and Panle Jia, "Estimating the effects of global patent protection in pharmaceuticals: a case study of quinolones in India," *American Economic Review*, 2006, 96 (5), 1477–1514.

Coe, David T, Elhanan Helpman, and Alexander W Hoffmaister, "North-south R & D spillovers," *The economic journal*, 1997, 107 (440), 134–149.

—, —, and —, "International R&D spillovers and institutions," *European Economic Review*, 2009, 53 (7), 723–741.

Colacito, Riccardo, Bridget Hoffmann, and Toan Phan, "Temperature and growth: A panel analysis of the United States," *Journal of Money, Credit and Banking*, 2019, 51 (2-3), 313–368.

Conte, Bruno, "Climate change and migration: the case of africa," 2022.

Cruz, Jose-Luis, "Global Warming and Labor Market Reallocation," Technical Report 2021.

Cruz, José Luis and Esteban Rossi-Hansberg, "The economic geography of global warming," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2021, (2021-130).

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken, "Temperature and income: reconciling new cross-sectional and panel estimates," *American Economic Review Papers and Proceedings*, 2009, 99 (2), 198–204.

—, —, and —, "Temperature shocks and economic growth: Evidence from the last half century," *American Economic Journal: Macroeconomics*, 2012, 4 (3), 66–95.

Dellink, Rob, Jean Chateau, Elisa Lanzi, and Bertrand Magné, "Long-term economic growth projections in the Shared Socioeconomic Pathways," *Global Environmental Change*, 2017, 42, 200–214.

Desmet, Klaus and Esteban Rossi-Hansberg, "On the spatial economic impact of global warming," *Journal of Urban Economics*, 2015, 88, 16–37.

- , Robert E. Kopp, Scott A. Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H. Strauss, “Evaluating the Economic Cost of Coastal Flooding,” *American Economic Journal: Macroeconomics*, April 2021, 13 (2), 444–86.
- Dietz, Simon, James Rising, Thomas Stoerk, and Gernot Wagner, “Economic impacts of tipping points in the climate system,” *Proceedings of the National Academy of Sciences*, 2021, 118 (34), e2103081118.
- Easterly, William, Michael Kremer, Lant Pritchett, and Lawrence H Summers, “Good policy or good luck?,” *Journal of monetary economics*, 1993, 32 (3), 459–483.
- Eaton, Jonathan and Samuel Kortum, “Trade in ideas Patenting and productivity in the OECD,” *Journal of international Economics*, 1996, 40 (3-4), 251–278.
- and —, “International technology diffusion: Theory and measurement,” *International Economic Review*, 1999, 40 (3), 537–570.
- and —, “Trade in capital goods,” *European Economic Review*, 2001, 45 (7), 1195–1235.
- Erosa, Andres, Tatyana Koreshkova, and Diego Restuccia, “How important is human capital? A quantitative theory assessment of world income inequality,” *The review of economic studies*, 2010, 77 (4), 1421–1449.
- Evenson, Robert E and Douglas Gollin, “Assessing the impact of the Green Revolution, 1960 to 2000,” *science*, 2003, 300 (5620), 758–762.
- Fernald, John, “A quarterly, utilization-adjusted series on total factor productivity,” in “in” Federal Reserve Bank of San Francisco 2014.
- Foster, Andrew D and Mark R Rosenzweig, “Learning by doing and learning from others: Human capital and technical change in agriculture,” *Journal of political Economy*, 1995, 103 (6), 1176–1209.
- and —, “Technical change and human-capital returns and investments: evidence from the green revolution,” *The American economic review*, 1996, pp. 931–953.
- and —, “Microeconomics of technology adoption,” *Annu. Rev. Econ.*, 2010, 2 (1), 395–424.

- Fried, Stephie, "Seawalls and stilts: A quantitative macro study of climate adaptation," *The Review of Economic Studies*, 2022, 89 (6), 3303–3344.
- Gollin, Douglas, Casper Worm Hansen, and Asger Mose Wingender, "Two blades of grass: The impact of the green revolution," *Journal of Political Economy*, 2021, 129 (8), 2344–2384.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski, "Optimal taxes on fossil fuel in general equilibrium," *Econometrica*, 2014, 82 (1), 41–88.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-run implications of investment-specific technological change," *The American economic review*, 1997, pp. 342–362.
- Grossman, Gene M and Elhanan Helpman, "Technology and trade," *Handbook of international economics*, 1995, 3, 1279–1337.
- , —, Ezra Oberfield, and Thomas Sampson, "Balanced growth despite Uzawa," *American Economic Review*, 2017, 107 (4), 1293–1312.
- Hall, Robert E and Charles I Jones, "Why do some countries produce so much more output per worker than others?," *The quarterly journal of economics*, 1999, 114 (1), 83–116.
- Hamilton, James D, "Why you should never use the Hodrick-Prescott filter," *Review of Economics and Statistics*, 2018, 100 (5), 831–843.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi, "Growth and structural transformation," *Handbook of economic growth*, 2014, 2, 855–941.
- Hsiang, Solomon M, Marshall Burke, and Edward Miguel, "Quantifying the influence of climate on human conflict," *Science*, 2013, 341 (6151), 1235367.
- Hsieh, Chang-Tai and Peter J Klenow, "Development accounting," *American Economic Journal: Macroeconomics*, 2010, 2 (1), 207–23.
- , —, and Ishan B Nath, "A global view of creative destruction," Technical Report, National Bureau of Economic Research 2022.

- Jones, Charles I, "The facts of economic growth," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 3–69.
- Jordà, Òscar, "Estimation and inference of impulse responses by local projections," *American economic review*, 2005, 95 (1), 161–182.
- Kahn, Matthew E, Kamiar Mohaddes, Ryan NC Ng, M Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang, "Long-term macroeconomic effects of climate change: A cross-country analysis," Technical Report, National Bureau of Economic Research 2019.
- , —, —, —, —, —, and —, "Long-term macroeconomic effects of climate change: A cross-country analysis," *Energy Economics*, 2021, 104, 105624.
- Kalkuhl, Matthias and Leonie Wenz, "The impact of climate conditions on economic production. Evidence from a global panel of regions," *Journal of Environmental Economics and Management*, 2020, 103, 102360.
- Keller, Wolfgang, "Geographic localization of international technology diffusion," *American Economic Review*, 2002, 92 (1), 120–142.
- , "International technology diffusion," *Journal of economic literature*, 2004, 42 (3), 752–782.
- Klenow, Peter J and Andres Rodriguez-Clare, "The neoclassical revival in growth economics: Has it gone too far?," *NBER macroeconomics annual*, 1997, 12, 73–103.
- and —, "Externalities and growth," *Handbook of economic growth*, 2005, 1, 817–861.
- Kremer, Michael, Jack Willis, and Yang You, "Converging to convergence," *NBER Macroeconomics Annual*, 2022, 36 (1), 337–412.
- Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Life cycle wage growth across countries," *Journal of Political Economy*, 2018, 126 (2), 797–849.
- Lemoine, Derek and Christian P Traeger, "Economics of tipping the climate dominoes," *Nature Climate Change*, 2016, 6 (5), 514–519.

- Lind, Nelson and Natalia Ramondo, "Global Innovation and Knowledge Diffusion," Technical Report, National Bureau of Economic Research 2022.
- and —, "Global Knowledge and Trade Flows: Theory and Measurement," Technical Report, National Bureau of Economic Research 2022.
- Lucas, Robert E, "Trade and the Diffusion of the Industrial Revolution," *American Economic Journal: Macroeconomics*, 2009, 1 (1), 1–25.
- Mankiw, N Gregory, David Romer, and David N Weil, "A contribution to the empirics of economic growth," *The quarterly journal of economics*, 1992, 107 (2), 407–437.
- Manuelli, Rodolfo E and Ananth Seshadri, "Human capital and the wealth of nations," *American economic review*, 2014, 104 (9), 2736–2762.
- McIntosh, Craig T and Wolfram Schlenker, "Identifying non-linearities in fixed effects models," *UC-San Diego Working Paper*, 2006.
- Moore, Frances C and Delavane B Diaz, "Temperature impacts on economic growth warrant stringent mitigation policy," *Nature Climate Change*, 2015, 5 (2), 127–131.
- Moscona, Jacob and Karthik Sastry, "Does directed innovation mitigate climate damage? Evidence from US agriculture," *Evidence from US Agriculture (August 2, 2021)*, 2021.
- and —, "Inappropriate Technology: Evidence from Global Agriculture," *Available at SSRN 3886019*, 2022.
- Mountford, Andrew and Harald Uhlig, "What are the effects of fiscal policy shocks?," *Journal of applied econometrics*, 2009, 24 (6), 960–992.
- Müller, Ulrich K, James H Stock, and Mark W Watson, "An econometric model of international growth dynamics for long-horizon forecasting," *Review of Economics and Statistics*, 2022, 104 (5), 857–876.
- Nath, Ishan, "Climate Change, The Food Problem, and the Challenge of Adaptation through Sectoral Reallocation," Technical Report 2021.

