WIFO ■ WORKING PAPERS
730/2026

Climate Change and Economic Dynamics: Temperature-dependent Shock Propagation

Christian Glocker
Thomas Url

Climate Change and Economic Dynamics: Temperature-dependent Shock Propagation

Christian Glocker, Thomas Url

Research assistance: Astrid Czaloun, Ursula Glauninger

WIFO Working Papers 730/2026
July 2026

Abstract

The gradual transition between different climate states, as evidenced by the smooth rise in temperatures, motivates conceptualizing climate change as a phenomenon influencing the propagation mechanism of traditional macroeconomic shocks, rather than as an independent shock. We formalize this idea through a theoretical model, which shows that climate change induces a structural shift in the economy by steepening the aggregate supply curve, thereby exacerbating the price effects of demand shocks while dampening the output response. Our empirical evidence is consistent with this prediction: higher temperatures raise the share of inflation variation attributable to demand shocks by up to 10 percentage points, underscoring the role of climate change in intensifying stagflationary dynamics rather than independently driving business cycle fluctuations.

E-Mail: christian.glocker@wifo.ac.at, thomas.url@wifo.ac.at

2026/1/W/0

Media owner (publisher), producer: Austrian Institute of Economic Research
1030 Vienna, Arsenal, Objekt 20 | Tel. (43 1) 798 26 01 0 | https://www.wifo.ac.at
Place of publishing and production: Vienna

CLIMATE CHANGE AND ECONOMIC DYNAMICS: TEMPERATURE-DEPENDENT SHOCK PROPAGATION

CHRISTIAN GLOCKER AND THOMAS URL

Abstract. The gradual transition between different climate states, as evidenced by the smooth rise in temperatures, motivates conceptualizing climate change as a phenomenon influencing the propagation mechanism of traditional macroeconomic shocks, rather than as an independent shock. We formalize this idea through a theoretical model, which shows that climate change induces a structural shift in the economy by steepening the aggregate supply curve, thereby exacerbating the price effects of demand shocks while dampening the output response. Our empirical evidence is consistent with this prediction: higher temperatures raise the share of inflation variation attributable to demand shocks by up to 10 percentage points, underscoring the role of climate change in intensifying stagflationary dynamics rather than independently driving business cycle fluctuations.

JEL codes: C33; E31; E32; Q54
Key words: Climate change; Business cycles; Inflation; Interacted panel VAR

C. Glocker: Austrian Institute of Economic Research, Arsenal Objekt 20, 1030 Vienna, Austria. Phone: +43 (0) 1 789 26 01-303, E-mail: *christian.glocker@wifo.ac.at*

T. Url: Austrian Institute of Economic Research, Arsenal Objekt 20, 1030 Vienna, Austria. Phone: +43 (0) 1 789 26 01-494, E-mail: *thomas.url@wifo.ac.at*

This version: July 2026

The authors would like to thank Maximilian Böck, Tobias Scheckel, Nawid Siassi, Stefan Schleicher and seminar participants at research seminars at the Austrian Institute of Economic Research (Wifo) and the Technical University Vienna (TU Wien) for valuable comments and helpful discussions. Excellent research assistance by Astrid Czaloun and Ursula Glauninger is gratefully acknowledged.

1. Introduction

In its latest assessment report, the Intergovernmental Panel on Climate Change (IPCC) has reaffirmed anthropogenic influence on the climate (Lee et al., 2023). This has already resulted in an increase in the global (surface) mean temperature of approximately 1.8 °C, a trend that is also evident in regional temperatures, as illustrated in Figure 1. The gradual increase has prompted active research in the field of economics. Several studies (e.g. Donadelli et al., 2017; Pretis, 2021; Mukherjee and Ouattara, 2021; Huber et al., 2023; Ciccarelli et al., 2024; Iliyasu and Sanusi, 2024) conceptualize climate change as an exogenous shock to the economic system at the business cycle frequency, akin to other "traditional" shocks, such as those resulting from economic policy changes, preference shifts, cost-push shocks, and the like. While this approach provides valuable insights, characterizing climate change (especially temperature anomalies) as a "shock" contradicts its observed smooth trajectory.

When climate scientists argue that the recent temperature increase should be considered a shock to the climate, they do so with a time horizon spanning thousands of years; see Osman et al. (2021) among others and Figure 2. Similarly, the perception of a climate-induced shock is also akin to the growth literature (or long-run macroeconomic analysis), in which time periods covering three to five decades are considered (Dell et al., 2012; Bilal and Stock, 2025). From a business cycle perspective, however, such a temperature increase represents a smooth and persistent transition from one climate state to another, signifying a *structural change*. This interpretation is motivated by the observation that the sustained rise in temperature over the past five decades far exceeds the average length of a business cycle. It is therefore controversial to conceptualize climate change as a "business cycle relevant shock", akin to those driving short term economic fluctuations. In this context, Ciccarelli and Marotta (2024) acknowledge that while climate-related shocks influence business cycles, their contribution to variations in output and prices appears limited, as indicated by the forecast error variance decomposition.

Against this background, we put forth the *general hypothesis* that climate change induces a structural change in the economic system, which in turn influences the propagation mechanism of traditional shocks. This is to say

FIGURE 1. Global temperature anomalies

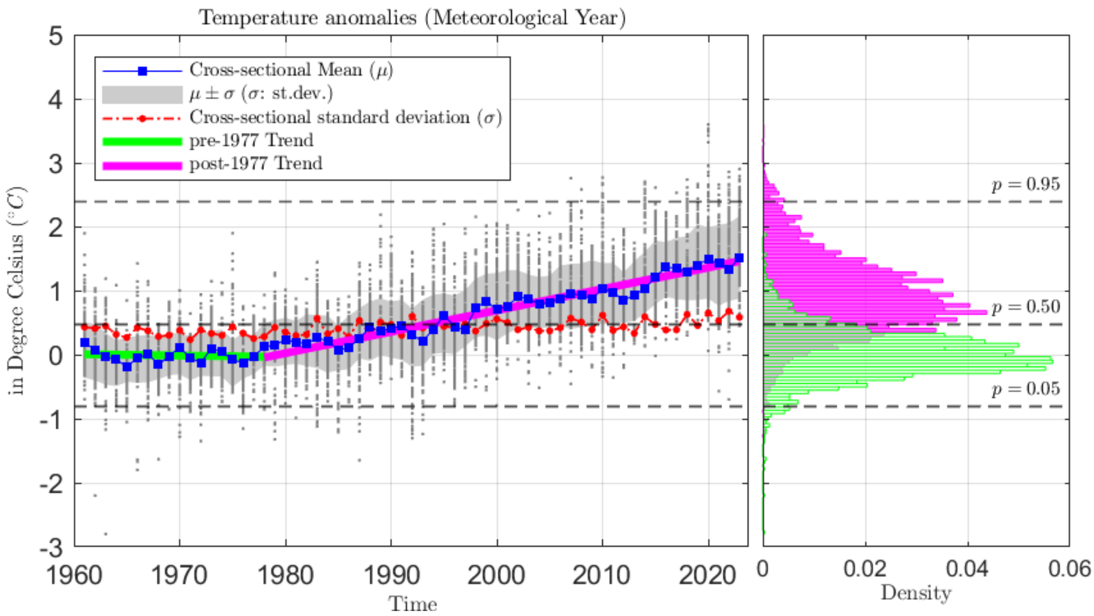

Note: The figure shows the (annual) FAOSTAT temperature anomaly (i.e. deviation) by country with respect to a baseline climatology, corresponding to the period 1951–1980. The series cover the period 1961–2023 for 198 countries. The multiple breakpoint test of Bai and Perron (1998, 2003) identifies a breakpoint for the global mean temperature anomaly with a break date in 1977; the green and magenta lines indicate the corresponding trend lines with histograms on the right. The 5th, 50th and 95th percentiles are derived from the joint distribution, i.e. over all countries and the whole period. While the pre-1977 distribution is still fairly symmetric, the post-1977 distribution is right-skewed.

that climate change does not constitute a source of shocks relevant for business cycle fluctuations in and of itself. Building on this, we put forth the *particular hypothesis* that climate change steepens the Phillips curve (or equivalently, the aggregate supply (AS) curve), thereby exacerbating the price effects of demand shocks while dampening the output response. The following three examples are intended to illustrate this basic idea.

Agricultural sector. *In a stable climate, a positive demand shock for wheat raises prices, incentivizes higher production, and stimulates overall economic activity through investment, employment, and income gains. Under climate change, however, rising frequencies of extreme weather events and crop failures depress productivity and restrict supply capacities (Barrios et al., 2008; Ortiz-Bobea et al., 2021). As a result, the same demand shock now triggers a stronger increase in prices but only limited output expansion, reflecting a steeper Phillips curve in agriculture.*

Figure 2. Global temperature anomalies (long series)

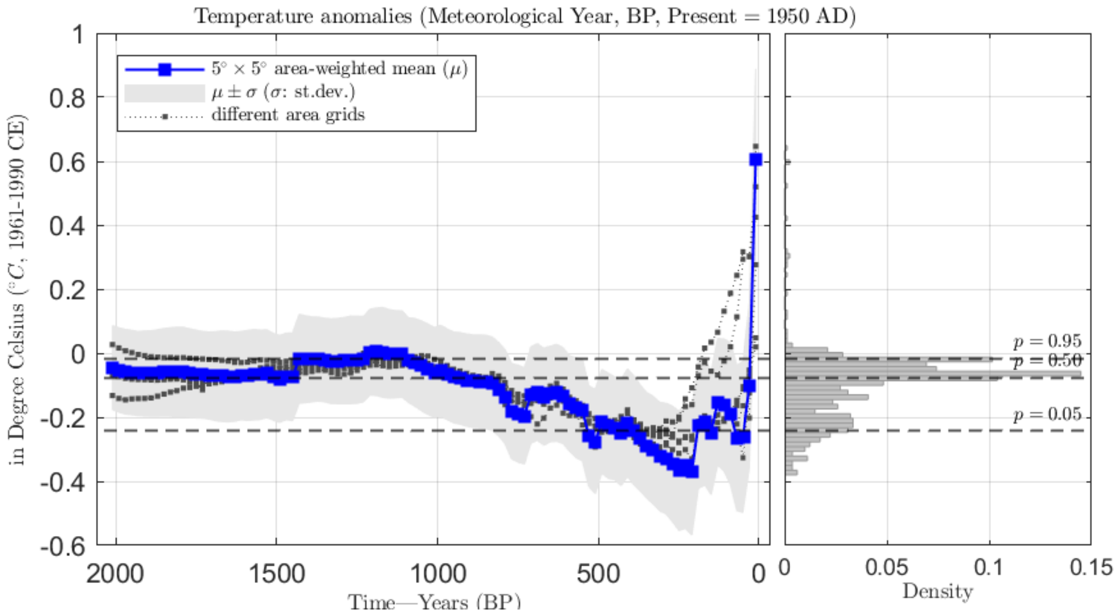

Note: The figure shows various measures for the global temperature anomaly (i.e. deviation) over a 2,000-year time-span with respect to a baseline climatology (CE), corresponding to the mean of the period 1961–1990. The time-axis identifies all ages in year BP (Before Present), where Present = 1950. The last (i.e., most recent) data point refers to the year 1940. The blue line shows the global stacked temperature anomaly for the $5^\circ \times 5^\circ$ grid area-weighted mean with its one-standard deviation uncertainty. The various black-dotted lines show global temperature anomalies based on alternative area weightings ($30^\circ \times 30^\circ$ grid, area-weighting using 10° latitude bins, etc.). The series are taken from Marcott et al. (2013) and available following this link.

Construction sector. *In construction, demand shocks typically boost output and employment. Yet, climate-related heat regulations, which mandate reduced working hours during excessively hot periods, increasingly constrain productivity (Zhang et al., 2023). While labor costs persist, output gains are muted due to frequent heat-induced disruptions (De Sario et al., 2023). Hence, under higher temperatures, a demand shock primarily raises costs and prices, while the output response remains weak, reflecting a steeper Phillips curve in construction.*

Services sector. *In services, particularly in urban areas, climate change manifests through higher cooling requirements. Rising temperatures elevate energy and adaptation costs to sustain productivity, including air conditioning and infrastructure adjustments (Fesselmeyer et al., 2024). Consequently, demand shocks amplify price pressures via higher marginal costs, while output responses remain largely unchanged, reflecting a steeper Phillips curve in services.*

These examples[1] illustrate how climate change can transform the effect of a positive demand shock into a source of stagflation. We use this as a starting point to examine the role of climate change in shaping the propagation mechanism of traditional macroeconomic shocks.

Measuring climate change. A critical inquiry at this juncture is how to effectively "measure" climate change for macroeconomic analysis. A natural definition of global climate encompasses the distribution of climatic outcomes–such as temperature, precipitation, wind speed, and ice cover–over time and space. While it is evident that a comprehensive characterization of global climate is unattainable, research indicates that a pivotal state variable encapsulating climate is the (mean) surface temperature (Mohnen et al., 1990; Allen et al., 2009; Bento et al., 2023). This offers a practical means for macroeconomic analysis as it reduces climate change to a single, quantifiable variable. Notably, sophisticated climate model simulations demonstrate that projections of various climate parameters, such as precipitation and regional temperatures, can be adequately approximated as functions of mean temperature (Manne et al., 1995; Held and Soden, 2006; Hassler et al., 2016). Bilal and Känzig (2026) also stress that there is a strong correlation between global temperature and the frequency of extreme climatic events. It remains crucial to distinguish between slow-moving, persistent changes in climate and transitory variations in weather. As can be seen in Figure 1, climate change refers to a series of small and persistent upward shifts in the statistical distribution of temperature anomalies, whereas transitory weather shocks are short term deviations from this trend. The temperature anomaly in Figure 1 deviates from the interpretation of persistent but trend-reverting temperature shocks proposed by Bilal and Känzig (2026). Global warming causes a sustained rise in temperature anomalies, and our interest lies in the structural change in the economic system triggered by this development.

We begin by formalizing our idea in a theoretical model. It features a two-country setting consisting of a small open economy and the "rest of the world" which together determine the anthropogenic influence on climate. The framework builds on the New Keynesian model (Galí, 2015, Chapter 3), extended to include an energy good produced with both fossil and renewable sources

[1] Bilal and Stock (2025) provide an in-depth survey of the literature on the impact of climate change across different economic sectors.

(Glocker and Wegmüller, 2024). Fossil energy use generates carbon emissions, which accumulate in the atmosphere (Golosov et al., 2014), raise global temperatures, and, in turn, reduce total factor productivity (TFP). The latter follows Nordhaus (1991, 2014), whose model highlights feedback loops between emissions, atmospheric concentrations, rising temperatures, and their economic damages. The central result is that this mechanism steepens the Phillips curve: higher temperatures make demand shocks trigger stronger price reactions but weaker output responses.