- Newell, Richard G, Brian C Prest, and Steven E Sexton, "The GDP-temperature relationship: implications for climate change damages," *Journal of Environmental Economics and Management*, 2021, 108, 102445.
- Nordhaus, William D, "An optimal transition path for controlling greenhouse gases," *Science*, 1992, 258 (5086), 1315–1319.
- and Andrew Moffat, "A survey of global impacts of climate change: replication, survey methods, and a statistical analysis," 2017.
- Owyang, Michael T, Valerie A Ramey, and Sarah Zubairy, "Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data," *American Economic Review*, 2013, 103 (3), 129–34.
- Parente, Stephen L and Edward C Prescott, "Barriers to technology adoption and development," *Journal of political Economy*, 1994, 102 (2), 298–321.
- and —, "A unified theory of the evolution of international income levels," in "Handbook of economic growth," Vol. 1, Elsevier, 2005, pp. 1371–1416.
- Pesaran, M Hashem and Ron Smith, "Estimating long-run relationships from dynamic heterogeneous panels," *Journal of econometrics*, 1995, 68 (1), 79–113.
- Pritchett, Lant and Lawrence H Summers, "Asiaphoria meets regression to the mean," Technical Report, National Bureau of Economic Research 2014.
- Ramey, Valerie A., "Macroeconomic Shocks and their Propagation.," in John B. Taylor and Harald Uhlig., eds., *Handbook of Macroeconomics. Vol. 2*, Amsterdam: Elsevier., 2016, pp. 71–162.
- Ramey, Valerie A and Sarah Zubairy, "Government spending multipliers in good times and in bad: evidence from US historical data," *Journal of political economy*, 2018, 126 (2), 850–901.
- Ramondo, Natalia and Andrés Rodríguez-Clare, "Trade, multinational production, and the gains from openness," *Journal of Political Economy*, 2013, 121 (2), 273–322.

- Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni, "Country-level social cost of carbon," *Nature Climate Change*, 2018, 8 (10), 895–900.
- Rudik, Ivan, Gary Lyn, Weiliang Tan, and Ariel Ortiz-Bobea, "The Economic Effects of Climate Change in Dynamic Spatial Equilibrium," 2021.
- Sachs, Jeffrey D and Andrew Warner, "Economic reform and the process of global integration," *Brookings papers on economic activity*, 1995, 1995 (1), 1–118.
- Schoellman, Todd, "Education quality and development accounting," *The Review of Economic Studies*, 2012, 79 (1), 388–417.
- Sheffield, Justin, Gopi Goteti, and Eric F Wood, "Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling," *Journal of climate*, 2006, 19 (13), 3088–3111.
- Somanathan, Eswaran, Rohini Somanathan, Anant Sudarshan, and Meenu Tewari, "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing," *Journal of Political Economy*, 2021, 129 (6), 1797–1827.
- Stock, James H and Mark W Watson, "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 2018, 128 (610), 917–948.
- Taylor, KE, RJ Stouffer, and GA Meehl, "An overview of CMIP5 and the experimental design," *Bull. Am. Meteorol. Soc*, 2012, 93, 485–498.
- Timmer, Marcel, Gaaitzen J de Vries, and Klaas De Vries, "Patterns of structural change in developing countries," in "Routledge handbook of industry and development," Routledge, 2015, pp. 65–83.
- Traeger, Christian P, "Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity," *Economic Theory*, 2014, 56 (3), 627–664.
- Weil, David N, "Health and economic growth," in "Handbook of economic growth," Vol. 2, Elsevier, 2014, pp. 623–682.

Weitzman, Martin L, "On modeling and interpreting the economics of catastrophic climate change," *The Review of Economics and Statistics*, 2009, 91 (1), 1–19.

Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang, "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," *Journal of Environmental Economics and Management*, 2018, 88, 1–17.

A A Simple Model of Global Growth

Consider N economies (countries), indexed by i , with endogenous firm entry and endogenous process innovation upon entry each period. The final goods production function is:

$$Y_{it} = \left(\int_0^{M_{it}} y_{jit}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}.$$

M_{it} is the mass of intermediate goods, indexed by j , which are available in country i , and $\sigma > 1$ is the corresponding elasticity of substitution. Intermediate goods are produced by single-product monopolistically competitive firms with the following technology:

$$y_{jit} = q_{jit} \ell_{jit}$$

where q_{jit} is process efficiency and ℓ_{jit} is production labor. Importantly, each intermediate producer lives for a single year. In each year, a new set of intermediate producers choose the process efficiency with which they enter. Entrants are subject to the following entry cost, denominated in units of labor:

$$F \cdot \exp \left(\frac{q_{jit}}{\mu_{it} \bar{Q}_{it-1}} \right)$$

where $F > 0$ and $\mu_{it} > 0$ follows a time-varying process. μ_{it} can be thought of as the efficiency of technology adoption within a given country and year. Given the exponential form, the cost of entering is convex in the level of process efficiency chosen. \bar{Q}_{it} is the geometric combination of domestic average process efficiency and that of technologically leading countries (e.g., OECD member countries), denoted as Q_t^* :

$$\bar{Q}_{it} = Q_{it}^{1-\omega} Q_t^{*\omega} \quad \text{where} \quad Q_{it} = \left(\int_0^{M_{it}} q_{jit}^{\sigma-1} dj / M_{it} \right)^{\frac{1}{\sigma-1}} \quad \text{and} \quad Q_t^* = \prod_{k \in \text{oecd}} Q_{kt}^{\alpha_k}.$$

Here $\omega \in (0,1)$ so that entrants in a country build on a combination of domestic and foreign technologies. And $\alpha_k = L_{kt} / L_t^{\text{oecd}}$ and $L_t^{\text{oecd}} = \sum_{k \in \text{oecd}}^N L_{kt}$. The higher is the combination of domestic and foreign technology (process efficiency) last year, the lower the cost of entering with a given process efficiency this year.

Labor in each country is used in production and entry:

$$\int_0^{M_{it}} \ell_{jit} dj + \int_0^{M_{it}} F \cdot \exp\left(\frac{q_{jit}}{\mu_{it} \bar{Q}_{it-1}}\right) dj = L_{it}.$$

L_{it} denotes the employment of country i in year t , which grows at the common exogenous rate n in each country:

$$L_{it} = (1 + n)L_{it-1}.$$

A.1 Equilibrium allocation

The final sector's problem delivers the usual demand functions for each variety:

$$y_{jit} = Y_{it} \left(\frac{P_{it}}{p_{jit}}\right)^\sigma \quad \text{where} \quad P_{it} \equiv \left(\int_0^{M_{it}} p_{jit}^{1-\sigma} dj\right)^{\frac{1}{1-\sigma}}.$$

Given the demand for its variety and the wage, the intermediate firm's problem delivers the usual pricing functions:

$$p_{jit} = \frac{\sigma}{\sigma - 1} \times \frac{w_{it}}{q_{jit}}.$$

Substituting this in the intermediate firm's profit function, we have:

$$\pi_{jit} = \frac{P_{it} Y_{it} q_{jit}^{\sigma-1}}{\sigma M_{it} Q_{it}^{\sigma-1}} - w_{it} F \exp\left(\frac{q_{jit}}{\mu_{it} \bar{Q}_{it-1}}\right)$$

where the free-entry condition implies:

$$\frac{P_{it} Y_{it} q_{jit}^{\sigma-1}}{\sigma M_{it} Q_{it}^{\sigma-1}} = w_{it} F \exp\left(\frac{q_{jit}}{\mu_{it} \bar{Q}_{it-1}}\right).$$

Taking the first-order condition of profits with respect to q_{jit} :

$$\mu_{it} \bar{Q}_{it-1} \times \frac{(\sigma - 1) P_{it} Y_{it} q_{jit}^{\sigma-2}}{\sigma M_{it} Q_{it}^{\sigma-1}} = w_{it} F \exp\left(\frac{q_{jit}}{\mu_{it} \bar{Q}_{it-1}}\right).$$

Substituting in the free-entry condition in this previous expression, we obtain the intermediate firm's choice of process efficiency:

$$q_{jit} = Q_{it} = (\sigma - 1)\mu_{it}\bar{Q}_{it-1} \quad \forall j.$$

Thus entrants choose higher process efficiency the higher is μ_{it} and the taller the shoulders they are building upon \bar{Q}_{it-1} . Note that, by symmetry, all j intermediate good producers choose the same process efficiency within a given country and year.

Integrating the free-entry condition over all firms and substituting in the choice of process efficiency as well as the aggregate budget constraint ($w_{it}L_{it} = P_{it}Y_{it}$), we obtain the equilibrium measure of varieties:

$$M_{it} = \frac{L_{it}}{\sigma F \exp(\sigma - 1)}.$$

This then implies that income per person is given by:

$$\frac{Y_{it}}{L_{it}} = \frac{w_{it}}{P_{it}} = (\sigma - 1)^2 \cdot \mu_{it} \cdot \bar{Q}_{it-1} \cdot \left[\frac{L_{it}}{\sigma^\sigma F \exp(\sigma - 1)} \right]^{\frac{1}{\sigma-1}}.$$

A.2 Balanced growth path

Given firm choices, the growth rate g_{it} of domestic process efficiency in country i is

$$1 + g_{it} = (\sigma - 1) \cdot \mu_{it} \cdot \left(\frac{Q_{t-1}^*}{Q_{it-1}} \right)^\omega.$$

If $\mu_{it} = \bar{\mu}_i \forall t$, including in the frontier countries, then it is easy to show that the growth rate of Q_{it} settles down to the constant growth rate of Q_t^* . That is, $g_i = g^*$.

The path of average process efficiency in country i along its balanced growth path is

$$Q_{it} = \left[\frac{(\sigma - 1) \cdot \bar{\mu}_i}{1 + g^*} \right]^{\frac{1}{\omega}} Q_t^*$$

Substituting this into the definition of Q_t^* for OECD countries, we find that:

$$1 + g^* = (\sigma - 1) \cdot \bar{\mu}^* \quad \text{where} \quad \bar{\mu}^* \equiv \prod_{k \in \text{oecd}} \bar{\mu}_k^{\alpha_k},$$

which can be substituted in the previous equation to obtain:

$$Q_{it} = (\bar{\mu}_i / \bar{\mu}^*)^{\frac{1}{\omega}} Q_t^*.$$

Note that $Q_{it}/Q_t^* \propto \bar{\mu}_i^{\frac{1}{\omega}}$ on the steady state growth path. A country's process efficiency relative to the frontier countries is increasing in its $\bar{\mu}^i$.

A country's income per capita can be expressed as:

$$Y_{it}/L_{it} = \frac{\sigma - 1}{\sigma} \cdot M_{it}^{\frac{1}{\sigma-1}} \cdot Q_{it}.$$

A country is richer the higher its mass of varieties and the higher its process efficiency. This can be translated in terms of exogenous variables as

$$Y_{it}/L_{it} = \frac{\sigma - 1}{\sigma} \cdot \left[\frac{L_{it}}{\sigma F \exp(\sigma - 1)} \right]^{\frac{1}{\sigma-1}} (\bar{\mu}_i / \bar{\mu}^*)^{\frac{1}{\omega}} Q_t^*. \quad (7)$$

Countries with more employment will generate more varieties because entry costs are denominated in terms of domestic labor.²¹ And, as mentioned, countries who are better at building on the previous year's technology (i.e., with higher $\bar{\mu}_i$) will tend to be richer.

Income per worker grows at the rate:

$$(1 + n)^{\frac{1}{\sigma-1}} (1 + g^*) - 1.$$

Using log first differences, the approximate growth rate is:

$$g_{Y/L} \approx \frac{1}{1 - \sigma} \cdot n + g^*$$

Thus all countries will grow at the same rate (in terms of both GDP and GDP per worker) in the long run if they have the same long run employment growth rate.

This model provides a stark point of contrast to "AK" models in which countries grow at permanently different rates if they have different investment rates in K and/or

²¹This is an example of a weak scale effect: the level of employment raises the *level* of income. In terms of varieties the model is in the spirit of the semi-endogenous growth models of Jones (1995) and Peretto (1998). It does not have the strong scale effect of the Romer (1990) model in which a higher level of employment raises the growth rate.

have different A levels (say due to differences in their climate). Here we could add A differences in the final goods production function and they would affect levels but not the growth rate of income per worker.

A.3 Transition Dynamics

Along a transition path, income per capita is given by:

$$Y_{it}/L_{it} = \frac{\sigma - 1}{\sigma} \cdot M_{it}^{\frac{1}{\sigma-1}} Q_{it} \quad \text{where} \quad Q_{it} = (\sigma - 1) \cdot \mu_{it} \cdot Q_{it-1}^{1-\omega} Q_{t-1}^{*\omega}.$$

So the transition dynamics for average process efficiency for non-OECD country i and for the OECD countries, respectively, is:

$$Q_{it+1} = (\sigma - 1) \cdot \mu_{it} \cdot Q_{it}^{1-\omega} Q_t^{*\omega} \quad \text{and} \quad Q_{t+1}^* = (\sigma - 1) \cdot \mu_t^* \cdot Q_t^*. \quad (8)$$

To characterize the speed at which countries converge to the common stationary growth rate g^* , once their μ_i settles down, one needs an estimate for ω . So suppose that we can proxy process efficiency Q_{it} by a country i 's TFP (in labor augmenting form) net of its "love of variety" component M_{it} (which is proportional to employment in country i). Then one could estimate equation (8) in logarithms by OLS with country fixed effects β_i :

$$\log(Q_{it+1}) = \beta_i + (1 - \omega) \log(Q_{it}) + \omega \log(Q_t^*) + u_{it} \quad (9)$$

However, the serial correlation coming from the unobserved μ_{it} could potentially bias an OLS estimate of ω . Therefore, we instead estimate ω by indirect inference. More precisely, we proceed in 6 steps:

1. We first obtain the biased OLS estimate $\hat{\omega}_{\text{empirical}}$ by estimating equation (9).
2. Then, given a value of σ , we choose a value of ω_0 and use it together with data on Q_{it} and equation (8) to obtain country-specific time series for μ_{it} .
3. With these μ_{it} series we estimate the AR(1) parameters μ_i , ρ_i and ζ_i for each country separately by OLS.
4. We draw shocks ϵ_{it} to μ_{it} from the normal distribution $\mathcal{N}(0, \zeta_i)$ to simulate the μ_{it} process for T periods (matching the length of our time series for Q_{it}), starting the

simulation with a random draw from the stationary distribution of μ_{it} : $\log(\mu_{i0}) \sim \mathcal{N}\left(\frac{\log(\mu_i)}{1-\rho_i}, \frac{\xi_i^2}{1-\rho_i^2}\right)$.

5. With the simulated time series of μ_{it} , we use equation (8) together with the empirical starting value Q_{i0} and our chosen value ω_0 to simulate the path of Q_{it} for T periods.
6. Finally, we estimate equation (9) with simulated data and compare the simulated and empirical estimates $\hat{\omega}_{\text{simulation}}$ and $\hat{\omega}_{\text{empirical}}$. To elicit the true value of ω , we iterate on our initial chosen value of ω_0 until the distance between $\hat{\omega}_{\text{simulation}}$ and $\hat{\omega}_{\text{empirical}}$ goes to zero (within a tolerance).

Assuming that Q_t^* is U.S. TFP in year t , we estimate ω according to this algorithm. We restrict our sample to countries with complete data between 1980 and 2019 and for which data quality is not an issue.²² Overall, we are left with a balanced panel of 103 countries. Finally, when estimating equation (9), we (a) use weights that correspond to each country's global employment share in a given year, (b) apply the Cochrane-Orcutt estimation procedure to adjust for serial correlation, and (c) either do or do not impose the constraint that the exponents on own and foreign technologies add up to 1.

With this strategy, the biased OLS estimate $\hat{\omega}$ we obtain is equal to 0.076 (0.006). And we find that this is generated by a *true* ω of 0.069. Thus, at least in our simulation, the bias is small and the OLS ω is not far from the true ω . The true ω of 0.07, combined with $\rho > 0$, implies that shocks to country technology adoption will have effects on GDP that will build and persist for a number of years before fading.

As a validation exercise, we use the simulated data produced in step 4 of the algorithm to calculate two cross-sectional moments (across 103 countries): (A) the standard deviation of average annual TFP growth and (B) the correlation of the logarithm of TFP between the beginning and ending periods of our simulation. Those moments are respectively equal to 1.95% and 0.898 when calculated on simulated data. If we instead compute these moments using real world data, we get values of 1.89% and 0.707, respectively.

²²The Penn World Tables classifies some countries as "outliers" because their data is of poor quality in some year. We exclude those countries from our sample, in addition to five other countries for which data quality is also an issue. The five other countries are Kuwait, the Central African Republic, Angola, Mongolia and Qatar.

B Monte Carlo Evidence on Growth vs. Level Effects

This section reports the details and results of the Monte Carlo investigation of biases in estimating levels versus growth effects. Recall the equations from the main text:

$$\Delta y_{it} = \rho \Delta y_{it-1} + \beta T_{it} + \theta_1 T_{i,t-1} + \theta_2 T_{i,t-2} + \mu_i + \mu_t + \eta_{it}.$$

$$T_{it} = \gamma T_{it-1} + \lambda_i + \lambda_t + \zeta_{it}, \quad \zeta_{it} \sim \mathcal{N}(0, \sigma^2).$$

Δy_{it} is GDP growth (based on log differences of GDP and stated in percent) in country i in year t , T_{it} is temperature in country i in year t , and the μ 's and λ 's represent country and year fixed effects. We are implicitly assuming that the log level of GDP is driven by a unit root permanent component as well as a component that is related to temperature.²³

Simple algebra shows that if temperature has only a transitory, one-period effect on the log level of GDP it must be the case that $\theta_1 = -\beta(1 + \rho)$ and $\theta_2 = \beta\rho$. That is, the coefficients on the lagged values of temperature in the GDP growth equation must reverse the previous effect on GDP growth. This algebra clarifies what Newell et al. (2021) mean by *sign reversal* when discussing their estimates that include lags of temperature (e.g. p. 4-5). In the special case in which there is no serial correlation of GDP growth ($\rho = 0$), temperature must enter as a first difference, i.e. $\theta_1 = -\beta$ and $\theta_2 = 0$. However, GDP growth is serially correlated in the data, so a more general formula is needed.

What happens if one estimates the model with the lagged temperature terms omitted, as in the baseline model on which BHM base their projections?