We test this prediction empirically using a model aligned with the structure of the theoretical framework. The negative feedback of temperature on TFP is proxied by measured surface temperature anomalies, motivating a structural vector autoregressive (SVAR) specification with an interaction term. This approach creates a non-linearity and allows impulse responses to vary with different levels of temperature anomalies. We identify global demand and supply shocks using sign restrictions derived from the theoretical model and estimate the SVAR over a panel of countries. Our results confirm the theoretical predictions. First, likelihood-ratio tests and information criteria favor modeling temperature as an interaction variable rather than as an exogenous variable in the SVAR, supporting our general hypothesis. Second, global demand and supply shocks produce the expected significant responses in domestic output and prices, most importantly, higher temperature anomalies magnify the price effects of demand shocks while dampening output reactions. By contrast, supply shock responses remain largely unaffected. These findings are consistent with our particular hypothesis: temperature anomalies alter the slope of the Phillips curve, affecting the transmission of demand but not supply shocks.

Related literature. Our contribution links to three strands. The first emphasizes structural determinants of business cycle propagation; the second studies inflationary effects of climate policies; the third investigates climate change and macroeconomic performance. Regarding the latter, Tol (2024) and Bilal and Stock (2025) provide two recent surveys assessing the long-run economic costs and non-market impacts from climate change. The central estimate of the economic impact of global warming is negative with a wide confidence interval. Research activity concentrates on (i) output effects (Dell et al., 2012; Zhao et al., 2018; Berg et al., 2024; Bilal and Känzig, 2026; Nath et al., 2024; Abalo et al., 2025) and (ii) productivity development (Donadelli et al., 2017; Acevedo et al., 2020; Zhang et al., 2023; Kumar and Maiti, 2024). Iliyasu and Sanusi (2024) document GDP losses and price increases as a consequence of temperature anomalies. All these studies are based on a broad set of countries and therefore use annual data. Only a few authors analyze the feedback between output and climatic outcomes using data at higher frequencies more related to the business cycle or to seasonal patterns (Doda, 2014; Cohen et al., 2022; Linsenmeier, 2024). In contrast, we do not attempt to estimate the social cost of carbon, rather we focus on the impact of higher temperature anomalies on the propagation mechanism of demand and supply shocks. Closer to our work, Ciccarelli and Marotta (2024) identify climate-related risks in a SVAR and show that physical risks act as negative demand shocks while transition risks shift supply, but they do not use weather outcomes as indicators of climate change and they do not allow for a non-linear response of the economy to temperature anomalies.

On the inflationary effect of climate policies, Cantelmo et al. (2024) stress monetary policy's stabilizing role in responding to climate-induced natural disaster shocks, while van der Ploeg and Withagen (2014) discuss the ambiguous effects of optimal carbon taxation. In a related line, Bettarelli et al. (2025) show that carbon mitigation policies generate substantial and persistent inflationary pressures. We extend this literature by showing that higher temperature anomalies are likely to magnify the price effect of expansionary demand shocks and simultaneously reduce their output effect.

Finally, regarding structural elements of propagation mechanisms, Abbritti and Weber (2018) highlight the role of labor market institutions in amplifying or dampening business cycle fluctuations, while Christoffel et al. (2009) show their importance for monetary policy transmission. Bai et al. (2026) stress the role of supply chain disruptions in this context. Crespo Cuaresma and Glocker (2023) emphasize production structure, particularly tradable–nontradable shares, for (fiscal-) demand shocks, and Towbin and Weber (2013) analyze how exchange rate regimes, debt structure, and import composition shape the propagation of global shocks into domestic economies. We suggest that temperature anomalies are another element shaping an economy's dynamic response to demand shocks.

The paper is structured as follows. Section 2 formalizes the research idea and Section 3 introduces a conforming theoretical model. Section 4 introduces the empirical model, discusses the connection between the theoretical and empirical models, and presents the results of the econometric analysis. Finally, Section 5 concludes.

2. Formalizing the research idea

To express the research idea in more detail, consider a dynamic system composed of a vector $\boldsymbol{y}_t$ of endogenous variables, whose dynamics are governed by the matrix lag-polynomial $\boldsymbol{B}(L) = \sum_{i=0}^{\mathcal{P}} \boldsymbol{B}_i L^i$ of order $\mathcal{P}$ with L being the lag operator and $\boldsymbol{B}_0 = \boldsymbol{I}$ (identity matrix). The system is subject to the exogenous shocks in $\boldsymbol{u}_t$ and reads

$$\boldsymbol{B}(L)\boldsymbol{y}_t = \boldsymbol{H}\boldsymbol{u}_t \tag{1}$$

where $\boldsymbol{H}$ governs the (contemporaneous) impact response of the shocks on the endogenous variables. Zolghadr-Asli et al. (2021); Casey et al. (2023); Ciccarelli et al. (2024) introduce climate change ($\mathcal{T}_t$) in this system by expanding $\boldsymbol{u}_t$ to $\boldsymbol{u}_t^* = [\boldsymbol{u}_t', \mathcal{T}_t]'$, so that equation (1) changes to

$$\boldsymbol{B}(L)\boldsymbol{y}_t = \boldsymbol{H}^* \begin{bmatrix} \boldsymbol{u}_t \\ \mathcal{T}_t \end{bmatrix} \tag{2}$$

where $\boldsymbol{H}^*$ extends $\boldsymbol{H}$ by one additional column. This setup assigns climate change a role analogous to that of traditional shocks in macroeconomic models. One might conceptualize the damage from a natural disaster as a comparable example of such a shock. The aforementioned studies use temperature anomalies as an external instrument for climate change and find that they have a statistically significant effect on key macroeconomic variables. However, as indicated by the forecast error variance decomposition (see, for instance, Ciccarelli and Marotta, 2024), these anomalies account for only a negligible share of the fluctuations in these variables. Considering this, we suggest that this result may reflect limitations in the empirical specification. Rather than treating climate change as a conventional shock, it may be more appropriate to consider it as a driver of structural change, influencing the coefficients in the

$\boldsymbol{B}(L)$ and $\boldsymbol{H}$ matrices. This perspective extends equation (1) to

$$\boldsymbol{B}(L;\mathcal{T}_t)\boldsymbol{y}_t = \boldsymbol{H}(\mathcal{T}_t)\boldsymbol{u}_t \tag{3}$$

where $\boldsymbol{B}(L;\mathcal{T}_t) = \sum_{i=0}^{\mathcal{P}} \boldsymbol{B}_i(\mathcal{T}_t)L^i$ with $\boldsymbol{B}_0(\mathcal{T}_t) = \boldsymbol{I}$. In this framework, climate change influences the dynamics of the system, rather than acting as an exogenous shock: it alters the propagation mechanisms of the shocks in $\boldsymbol{u}_t$. To explore this more formally, we consider a first-order Taylor approximation of the coefficient matrices with respect to the climate change variable $\mathcal{T}_t$, evaluated around a mean level $\bar{\mathcal{T}}$. Substituting these approximations into equation (3), we obtain

$$\boldsymbol{y}_t + \sum_{i=1}^{\mathcal{P}} \left(\boldsymbol{B}_i(\bar{\mathcal{T}}) + \frac{\partial \boldsymbol{B}_i}{\partial \mathcal{T}_t}\hat{\mathcal{T}}_t \right) \boldsymbol{y}_{t-i} \approx \left(\boldsymbol{H}(\bar{\mathcal{T}}) + \frac{\partial \boldsymbol{H}}{\partial \mathcal{T}_t}\hat{\mathcal{T}}_t \right) \boldsymbol{u}_t \tag{4}$$

where $\hat{\mathcal{T}}_t = \mathcal{T}_t - \bar{\mathcal{T}}$ is the temperature anomaly.[2] Equation (4) represents a multivariate dynamic model in which $\hat{\mathcal{T}}_t$ acts as an interaction variable. The variable $\hat{\mathcal{T}}_t$ introduces a non-linearity in the model and exerts an influence on the transmission channel of traditional business cycle shocks; however, it does not manifest as an exogenous (business cycle) shock. In what follows, we present a theoretical model that is consistent with equation (3).