To answer this question, we conduct some simple Monte Carlo simulations. We create a panel of 150 countries, each with 60 years of data. We calibrate the model so that temperature has a temporary, contemporaneous effect on the level of GDP. We set β to -1 and the autocorrelation parameter for GDP growth, ρ , to 0.2 based on regressions on our data set.²⁴

²³The nonstationarity of GDP does not imply that all shocks to GDP have permanent level effects. GDP is likely affected by both permanent and temporary driving forces. For example, a permanent change in technology likely leads to a permanent change in GDP and its gradual diffusion could lead to serial correlation in GDP growth rates. Monetary policy shocks are examples of driving forces that have temporary effects.

²⁴Note that we measure GDP growth as a percent, so our coefficients on temperature are typically 100

Table A-1 shows the results of estimating several specifications on the simulated data. We begin by considering the case in which $\gamma=0$, so that there is no serial correlation in temperature. The first column shows the result of estimating the BHM model with no lags. Interestingly, even when no lags of temperature (or GDP growth) are included, the estimates of β are centered around the true value of -1. There is no bias in this case because the omitted lagged temperature variables are uncorrelated with current temperature since deviations from mean are i.i.d. However, this contemporaneous estimated growth effect tells us *nothing* about how GDP growth will respond in the future, so one cannot infer permanent growth effects. In fact, the temperature results completely reverse if we include lagged temperature in the regression. Column 2 shows that the parameter on the first lag of temperature reverses the effect of contemporaneous temperature, such that the sum of the parameters on the three temperature variables is zero, implying no lasting effects. However, Column 2, which does not control for lagged GDP growth, recovers biased coefficients for the first and second lags of temperature. Column 3 shows that only with an additional control for lagged GDP growth can the regression recover estimates close to the true coefficients for both contemporaneous and lagged temperature.²⁵

Columns 4 through 6 of Table A-1 estimate the same regressions as Columns 1 through 3, but on simulated data in which temperature follows a first-order autoregressive process (AR(1)). The primary difference here is that the regression of GDP growth on contemporaneous temperature with no lags included is downward biased by 50 percent (Column 4). The bias occurs in this case because the omitted lags of temperature are correlated with contemporaneous temperature. Once the two temperature lags and the one lag of GDP growth are included, as in Column 6, the coefficient on temperature is unbiased, as are the coefficients on lagged temperature.

This Monte Carlo experiment illustrates two main points. First, even without serial correlation of GDP or temperature, the BHM baseline specification constrains temperature to have a growth effect because it rules out reversals that turn the effect into a temporary effect on GDP levels. Lagged values of temperature must be included in order to detect the reversal effect. Second, the presence of serial correlation of GDP growth

times those of most others in the literature.

²⁵The estimate of ρ on lagged GDP growth displays the well-known downward bias of autoregressive parameters in finite samples. The bias is approximately $-(1 + 3\rho)/(\# \text{ of observations in the time dimension})$. Our simulations have 60 years for each country, so the bias is predicted to be 0.027.

Table A-1: Monte Carlo Illustration of Bias from Omitting Temperature Lags

	Dependent Variable: GDP Growth in year t					
	(1)	(2)	(3)	(4)	(5)	(6)
	$\gamma = 0$	$\gamma = 0$	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 0.5$	$\gamma = 0.5$
Temperature $_{it}$	-0.995 (0.128)	-0.995 (0.130)	-0.996 (0.128)	-0.517 (0.113)	-0.995 (0.131)	-0.996 (0.128)
Temperature $_{i,t-1}$		1.003 (0.130)	1.177 (0.129)		1.00 (0.145)	1.175 (0.143)
Temperature $_{i,t-2}$		-0.0028 (0.132)	-0.182 (0.130)		-0.0029 (0.132)	-0.180 (0.130)
GDP Growth $_{i,t-1}$			0.179 (0.011)			0.179 (0.011)

Notes: Simulated data for 150 countries with 60 years of data each. The true parameter on contemporaneous temperature, β , is -1. The true parameters on the two lags of temperature are 1.2 and -0.2, respectively. γ is the autocorrelation coefficient on temperature and varies across specifications. The true parameter on lagged GDP growth is 0.2. Standard errors in parentheses. The downward bias in the estimate of this latter parameter is well-known for finite samples.

and temperature implies that simple first-difference versus levels specifications are not appropriate, so more lags are likely to be necessary. How many lags should be included depends on the serial correlation properties of GDP growth and temperature and whether there are lagged effects of temperature.

C Literature on Globally-Interconnected Growth

The evidence we present in Section 2 of this paper adds to an already-established body of evidence that has led to a consensus among growth economists that country growth rates are tethered together in the long run (i.e., $\omega > 0$). In this section, we summarize the three strands of literature that underlie this consensus.

Conditional convergence

The consistent finding in the cross-country growth regression literature is that per capita incomes tend to converge to parallel growth paths (or sometimes even the same growth path). That is, countries converge towards relative steady state income levels determined by persistent fundamentals affecting their long run investment rates in physical capital, human capital, and technology. Classic cites in this regard include Barro and Sala-i Martin (1992), Mankiw, Romer and Weil (1992), and Easterly, Kremer, Pritchett and Summers (1993). Their findings hold up in more recent studies such as Pritchett and Summers (2014), Barro (2015), and Kremer, Willis and You (2022).

The dominant explanation for this pattern is that technology diffuses across countries, so that countries experience the same long run growth rate if they are sufficiently open to the international flow of ideas. This view is advocated by Mankiw et al. (1992), Barro (1995), Parente and Prescott (1994, 2005), Grossman and Helpman (1995), Sachs and Warner (1995), Klenow and Rodriguez-Clare (2005), Acemoglu, Aghion and Zilibotti (2006), Acemoglu (2008), Lucas (2009), Alvarez, Buera and Lucas (2013), Buera and Oberfield (2020), Cai, Li and Santacreu (2022), Lind and Ramondo (2022b), Hsieh, Klenow and Nath (2022), and many others.

Development accounting

A large literature estimates *level* effects of country differences in investment rates in human and physical capital. That is, such differences help account for differences in levels of development rather than generating persistent differences in country growth rates.

One of the first and most influential in this vein was Mankiw et al. (1992). Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999) homed in on how schooling contributed to income differences. Erosa et al. (2010), Schoellman (2012), and Manuelli and Seshadri (2014) emphasized differences in the quality of schooling across countries. Weil

(2014) examined the role of health differences, and Lagakos et al. (2018) human capital accumulated on the job.

Caselli (2005), Hsieh and Klenow (2010), and Jones (2016) provide surveys of this literature. Again, these studies provide evidence that investment rate differences have level effects on country incomes, rather than causing country growth rates to diverge.

Technology diffusion

Many studies provide direct or at least indirect evidence of technology diffusing across countries. The evidence covers categories like patents, trade, foreign direct investment (FDI), hybrid seeds, and generic drugs:

Eaton and Kortum (1996, 1999) show that firms frequently patented the same invention in many different OECD countries at once in the era before the European Patent Office. Patenting is costly, so this indicates that firms routinely tried to protect their intellectual property from being used by competitors selling in foreign markets. More recently, Jones (2016) stresses that over half of patents in the United States are filed by companies and individuals based outside the U.S. Akcigit, Ates and Impullitti (2018) use this data to estimate the joint contribution of research in the U.S. and Europe to their common growth rate.

Eaton and Kortum (2001) document that all but a few countries import most of their equipment from other countries. Since Greenwood et al. (1997) much of U.S. growth has been traced to equipment-embodied technical change. Grossman, Helpman, Oberfield and Sampson (2017) is a recent paper in the same spirit. Coe, Helpman and Hoffmaister (1997, 2009) find that importing goods from R&D-intensive economies is associated with higher productivity, consistent with technology embodied in rich-country exports. See also Keller (2002), and Keller (2004) for a survey.

Firms can also transfer technology through FDI, i.e., operating plants in other countries. Natalia Ramondo provides some of the best evidence in a series of papers with collaborators: Ramondo and Rodríguez-Clare (2013), Arkolakis, Ramondo, Rodríguez-Clare and Yeaple (2018), Alviarez, Cravino and Ramondo (2020), and Lind and Ramondo (2022a,b).