3. The theoretical model

The model economy consists of a two-country framework with a small open economy (the "home" country) and the rest of the world (ROW), jointly forming a global economy that shapes anthropogenic climate dynamics. The framework extends the New Keynesian model (Galí, 2015, Chapter 3) to include an energy good produced from both fossil and renewable sources (Glocker and

[2]While $\hat{\mathcal{T}}_t$ is defined as a deviation of the mean surface temperature from its historical baseline, the underlying physical and economic mechanism primarily operates through the frequency and intensity of extreme heat events rather than through modest shifts in the mean per se. According to Ruff and Neelin (2012), in a Gaussian or near-Gaussian weather regime, an increase in the mean of the temperature distribution mechanically raises the probability $\Pr(\mathcal{T} > \mathcal{T}_{\text{critical}})$ that (daily) temperatures exceed critical thresholds at which labor productivity, health, and infrastructure are severely impaired. Consequently, our use of the temperature anomaly should be interpreted as a reduced-form statistic for shifts in the entire distribution of temperature, and in particular for the increase in the right-tail probability of "hot days". Formally, for the temperature distribution $\mathcal{T}$ can be well approximated by a distribution $F_{\mathcal{T}}(\,\cdot\,|\mu_t, \sigma_t)$ whose mean μ_t co-moves tightly with $\Pr(\mathcal{T} > \mathcal{T}_{\text{critical}}) = 1 - F_{\mathcal{T}}(\mathcal{T}_{\text{critical}}|\mu_t, \sigma_t)$, so that using $\hat{\mathcal{T}}_t$ as state variable is equivalent to parameterizing the evolution of $\Pr(\mathcal{T} > \mathcal{T}_{\text{critical}})$.

Wegmüller, 2024). Nominal price rigidities represent the key friction. An energy composite is produced from fossil and renewable raw inputs, with fossil energy generating carbon emissions. Energy is demanded by both households and firms, while monetary policy follows a flexible exchange rate regime. The pivotal feature of the model is the feedback from carbon emissions to TFP, where the global temperature anomaly serves as a proxy for climate change (Nordhaus, 2014; Fernández-Villaverde et al., 2024; Barrage and Nordhaus, 2024).

The analysis focuses on the home economy, while ROW variables (denoted by *) and global aggregates (denoted by **) are introduced explicitly when relevant. Model implications are derived in response to a foreign demand and a foreign supply shock.

3.1. **Energy producers.** Energy producers operate in perfectly competitive markets. They combine imported fossil energy $e_{f,t}$, purchased at foreign price $P^*_{e,t}$, with domestic renewable energy $e_{r,t}$, produced from renewable capital k_t rented from households at the price $P_{k,t}$. One unit of renewable capital produces one unit of renewable energy: $e_{r,t} = k_{t-1}$. Final energy e_t is produced according to

$$e_t = \xi_E \left((1-\zeta) e_{f,t}^{\frac{\kappa-1}{\kappa}} + \zeta k_{t-1}^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}}, \tag{5}$$

where ξ_E denotes technology, κ the elasticity of substitution, and $\zeta \in [0,1]$ the renewable energy share. Profits are given by $P_{e,t}e_t - S_t P^*_{e,t} e_{f,t} - P_{k,t}k_{t-1}$, where S_t is the nominal exchange rate (expressing the foreign currency in terms of the domestic currency). Optimal input choice yields relative demand

$$\frac{k_{t-1}}{e_{f,t}} = \frac{1-\zeta}{\zeta} \left(\frac{P_{k,t}}{S_t P^*_{e,t}} \right)^{-\kappa}. \tag{6}$$

The corresponding domestic energy price index is

$$P_{e,t} = \left((1-\zeta)(S_t P^*_{e,t})^{1-\kappa} + \zeta P_{k,t}^{1-\kappa} \right)^{\frac{1}{1-\kappa}}, \tag{7}$$

The foreign energy price $P^*_{e,t}$ follows an exogenous AR(1) process with $u_{P^*_e,t}$ being the foreign supply shock.

3.2. **Fossil energy, carbon emissions and temperatures.** Fossil energy use generates carbon emissions that accumulate in the global carbon stock M_t^{**}

$$M_t^{**} = (1-\varphi)M_{t-1}^{**} + e_{f,t} + e_{f,t}^*, \tag{8}$$

where φ is the decay rate. ROW fossil use $e_{f,t}^*$ is assumed to depend positively on foreign demand Y_t^* and negatively on $P_{e,t}^*$. Following Hassler and Krusell (2012), global carbon concentration determines the temperature anomaly $\hat{\mathcal{T}}_t$, defined as the deviation from its steady state $\bar{\mathcal{T}}$

$$\hat{\mathcal{T}}_t = \mathcal{T}\left(\hat{M}_t^{**}\right) - \bar{\mathcal{T}}, \quad \hat{M}_t^{**} = \ln(M_t^{**}/\bar{M}^{**}).$$

We assume temperature anomalies are uniform across the two countries. Since climate responses to emissions unfold only over decades–consider Nordhaus (1991); Ricke and Caldeira (2014) and the discussion in Section 1–we treat temperature anomalies as exogenous in the short run, i.e. $\hat{\mathcal{T}}_t = \hat{\mathcal{T}}$. We assume that TFP ($\xi_A$) depends negatively on temperature anomalies (Letta and Tol, 2019; Kumar and Maiti, 2024; Donadelli et al., 2022; Drescher and Janzen, 2025; van der Ploeg and Rezai, 2026)[3]

$$\xi_A(\hat{\mathcal{T}}) = e^{-\alpha\hat{\mathcal{T}} + \bar{\xi}_A}, \tag{9}$$

where α captures the semi-elasticity of TFP to temperature anomalies. Estimates suggest that a one-degree rise in temperature reduces TFP by about 3.2 percent (Kumar and Maiti, 2024), with stronger effects in advanced economies (Donadelli et al., 2022). Additional evidence links higher temperatures to adverse labor productivity and occupational health outcomes, further corroborating the negative impact of climate change on productivity (Letta and Tol, 2019; Henseler and Schumacher, 2019).

3.3. **Households.** We consider a continuum of infinitely lived households that derive utility from consumption C_t and disutility from labor h_t. Preferences are represented by $u(C_t, h_t) = \frac{C_t^{1-\sigma}-1}{1-\sigma} - \chi\frac{h_t^{1+\phi}}{1+\phi}$, where σ denotes relative risk

[3]Importantly, these studies document that economic damages are highly non-linear in temperature, with sizeable effects arising once (daily) temperatures exceed critical thresholds for human physiology and physical infrastructure. In practice, a rise in the mean anomaly $\hat{T}$ over a season implies a higher share of days with $\mathcal{T} > \mathcal{T}_{\text{critical}}$, so that $\xi(\hat{\mathcal{T}})$ should be interpreted as decreasing in the probability and intensity of such threshold exceedances, for which $\hat{\mathcal{T}}$ serves as a sufficient statistic.

aversion, ϕ the inverse Frisch elasticity, and χ is a scaling parameter. Households hold domestic (B_t) and foreign bonds (B_t^*), with foreign assets valued at exchange rate S_t. Limited capital mobility is introduced via a quadratic cost $\psi_{B,t} = \frac{\psi_B}{2}\left(\frac{S_t}{P_t}B_t^*\right)^2$. They also invest I_t in renewable energy capital k_t, rented to energy producers at price $P_{k,t}$. Capital depreciates at rate δ

$$k_t = (1-\delta)k_{t-1} + I_t. \tag{10}$$

Nominal income consists of labor earnings $W_t h_t$ plus returns on bond and capital holdings. The budget constraint is

$$P_t(C_t + I_t - \psi_{B,t}) + B_t + S_t B_t^* = R_{t-1}B_{t-1} + R^* S_t B_{t-1}^* + W_t h_t + P_{k,t}k_{t-1}. \tag{11}$$

Households maximize discounted life-time utility $E_0 \sum_{t\geq 0} \beta^t u(C_t, h_t)$ subject to (10) and (11), with discount factor β. The optimality conditions are

$$1 = E_t\left[\Lambda_{t,t+1}\frac{R_t}{1+\pi_{t+1}}\right], \tag{12}$$

$$1 - \psi_B \frac{S_t}{P_t}B_t^* = E_t\left[\Lambda_{t,t+1}\frac{R^*}{1+\pi_{t+1}}\frac{S_{t+1}}{S_t}\right], \tag{13}$$

$$\frac{W_t}{P_t} = \chi h_t^{\phi} C_t^{\sigma}, \tag{14}$$

$$1 = E_t\left[\Lambda_{t,t+1}\left(1-\delta+\frac{P_{k,t+1}}{P_{t+1}}\right)\right], \tag{15}$$

where $\Lambda_{t,t+k} = \beta^k (C_t/C_{t+k})^{\sigma}$ is the stochastic discount factor, and $1+\pi_t = P_t/P_{t-1}$ is the gross-inflation rate. We assume that consumption C_t is composed of non-energy (c_t) and energy ($e_{H,t}$) consumption goods according to

$$C_t = \min\left\{\frac{c_t}{1-\vartheta}, \frac{e_{H,t}}{\vartheta}\right\} \tag{16}$$

and ϑ measures the share of energy in total household consumption. To find an optimal allocation between c_t and $e_{H,t}$, the household solves $\max_{c_t, e_{H,t}} C_t$ using equation (16) and the constraint $P_t C_t = P_{c,t}c_t + P_{e,t}^H e_{H,t}$. The optimality condition is given by

$$c_t = \frac{1-\vartheta}{\vartheta}e_{H,t} \tag{17}$$

which implies the following demand functions: $c_t = (1-\vartheta)C_t$ and $e_{H,t} = \vartheta C_t$, and from which we derive the following consumer price index (CPI, P_t)

$$P_t = (1-\vartheta)P_{c,t} + \vartheta P^H_{e,t} \tag{18}$$

and we refer to $1+\pi_{c,t} = P_{c,t}/P_{c,t-1}$ as the CPI core inflation rate. We assume that non-energy consumption is a Cobb-Douglas bundle of domestic ($c_{D,t}$) and foreign ($c_{F,t}$) goods:

$$c_t \propto c^{\varrho}_{D,t} c^{1-\varrho}_{F,t} \tag{19}$$

with prices $P_{c,t}$, $P_{D,t}$, and $P_{F,t}$, the latter denominated in domestic currency terms. The household solves $\max_{c_{D,t},c_{F,t}} c_t$ using equation (19) and the constraint $P_{c,t}c_t = P_{D,t}c_{D,t} + P_{F,t}c_{F,t}$. The resulting demand function for domestic goods is

$$c_{D,t} = \varrho \frac{P_{c,t}}{P_{D,t}} c_t \tag{20}$$

and the corresponding price index is given by

$$P_{c,t} = P^{\varrho}_{D,t} P^{1-\varrho}_{F,t}. \tag{21}$$

Imported goods are provided in competitive markets and the foreign currency price is normalized to one so that: $P_{F,t} = S_t$.

Finally, in line with the New Keynesian literature, we assume that the domestic non-energy consumption goods $c_{D,t}$ are given by a constant elasticity of substitution (CES) aggregation over a continuum of goods $c_{D,t}(i)$ indexed on the interval $i \in [0,1]$ with a sector-specific elasticity of substitution $\varepsilon > 1$ as follows: $c_{D,t} = \left(\int_0^1 c_{D,t}(i)^{\frac{\varepsilon-1}{\varepsilon}} di\right)^{\frac{\varepsilon}{\varepsilon-1}}$, where $c_{D,t}(i)$ represents the quantity of good i obtained by the household in period t. We define $P_{D,t}(i)$ as the price of the domestic non-energy consumption good i. It follows that the demand function for the individual good i is given by $c_{D,t}(i) = (P_{D,t}(i)/P_{D,t})^{-\varepsilon} c_{D,t}$ and the corresponding price index is $P_{D,t} = \left(\int_0^1 P_{D,t}(i)^{1-\varepsilon} di\right)^{\frac{1}{1-\varepsilon}}$.

3.4. **Intermediate goods.** Intermediate good j is produced with labor $h_t(j)$ and firm-level energy $e_{F,t}(j)$ under Leontief technology:

$$y_t(j) = \xi_A(\hat{\mathcal{T}}) \min\left\{\frac{h_t(j)}{1-\varsigma}, \frac{e_{F,t}(j)}{\varsigma}\right\}, \tag{22}$$

where ς is the fixed energy share. Given equation (22), marginal costs are given by

$$MC_t(\hat{\mathcal{T}}) = \frac{1}{\xi_A(\hat{\mathcal{T}})}\left((1-\varsigma)W_t + \varsigma P_{e,t}\right). \tag{23}$$

We assume perfectly competitive factor input markets, which implies that firms act as price takers in obtaining their input factors to production, hence nominal marginal costs are the same across firms $MC_t(j,\hat{\mathcal{T}}) = MC_t(\hat{\mathcal{T}})$. Intermediate goods are sold under monopolistic competition with Calvo price stickiness, described in detail in Galí (2015). Each firm re-optimizes its price with probability $1-\xi_p$, otherwise maintaining past pricing. The optimal reset price $\tilde{P}_{D,t}$ satisfies

$$E_t \sum_{k\geq 0} \xi_p^k \Lambda_{t,t+k} y_{D,t+k}(j)\left(\tilde{P}_{D,t} - \frac{\varepsilon}{\varepsilon-1} MC_{t+k}(\hat{\mathcal{T}})\right) = 0.$$

and the domestic price index evolves according to

$$P_{D,t} = \left(\xi_p P_{D,t-1}^{1-\varepsilon} + (1-\xi_p)\tilde{P}_{D,t}^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}. \tag{24}$$

3.5. **Policy and equilibrium.** Monetary policy follows a Taylor-type rule

$$i_t = \bar{\imath} + \phi_\pi \pi_t + \phi_y \hat{y}_t \tag{25}$$

with $i_t = \ln(R_t)$ and $\hat{y}_t$ denotes the (log-)deviation of output from its steady state. In equilibrium, energy supply equals energy demand $e_t = e_{H,t} + e_{F,t} = \vartheta C_t + \frac{\varsigma}{\xi_A(\hat{\mathcal{T}})} y_t$. Final output is $P_{D,t} y_t$ in nominal terms, while total absorption equals domestic and foreign consumption and investment adjusted for energy and non-energy imports. The resulting resource constraint reads

$$y_t + \frac{S_t P_{e,t}^*}{P_{D,t}} e_{f,t} = \varrho \frac{P_{c,t}}{P_{D,t}}\left[\left(1-\vartheta+\frac{P_{e,t}}{P_{c,t}}\frac{\vartheta}{\varrho}\right)C_t + I_t\right] + \frac{S_t}{P_{D,t}} X_t^*, \tag{26}$$

with foreign demand specified as $X_t^* = (P_{D,t}/S_t P^*)^{-\theta^*} Y_t^*$, where Y_t^* follows an AR(1) process with $u_{Y^*,t}$ being the foreign demand shock.