The use of hybrid seeds, with substantial impact on agricultural productivity, can be traced directly to foreign genetic ancestors in many countries. Foster and Rosenzweig

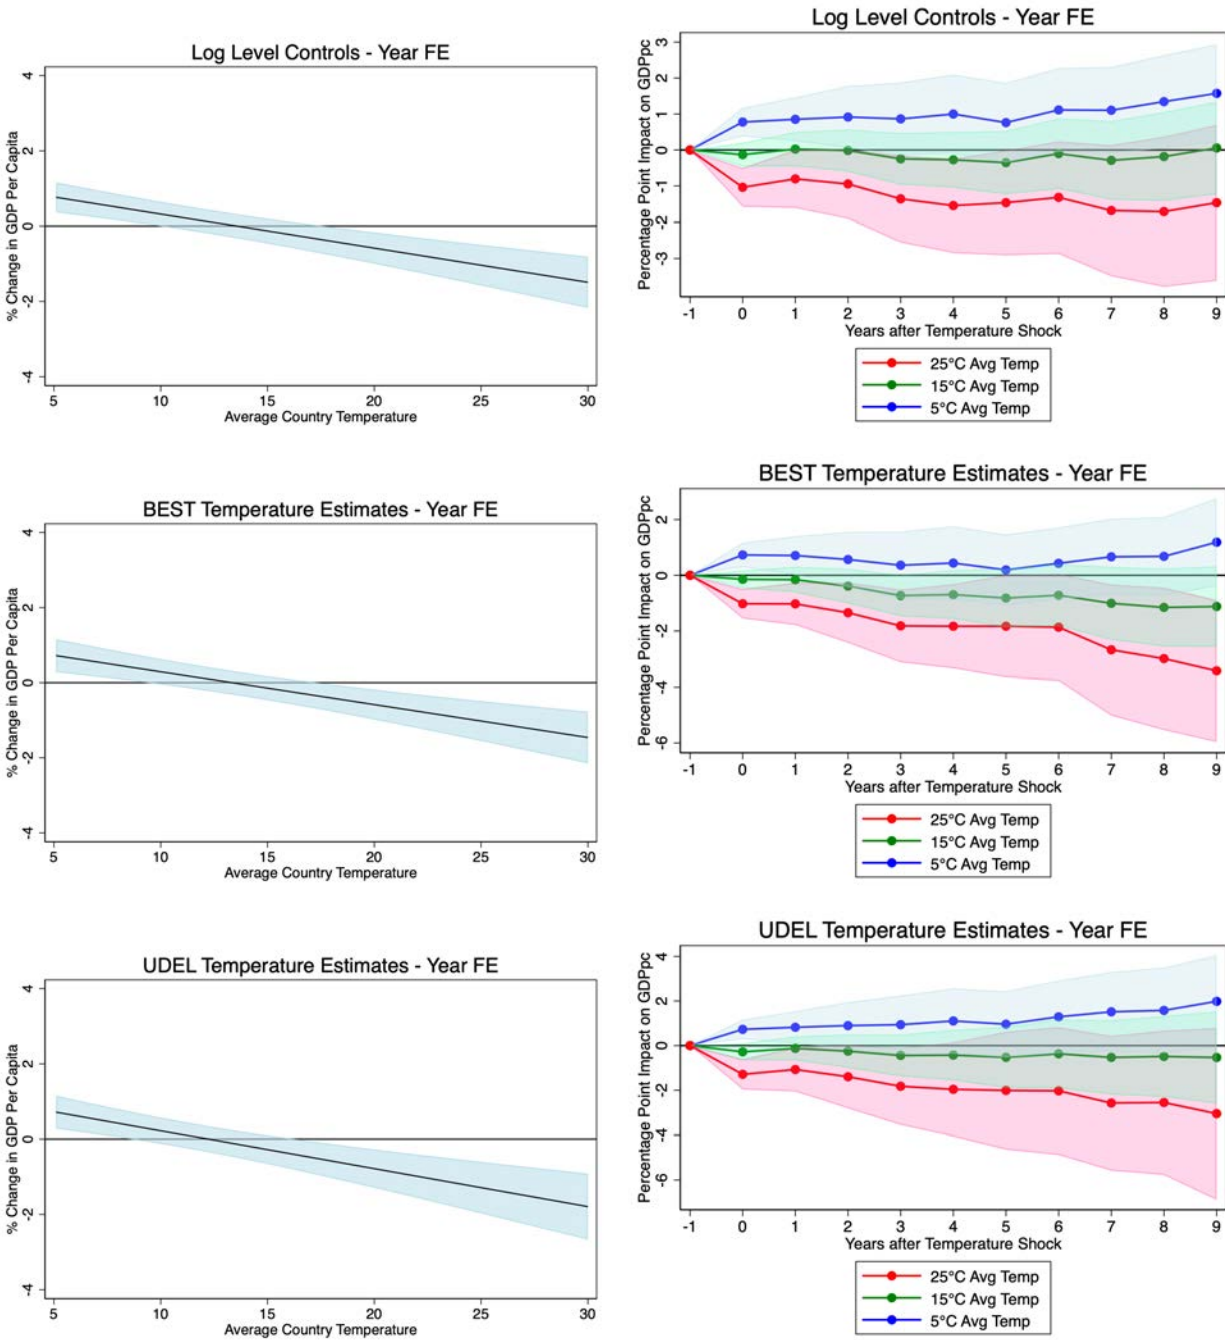
(1995, 1996) study India in particular, and they provide a survey in Foster and Rosenzweig (2010). Evenson and Gollin (2003) and Gollin, Hansen and Wingender (2021) provide evidence for many countries.

Alfonso-Cristancho et al. (2015) compile statistics on generic drug production across the world. The World Trade Organization Trade-Related Aspects of Intellectual Property Rights (TRIPS) agreement aimed to deal with generic drugs and other flows of intellectual property. See Chaudhuri, Goldberg and Jia (2006) for how TRIPS impacted the generic drug industry in India.

Some papers analyze ways in which technology developed in advanced economies may not be appropriate for emerging economies. Still, they obtain that a fraction of technologies flow, resulting in level differences rather than growth rate differences across countries. Examples include Basu and Weil (1998), Acemoglu and Zilibotti (2001), Alviarez et al. (2020), and Moscona and Sastry (2022).

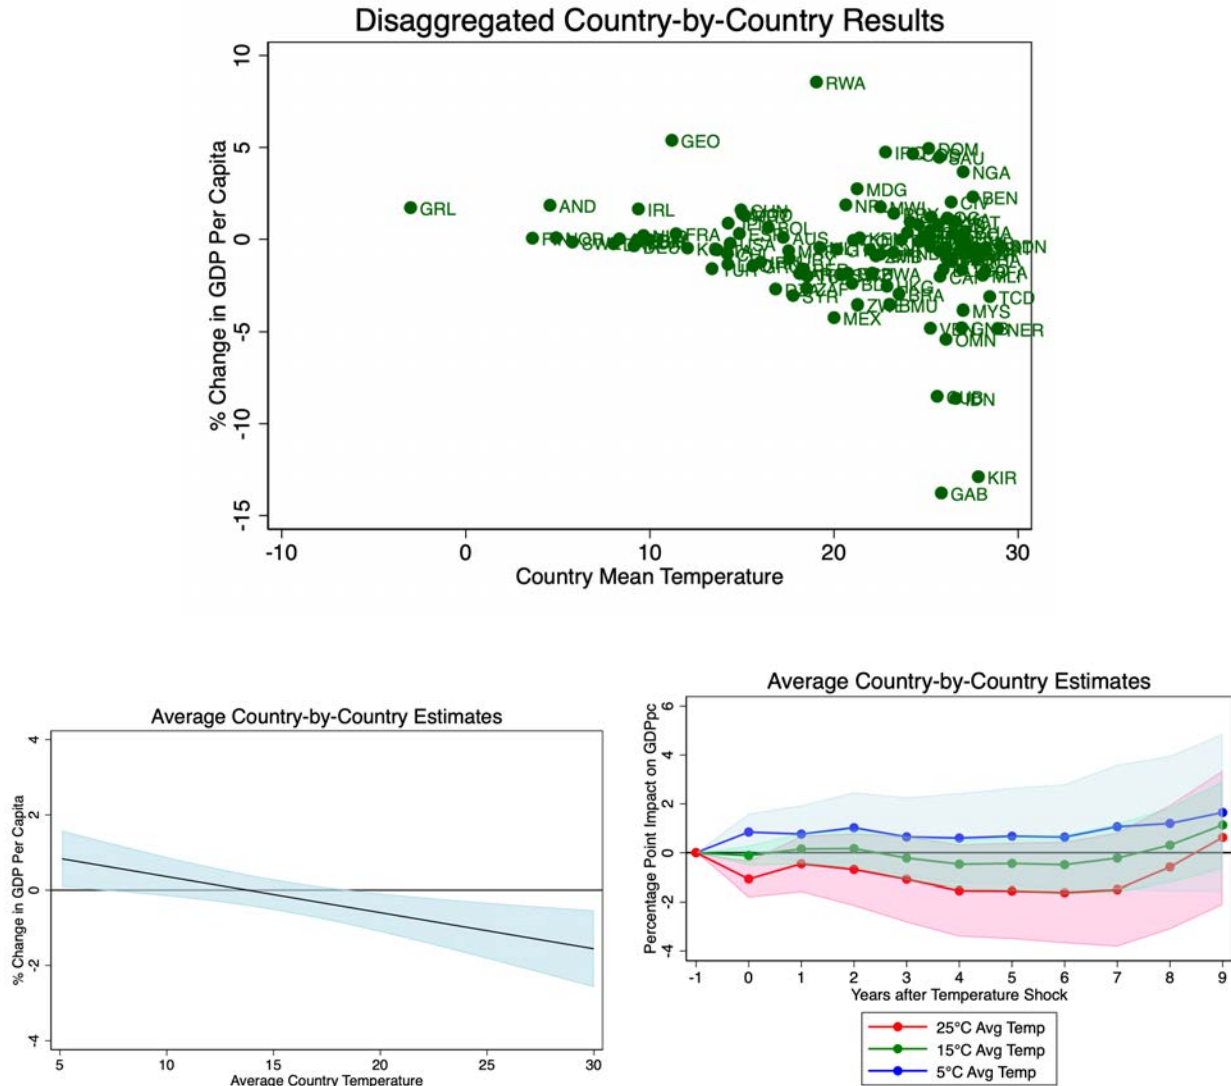
D Empirical and Projection Robustness Results