The balance of payments condition is

$$S_t(B_t^* - R^* B_{t-1}^*) = P_{D,t} y_t + S_t P_{e,t}^* e_{f,t} - P_{c,t}\left[\left(1-\vartheta+\frac{P_{e,t}}{P_{c,t}}\frac{\vartheta}{\varrho}\right)C_t + I_t\right]. \tag{27}$$

In equations (26) and (27), the term $P_t \psi_{B,t}$ is omitted since it cancels in log-linearized equilibrium under zero steady state net foreign assets.

Table 1. Calibration

Ident.	Value	Description	Range
β	0.971	Discount factor	$[0.95 - 0.99]$
ϕ	4	Inverse of Frisch labor supply elasticity	$[1 - 6]$
ξ_p	0.7	Degree of price stickiness	$[0.65 - 0.9]$
ψ_B	0.02	Adjustment costs for net foreign assets	$[0.01 - 0.1]$
ϱ	0.7	Share of domestically produced goods in domestic absorption	$[0.5 - 0.9]$
ϑ	0.08	Share of energy consumption in total household consumption	$[0.04 - 0.15]$
ς	0.12	Share of energy usage in production	$[0.05 - 0.2]$
ϕ_π	1.5	Reaction of monetary policy to inflation	$[1.1 - 2.5]$
ϕ_y	0.5	Reaction of monetary policy to output	$[0.0 - 0.9]$
ζ	0.4	Share of renewable energy	$[0.1 - 0.9]$
φ	0.02	Rate of decay of atmospheric carbon stock	$[0.01 - 0.1]$
κ	1.5	Elasticity of substitution between fossil and renewable raw energy	$[0.01 - 2.5]$
$\rho_{P_e^*}$	0.7	Auto-correlation of exogenous energy price shock	$[0.4 - 0.9]$
ρ_{Y^*}	0.7	Auto-correlation of exogenous global demand shock	$[0.4 - 0.9]$
θ^*	0.5	Elasticity of substitution between domestic and foreign goods in the foreign economy	$[0.1 - 1.1]$
α	0.03	Temperature anomaly sensitivity of TFP	$[0.01 - 0.06]$

3.6. **Baseline calibration and solution of the model.** We assume that trade is balanced in the steady state and that both foreign and domestic bonds are in zero net supply. The foreign price of imported goods (P_t^*) and the steady state price of imported energy ($P_{e,t}^*$) are normalized to one. The steady state shares of household consumption and investment in output are set to 0.95 and 0.05, respectively.

The shares of household energy in total consumption (ϑ) and of energy in firms' costs (ς) are calibrated from the weights of energy in the CPI and PPI, using averages across the countries considered in Section 4.1.4. The elasticity of substitution between fossil and renewable energy is set above unity, conforming with long-run estimates reported by Jo (2025). Table 1 details the calibration choices, which largely follow standard practice.

To calibrate ROW carbon emissions relative to output and fossil energy prices ($e_{f,t}^* = e_f^*(Y_t^*, P_{e,t}^*)$), we follow Doda (2014); Onofrei et al. (2022), who document that emissions are procyclical and more volatile than GDP. For instance, Doda (2014) estimates an elasticity of 1.6 for high-emission countries and 0.92 more broadly. In line with these findings, we set the elasticity between output and emissions to 1.2, and between fossil energy prices and emissions to -0.4, reflecting limited short run substitutability.

Equation (24), which governs the evolution of domestic goods' prices ($P_{D,t}$), and the first-order condition for the optimal price ($\tilde{P}_{D,t}$) yield the New Keynesian Phillips curve

$$\pi_{D,t} = \beta E_t \pi_{D,t+1} + \kappa_p \hat{mc}_t \tag{28}$$

where $\kappa_p = \frac{(1-\xi_p\beta)(1-\xi_p)}{\xi_p}$, and $mc_t = MC_t(\hat{\mathcal{T}})/P_{D,t}$ is real marginal costs. For later relevance, we note that the (log-)linearization implies that $\hat{mc}_t = dmc_t/\overline{mc}$, that is, the change in marginal costs[4] (dmc_t) relative to its steady state given by $\overline{mc} = \frac{\epsilon-1}{\epsilon}$. Alternatively, the change in marginal costs (dmc_t) can be proxied by the (absolute) deviation from the steady state: $dmc_t \approx mc_t - \overline{mc}$, which then implies $\hat{mc}_t \approx (mc_t - \overline{mc})/\overline{mc}$.

The complete list of the equations governing the equilibrium dynamics is provided in the Appendix.

3.7. **Equilibrium, model solution, and dynamic simulations.** We solve the model under rational expectations by computing a log-linearization around its steady state (see Appendix). This yields a first-order stochastic difference equation system in line with equation (3) with $\mathcal{P} = 1$, where $\boldsymbol{y}_t$ denotes the vector of endogenous variables and $\boldsymbol{u}_t = \left[u_{P_e^*,t}, u_{Y^*,t}\right]'$ is the vector of exogenous shocks. The assumption made in Section 3.2 implies that $\mathcal{T}_t$ is not part of the endogenous variables, but instead a structural parameter, hence shaping the coefficient matrices $\boldsymbol{B}_1(\hat{\mathcal{T}}_t)$ and $\boldsymbol{H}(\hat{\mathcal{T}}_t)$. In the subsequent analysis, we investigate how temperature anomalies modify the transmission of foreign demand and supply shocks.

3.7.1. *Temperature anomaly and the slope of the Phillips curve.* To examine the sensitivity of the Phillips curve's slope (given by $d\pi_{D,t}/dy_t$) with respect to the temperature anomaly, we first investigate the role of marginal costs. By substituting the real wage ($w_t = W_t/P_t$) and hours worked (h_t) into the expression for marginal costs, and using the household labor supply equation (14), the production function (22), and the equation for marginal costs (23), we obtain $mc_t = \left(\frac{1-\varsigma}{\xi_A(\hat{\mathcal{T}})}\right)^2 \chi y_t + \varsigma p_{e,t}$, under the assumption that $\sigma = 0$ (the household utility function is linear in consumption (C_t) and $p_{e,t} = P_{e,t}/P_{D,t}$. Further assuming $\zeta = 0$ (no renewable raw energy used) and $\kappa = 0$ (under $\zeta = 0$, the

[4]Note that dmc_t emerges once taking the total differential of equation (24) and the first-order condition for the optimal price ($\tilde{P}_{D,t}$) and simplifying.

assumption of $\kappa = 0$ has no further consequences), we can treat $p_{e,t}$ as a cost-push shock related to fossil energy. Considering the change in (real) marginal costs relative to the steady state, we derive $\frac{dmc_t}{\overline{mc}} = \left(\frac{1-\varsigma}{\xi_A(\hat{\mathcal{T}})}\right)^2 \frac{\chi\epsilon}{\epsilon-1} dy_t$, when $dp_{e,t} = 0$. This expression can be rearranged to yield $\frac{d(mc_t/\overline{mc})}{dy_t} = \left(\frac{1-\varsigma}{\xi_A(\hat{\mathcal{T}})}\right)^2 \frac{\chi\epsilon}{\epsilon-1}$, from which we derive the following slope sensitivity of marginal costs with respect to the temperature anomaly

$$\frac{d^2(mc_t/\overline{mc})}{dy_t d\hat{\mathcal{T}}} = 2\alpha\chi\frac{\epsilon}{\epsilon-1}\left(\frac{1-\varsigma}{\xi_A(\hat{\mathcal{T}})}\right)^2 > 0 \tag{29}$$

This equation highlights that the slope of the marginal cost curve increases with rising temperature anomalies.[5] Turning now to the (New Keynesian) Phillips curve, we first of all note that $\pi_{D,t} = \pi_{D,t} - \bar{\pi}_D \approx d\pi_{D,t}$ where the latter approximates the deviation of inflation from its steady state value ($\bar{\pi}_D = 0$). Holding expected inflation constant at its steady state level ($E_t d\pi_{D,t+1} = E_t\pi_{D,t+1} = 0$) we have that $\frac{d\pi_{D,t}}{d(mc_t/\overline{mc})} = \kappa_p$, and consequently

$$\begin{aligned}
\frac{d}{d\hat{\mathcal{T}}}\left(\left.\frac{d\pi_{D,t}}{dy_t}\right|_{\mathrm{PC}}\right) &= \frac{d}{d\hat{\mathcal{T}}}\left(\frac{d\pi_{D,t}}{d(mc_t/\overline{mc})}\frac{d(mc_t/\overline{mc})}{dy_t}\right) \\
&= \kappa_p\frac{d^2(mc_t/\overline{mc})}{dy_t d\hat{\mathcal{T}}} \\
&= 2\alpha\chi\kappa_p\frac{\epsilon}{\epsilon-1}\left(\frac{1-\varsigma}{\xi_A(\hat{\mathcal{T}})}\right)^2 > 0 \qquad (30)
\end{aligned}$$

This indicates that an increase in the temperature anomaly results in a steeper Phillips curve. This link arises from the assumption that elevated temperatures depress productivity, which, according to equation (29), amplifies the responsiveness of marginal costs to output. While this effect– higher temperature anomalies resulting in a steeper Phillips curve–originates from an apparently simple assumption about TFP, it has profound consequences: a steeper Phillips curve implies that inflation becomes more sensitive to output fluctuations, which directly conditions the propagation of (demand) shocks. A key parameter in this context is α, which measures the impact of the temperature anomaly on total factor productivity $\xi_A(\hat{\mathcal{T}})$, and consequently on output

[5] Notably, with a constant returns to scale production function, as is the case in our model, we have $\partial mc_t(y_t)/\partial y_t = 0$ (see for instance Mas-Colell et al., 1995, Chapter 5). This should not be confused with the expression presented in equation (29), which incorporates additional model-specific equations to account for equilibrium dynamics, ultimately leading to $\partial mc_t(y_t)/\partial y_t \neq 0$ due to feedback effects.

and marginal costs. This mechanism reflects the TFP channel. Another important parameter here is the Calvo price stickiness parameter ξ_p embedded in κ_p, which shapes the slope of the Phillips curve on top of α.[6]

Given that the implications of equation (30) are derived from a simplified framework, the subsequent section considers a more general setting.

3.7.2. *Theoretical implications of temperature anomaly for demand and supply shocks.* We now assess how temperature anomalies influence the transmission of demand and supply shocks through quantitative simulations of the model with the results shown in Figure 3. The sub-panels in the top row show the impulse response functions (IRFs) for inflation and output in response to the foreign demand shock ($u_{Y^*,t}$), while those in the bottom row show the IRFs to the foreign supply shock ($u_{P_e^*,t}$). Both shocks are expansionary. The demand shock is specified such that Y_t^* increases by one percent, while the supply shock is specified such that $P_{e,t}^*$ decreases by one percent. Both shocks raise domestic output while the inflation response is opposite.

Most importantly, since the IRFs are continuous functions of $\hat{\mathcal{T}}$, we can plot them for various values of $\hat{\mathcal{T}}$. The black dotted lines in the first two sub-panels of Figure 3 show the IRFs when $\hat{\mathcal{T}} = 0$, while the blue dashed lines denote those when $\hat{\mathcal{T}} > 0$. The first scenario refers to a "normal" climate state, while the second corresponds to one of above-average temperatures. In response to the demand shock, the IRFs of inflation and output show a high sensitivity to $\hat{\mathcal{T}}$: inflation is higher when $\hat{\mathcal{T}} > 0$, while output is lower. In other words, the output effects diminish at higher temperatures, while the inflation effects increase.

The fourth sub-panel in the first row of Figure 3 provides a robustness check for a more general parameter space. In particular, we simulate the model over a wide range of different values for the structural parameters. To this end, we attach a uniform distribution to each parameter and define upper and lower bounds as indicated in the fourth column (Range) in Table 1. We simulate the model 2000 times[7] and show the difference of the IRFs for the following

[6]To underscore the simplicity of the underlying intuition, we also derive the mechanism in the Appendix using the AS-AD model following Blanchard and Johnson (2017).

[7]The implementation is as follows: for the same draw of the structural parameters, we solve the model for (i) the scenario $\hat{\mathcal{T}} = 0$, and (ii) the scenario $\hat{\mathcal{T}} > 0$. We then compute IRFs for each scenario, of which we then compute the difference in the IRFs across the two scenarios, as depicted in Figure 3. By this procedure we can uniquely attach the difference

Figure 3. Temperature sensitivity of demand and supply shocks

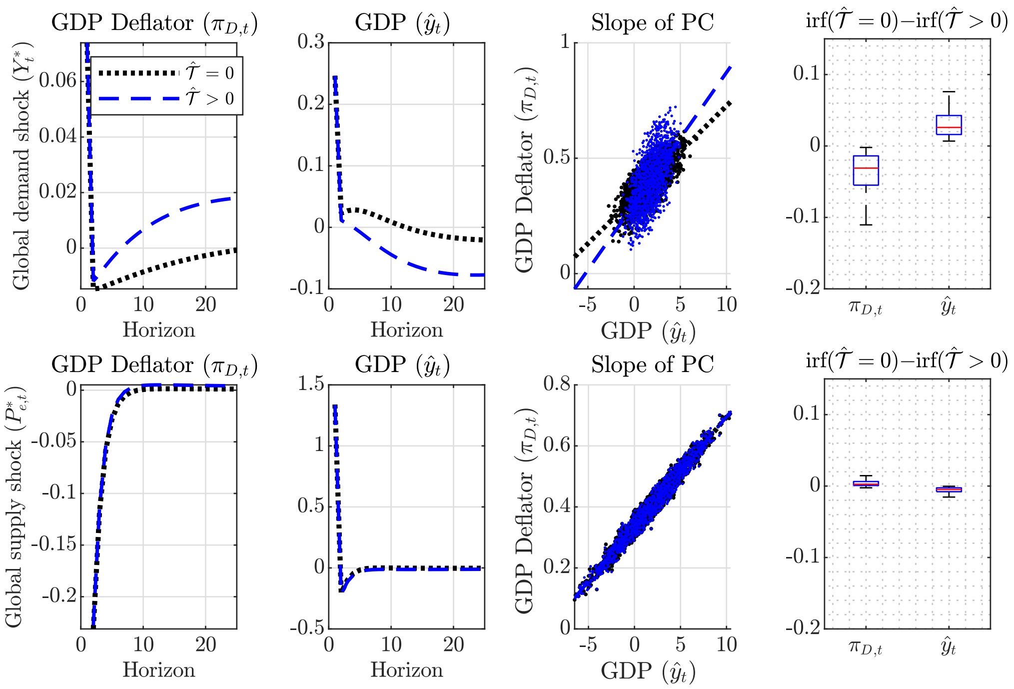

Note: The figure shows the sensitivity of the IRFs of inflation and output to a global demand (upper sub-panels) and a global supply shock (lower sub-panels) with respect to the temperature anomaly. The sub-panels in the fourth column show the sensitivity with respect to a more general calibration (average effect over the first 6 steps of the IRFs). The red line in the boxplot refers to the median.