Figure A-1: Robustness Results



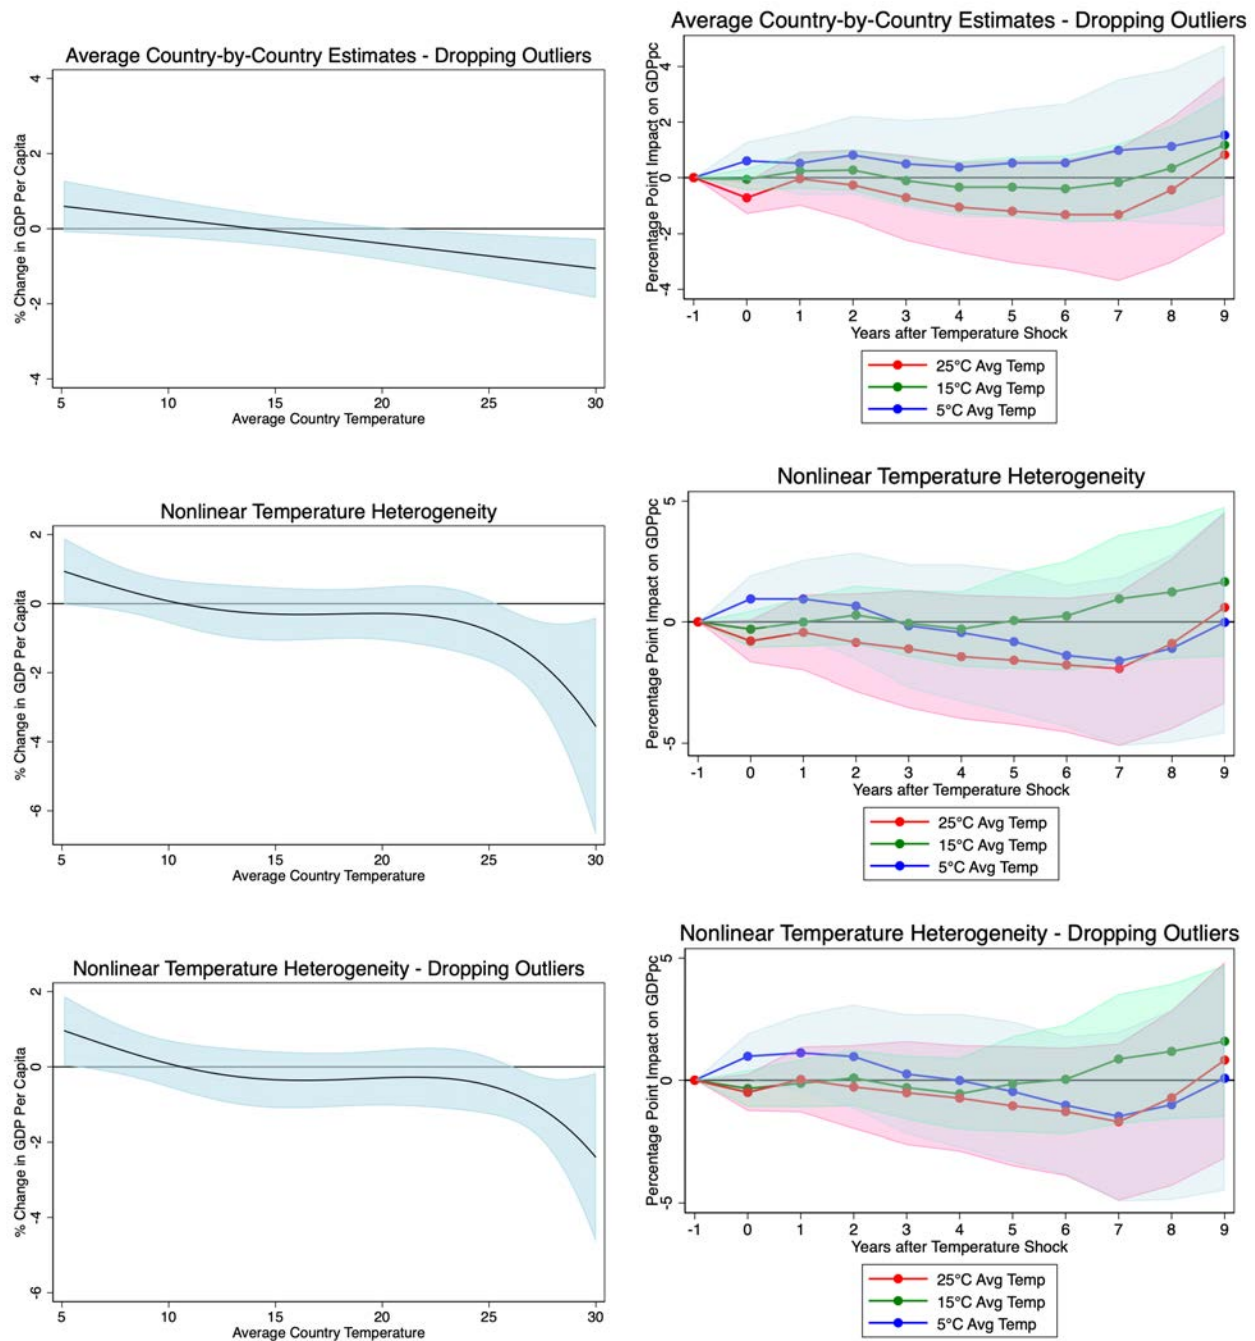
Notes: Graphs show local projections estimates of the persistent effects of an unanticipated 1°C temperature shock on GDP using Equation 6. Each graph on the left shows the contemporaneous effects of a 1°C temperature shock on GDP. The effect is allowed to vary with long-run average historical country temperature, which is shown on the x-axis. Each graph on the right shows the impulse response function of GDP to the initial shock to a 10-year horizon. Panel (a) shows estimates that use log levels of the lagged control variables rather than first-differences. Panels (b) and (c) show estimates using the Berkeley Earth Surface Temperature dataset and University of Delaware Temperature dataset, respectively, each with year fixed effects. In each impulse response function, the blue, green, and red lines represent cold (5°C), moderate (15°C), and hot (25°C) countries, respectively, and the 95% confidence intervals are shown with corresponding color shading. GDP data are from the World Development Indicators.

Figure A-2: Country-by-Country Local Projections Results



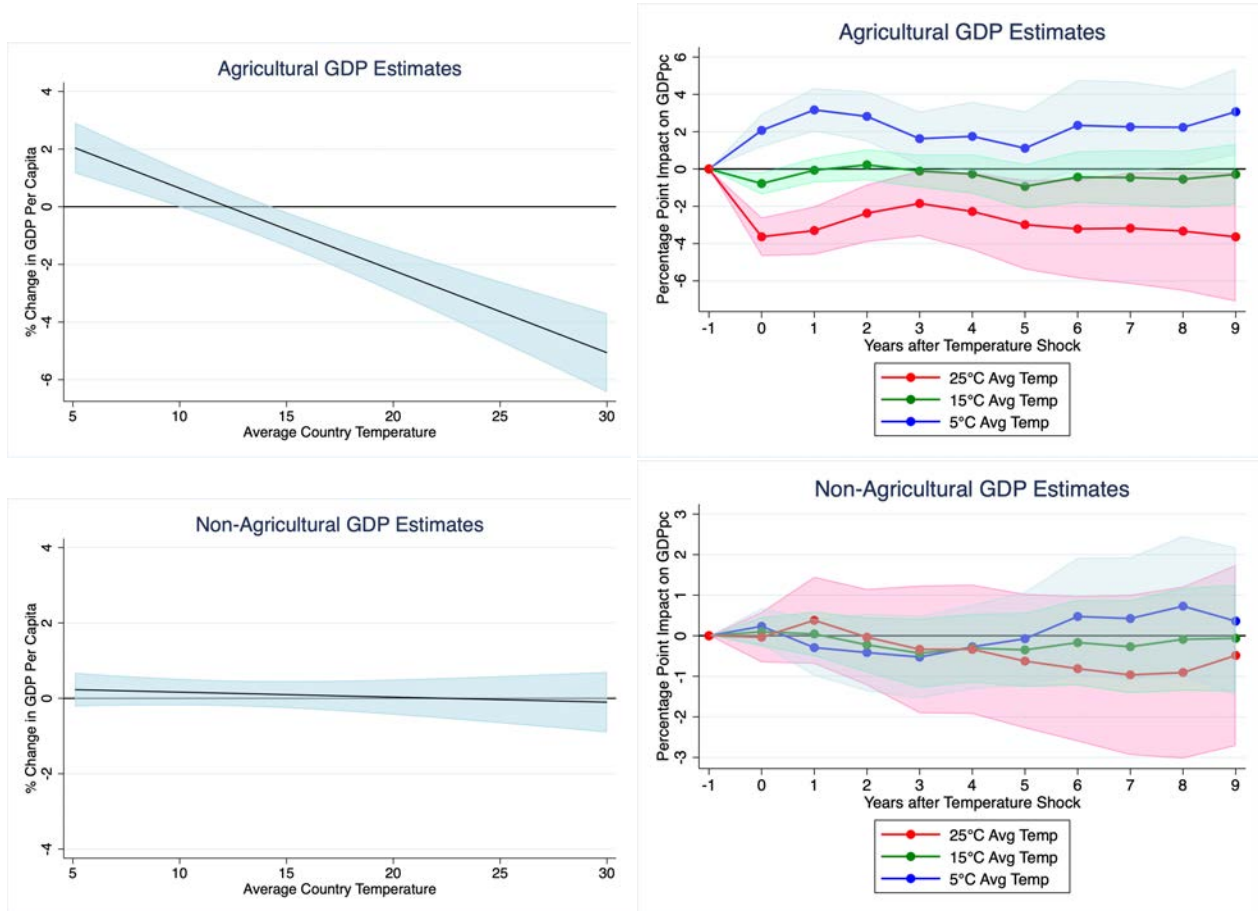
Notes: Panel (a) shows local projections estimates of the contemporaneous effects of an unanticipated 1°C temperature shock on GDP using Equation 6, estimated one country at a time for all 112 countries with at least 50 years of data in the sample. In panel (b), the graph on the left shows the average effects across the range of long-run average temperatures experienced throughout the world, which is shown on the x-axis. The 95% confidence interval is shaded in blue. The graph on the right of panel (b) shows the impulse response function of GDP to the initial shock to a 10-year horizon. In each impulse response function, the blue, green, and red lines represent cold (5°C), moderate (15°C), and hot (25°C) countries, respectively, and the 95% confidence intervals are shown with corresponding color shading. Temperature data are from the Global Meteorological Forcing Dataset. GDP data are from the World Development Indicators.

Figure A-3: Country-by-Country Local Projections Robustness Checks



Notes: Panel (a) shows local projections estimates of the contemporaneous effects of an unanticipated 1°C temperature shock on GDP using Equation 6, estimated one country at a time for all 112 countries with at least 50 years of data in the sample. In panel (b), the graph on the left shows the average effects across the range of long-run average temperatures experienced throughout the world, which is shown on the x-axis. The 95% confidence interval is shaded in blue. The graph on the right of panel (b) shows the impulse response function of GDP to the initial shock to a 10-year horizon. In each impulse response function, the blue, green, and red lines represent cold (5°C), moderate (15°C), and hot (25°C) countries, respectively, and the 95% confidence intervals are shown with corresponding color shading. Temperature data are from the Global Meteorological Forcing Dataset. GDP data are from the World Development Indicators.

Figure A-4: Agricultural and Non-Agricultural GDP



Notes: Graphs show local projections estimates of the persistent effects of an unanticipated 1°C temperature shock on GDP using Equation 6. Each graph on the left shows the contemporaneous effects of a 1°C temperature shock on GDP. The effect is allowed to vary with long-run average historical country temperature, which is shown on the x-axis. Each graph on the right shows the impulse response function of GDP to the initial shock to a 10-year horizon. Panel (a) shows the effects for agricultural GDP, and panel (b) shows the effects for nonagricultural GDP. Data on agricultural and nonagricultural GDP is compiled from Herrendorf, Rogerson and Valentinyi (2014), the UN National Accounts database, and the University of Groningen 10-Sector Database (Timmer, de Vries and De Vries, 2015). In each impulse response function, the blue, green, and red lines represent cold (5°C), moderate (15°C), and hot (25 °C) countries, respectively, and the 95% confidence intervals are shown with corresponding color shading. Temperature data are from GMFD, and GDP data are from the World Development Indicators.


Table A-2: Heterogeneous Effects of Temperature Shock on GDP
Country-by-Country Local Projections Estimates

	Dependent Variable: $\beta_{GDP}^{h=0}$			
	(1)	(2)	(3)	(4)
Country Mean Temperature	-0.096** (0.032)	-0.12** (0.044)	-0.13** (0.046)	-0.14** (0.048)
Dummy for Original OECD		-0.61 (0.54)		
Mean Agricultural Share of GDP			4.02 (2.77)	
Dummy for Poor Country in 1980				1.27 (0.74)
Constant	1.32* (0.53)	1.91* (0.93)	1.36* (0.56)	1.53* (0.59)
N	112	112	111	112


Notes: Table show how the effects of a 1°C temperature shock on contemporaneous GDP vary with country characteristics. The dependent variable in each regression is the coefficient $\beta_{GDP}^{h=0}$ estimated using Equation 6 for one country at a time. The independent variables in each regression include long-run average temperature in each country, and a variety of measures of levels of development and agricultural specialization.

Figure A-5: Model-Based Interpretation of Empirical Results
US TFP Control Specification

(a) Empirical vs. Simulated GDP Impulse Response Function




(b) Transition Dynamics with $\omega = 0.08$




Notes: The red line in panel (a) shows the empirical impulse response function of the path of GDP following an unanticipated 1°C shock to temperature in year 0, estimated using Equation 6 with a contemporaneous control for US TFP instead of year fixed effects, with the 95% confidence interval shaded in pink. The black line shows a model simulation with $\omega = 0.21$ of the impulse response function following a shock with magnitude calibrated to match the contemporaneous effect in year 0, and persistence calibrated to match the impulse response function of temperature shown in Figure 6b. Panel (b) shows a model simulation of the medium-term growth trajectory following a permanent shock starting in year 0 with $\omega = 0.21$ in orange, and $\omega = 0$ for comparison in green.

Figure A-6: Projected Impacts of Unabated Global Warming in Example Countries - US TFP Control Specification

(a) India



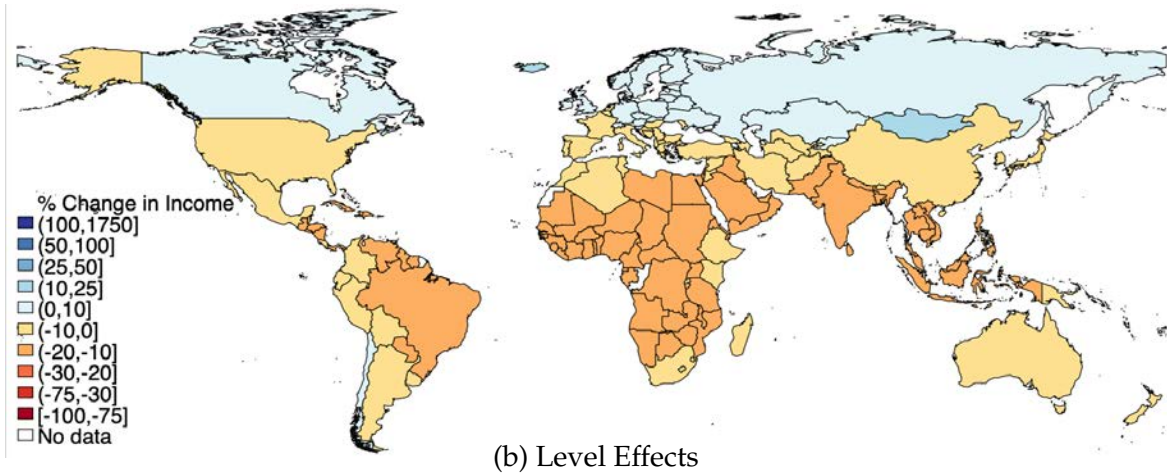
(b) Sweden



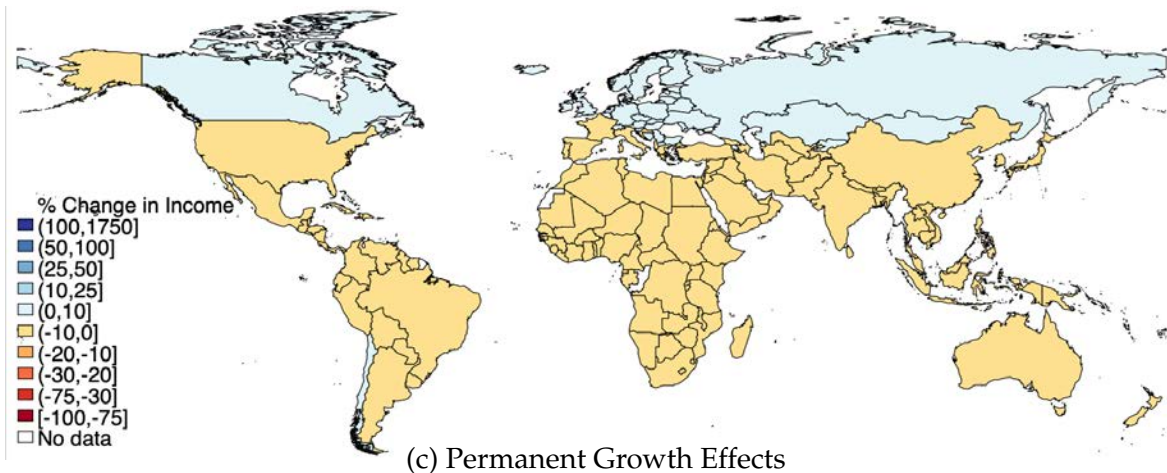
Notes: Graphs show the projected effects of unabated global warming on the trajectory of GDP under different projection methods for two example countries, India and Sweden. “Persistent growth effects” projections in orange use the 10-year cumulative response ratio shown in Figure 6c, from the specification with US TFP controls, to calibrate the effect of each degree of projected warming. “Level effects” projections in green use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” projections in red use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates. Corresponding projections for the specification with year fixed effects instead of US TFP controls are shown in Figure 9.

Figure A-7: Projected Impacts of Unabated Global Warming on Country-Level GDP

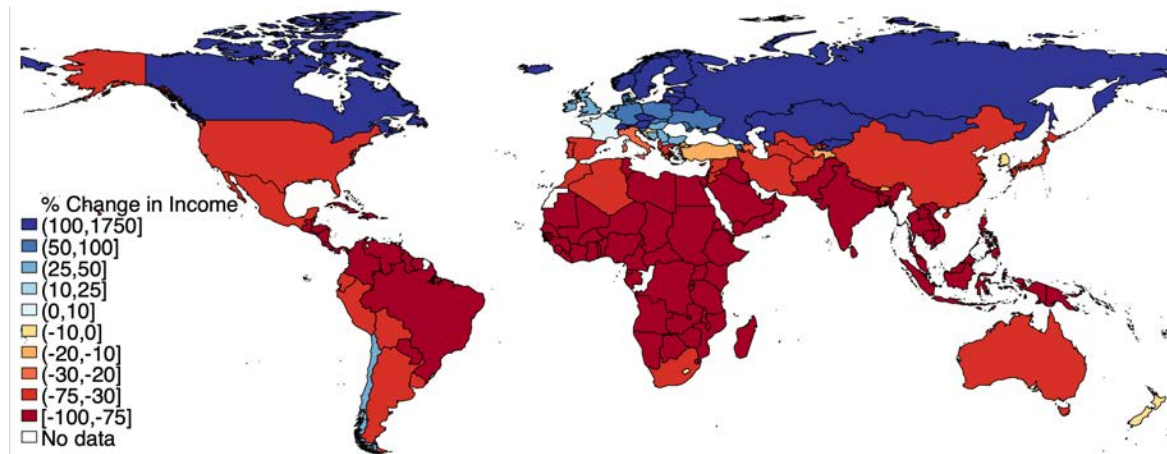
(a) Persistent Growth Effects - US TFP Control Specification



(b) Level Effects



(c) Permanent Growth Effects



Notes: Maps show the projected effects of unabated global warming on end-of-century country level GDP under different projection methods. “Persistent growth effects” estimates in panel (a) use the 10-year cumulative response ratio shown in Figure 6c, from the specification with contemporaneous US TFP controls, to calibrate the long-run level effect of each degree of projected warming. “Level effects” projections in panel (b) use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates.

**Table A-3: Projected Effects of Global Warming
Moderate Emissions Scenario**

Region	Persistent Growth Effects	Level Effects	Permanent Growth Effects
<i>Panel A - Year Fixed Effect Specification</i>			
Global GDP	-5.6	-1.0	-22.5
Global Population Average	-8.6	-1.8	-46.6
Sub-Saharan Africa	-11.3	-2.4	-67.4
Middle East & North Africa	-10.8	-2.1	-60.6
Asia	-9.5	-1.9	-55.1
South & Central America	-8.5	-1.6	-52.7
North America	-4.1	-0.5	-10.8
Europe	1.8	0.5	45.3
<i>Panel B - US TFP Control Specification</i>			
Global GDP	-3.2	-0.9	-22.5
Global Population Average	-5.1	-1.5	-46.6
Sub-Saharan Africa	-6.9	-2.1	-67.4
Middle East & North Africa	-6.2	-1.8	-60.6
Asia	-5.6	-1.7	-55.1
South & Central America	-4.9	-1.4	-52.7
North America	-2.0	-0.4	-10.8
Europe	0.8	0.4	45.3

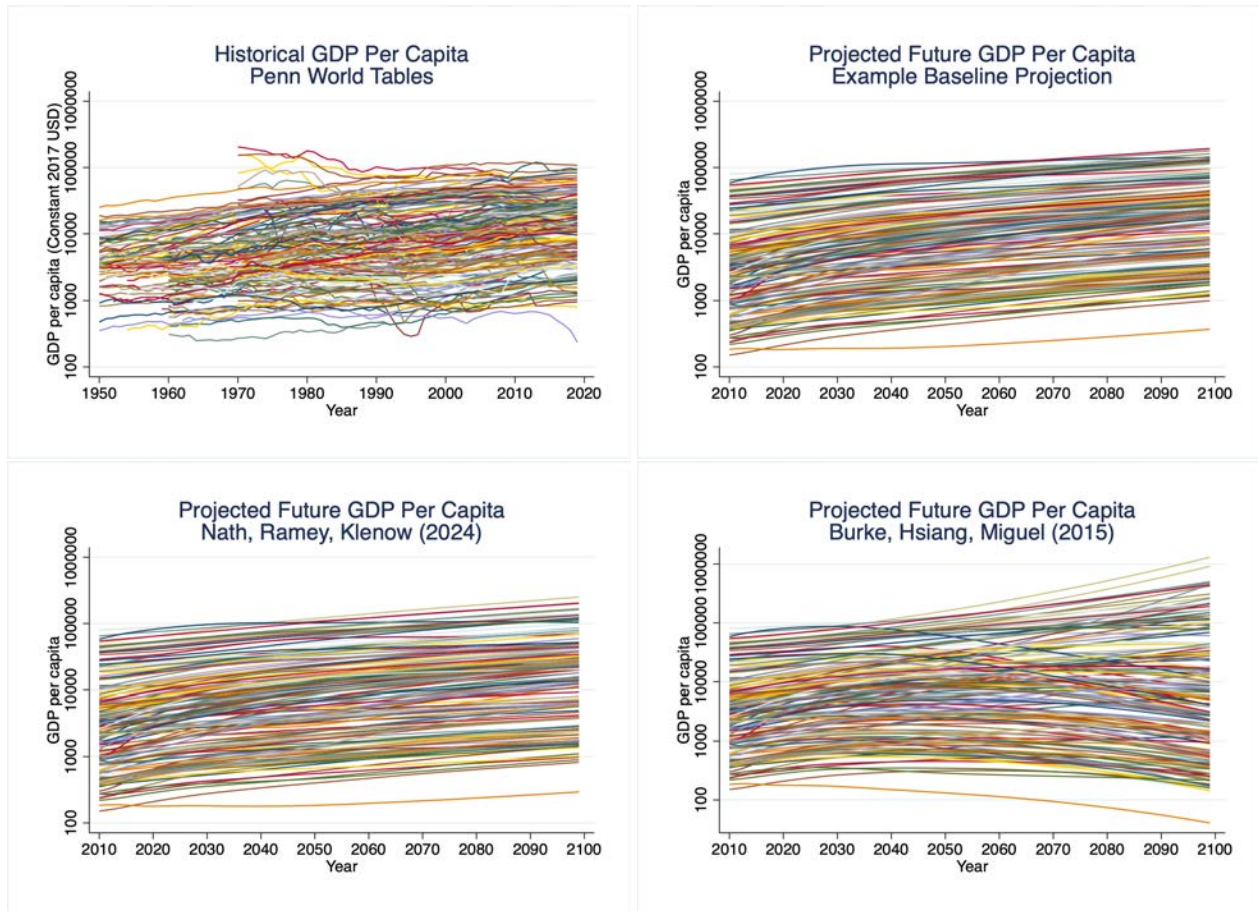
Notes: Table show the projected effects, in percent changes, of 2°C global warming on end-of-century GDP under different projection methods. The 2°C approximately corresponds to Representative Concentration Pathway 4.5. Country level temperature projections come from Burke et al. (2015), and are proportionately scaled down from the 3.7°C global average temperature increase in RCP 8.5. “Persistent growth effects” projections use the 10-year cumulative response ratio shown in Figure 6c to calibrate the effect of each degree of projected warming. “Level effects” projections use only the estimated contemporaneous effect of a 1°C shock, and allow for no persistence or accumulating effects. “Permanent growth effects” projections use estimates from Burke, Hsiang and Miguel (2015), and allow for the effects of temperature to permanently alter country-level growth rates. Panel (a) shows results for the local projections specification with year fixed effects, and panel (b) shows results for the specification with a contemporaneous control for US TFP instead.

**Table A-4: Estimated Effects of Historical Warming
From 1950s to 2010s**

	Year FE	US TFP Control
Global GDP	-0.8	-0.4
Global Population Average	-2.9	-1.6
Middle East & North Africa	-5.4	-2.8
Sub-Saharan Africa	-5.3	-3.1
South & Central America	-4.7	-2.6
Asia	-3.2	-1.8
North America	-0.3	-0.1
Europe	2.1	0.9

Notes: Table show the estimated effects, in percent changes, of the global warming that has occurred between 1950 and 2019. Projections are made using the average country-level decadal temperature from 1950 to 1959 as the starting period and 2010 to 2019 as the ending period. Temperature data is from the Berkeley Earth Surface Temperature dataset. Projections use the 10-year cumulative response ratio shown in Figure 6c to calibrate the effect of each degree of observed warming. Column 1 shows results for the local projections specification with year fixed effects, and Column 2 shows results for the specification with a contemporaneous control for US TFP instead.

Figure A-8: Visualizing Income Level Differences and Growth Convergence



Notes: Top left panel shows historical GDP per capita across the world from 1950 to 2019, as measured by the Penn World Tables. Top right panel shows an example future projection of global incomes in the 21st century from Dellink et al. (2017). The projection is from Shared Socioeconomic Pathway (SSP) Four, which is labeled "Moderate Growth, Low Convergence." The bottom left panel shows a version of the SSP 4 projections that incorporate estimated climate damages from this paper, and the bottom right panel shows a version of the SSP 4 projections that incorporate estimated climate damages from Burke, Hsiang and Miguel (2015).