scenario: $\hat{\mathcal{T}} = 0$ relative to $\hat{\mathcal{T}} > 0$. We consider the average effect over the first six steps of the IRFs. The difference between the two scenarios is depicted in the right sub-panel in Figure 3. We observe a pronounced sensitivity of the model's variables to the temperature anomaly in the case of the demand shock. When $\hat{\mathcal{T}} > 0$, the impact of demand shocks on inflation is more pronounced, while the output response is muted. We perform the same analysis for the supply shock, the results of which are presented in the fourth sub-panel of the second row. As can be observed, the sensitivity of the IRFs to $\hat{\mathcal{T}} > 0$ is markedly lower in comparison to the demand shock. Most notably, the boxes largely overlap with zero, indicating that there is no unique effect regarding the extent to which $\hat{\mathcal{T}}$ shapes the IRFs in the case of a supply shock.

in the response to changes in $\hat{\mathcal{T}}$, while at the same time preserving flexibility in the model calibration over a wide parameter space.

This pattern is also evident in the two sub-panels of the third column, which depict the slope of the Phillips curve.[8] The upper sub-panel illustrates its sensitivity to the demand shock, while the lower sub-panel corresponds to the supply shock. As shown, the temperature anomaly plays a crucial role in shaping the transmission of the demand shock: a higher temperature anomaly leads to a steepening of the Phillips curve, resulting in a stronger inflationary response but a weaker output reaction. This effect, however, is absent in the case of the supply shock, as the temperature anomaly influences only the supply side of the model (see equation (30)).

4. The econometric model

We validate our theoretical results using an empirical model. To this end, we use the solution format of the theoretical model and extend it in several steps to create an estimable empirical model. We then examine the influence of the temperature anomaly on the transmission mechanism of global demand and supply shocks.

4.1. **From the theoretical to the empirical model.** The solution of the theoretical model is given by equation (3) with $\boldsymbol{B}(L;\mathcal{T}) = \boldsymbol{I} - \boldsymbol{B}(\mathcal{T})L$ where $\boldsymbol{I}$ is the identity matrix, $\boldsymbol{B}(\mathcal{T})$ a matrix whose coefficients depend on $\mathcal{T}$ and $\boldsymbol{u}_t$ is the vector of structural shocks defined in Section 3.7.

We partition the vector of endogenous variables $\boldsymbol{y}_t = [\boldsymbol{y}'_{c,t},\, \boldsymbol{y}'_{g,t}]'$ to examine more clearly the solution of the theoretical model. The vector $\boldsymbol{y}_{g,t} = [P^*_{e,t}, Y^*_t]'$ collects the foreign/global variables, while $\boldsymbol{y}_{c,t} = [\ldots,\, \pi_t,\, \hat{y}_t,\, i_t]'$ (dimension $n \times 1$) the home-country variables. This partitioning motivates the following form of equation (3) from the theoretical model

$$(31) \qquad \begin{bmatrix} \underset{(n\times 1)}{\boldsymbol{y}_{c,t}} \\ \underset{(2\times 1)}{\boldsymbol{y}_{g,t}} \end{bmatrix} = \begin{bmatrix} \underset{(n\times n)}{\boldsymbol{B}_{11}(\mathcal{T})} & \underset{(n\times 2)}{\boldsymbol{B}_{12}} \\ \underset{(2\times n)}{\boldsymbol{0}} & \underset{(2\times 2)}{\boldsymbol{B}_{22}} \end{bmatrix} \begin{bmatrix} \underset{(n\times 1)}{\boldsymbol{y}_{c,t-1}} \\ \underset{(2\times 1)}{\boldsymbol{y}_{g,t-1}} \end{bmatrix} + \begin{bmatrix} \underset{(n\times 2)}{\boldsymbol{H}_1(\mathcal{T})} \\ \underset{(2\times 2)}{\boldsymbol{H}_2} \end{bmatrix} \underset{(2\times 1)}{\boldsymbol{u}_t}$$

which is recognized as a structural VAR (SVAR) representation of the theoretical model. The ensuing sections delineate three steps for establishing the (estimable) empirical model from the theoretical counterpart.

[8]To produce these two sub-panels, we followed Gasteiger and Grimaud (2023) and simulated the model for 2000 periods under the demand shock (upper sub-panel) only and the supply shock (lower sub-panel) only.

4.1.1. *Block exogeneity: the global–domestic relationship.* Equation (31) exhibits a block-exogenous structure, because the lower-left block of the coefficient matrix is zero. This implies that home-country variables do not influence global dynamics, while global shocks do propagate to the home-country through $\boldsymbol{B}_{12}$. Block exogeneity is imposed by our theoretical model but remains an empirically testable restriction. In our estimation, it is consistently supported by information criteria.[9]

4.1.2. *Interaction term: temperature effects in the country block.* The elements of the matrix $\boldsymbol{B}_{11}(\mathcal{T})$ are functions of the temperature anomaly $\mathcal{T}$. However, we cannot identify an explicit functional form from the solution of the theoretical model. Hence we approximate $\boldsymbol{B}_{11}(\mathcal{T})$ using a first-order Taylor expansion, as proposed in equation (4) which implies

$$\text{IVAR:} \quad \boldsymbol{B}_{11}(\mathcal{T}_t) = \boldsymbol{\Gamma}(\bar{\mathcal{T}}) + \boldsymbol{\mathcal{F}}\hat{\mathcal{T}}_t \tag{32}$$

where $\boldsymbol{\Gamma}$ and $\boldsymbol{\mathcal{F}}$ are $n \times n$ matrices and we substituted $\mathcal{T}$ with its empirical counterpart $\mathcal{T}_t = \bar{\mathcal{T}} + \hat{\mathcal{T}}_t$. Extending equation (32) with $\boldsymbol{y}_{c,t-1}$ gives

$$\boldsymbol{B}_{11}(\mathcal{T}_t)\boldsymbol{y}_{c,t-1} = \left(\boldsymbol{\Gamma}(\bar{\mathcal{T}}) + \boldsymbol{\mathcal{F}}\hat{\mathcal{T}}_t\right)\boldsymbol{y}_{c,t-1} \tag{33}$$

where $\hat{\mathcal{T}}_t\boldsymbol{y}_{c,t-1}$ establishes an interaction term that gives rise to an interacted vector autoregressive (IVAR) model (compare Towbin and Weber, 2013). Its response coefficients are allowed to change deterministically with $\hat{\mathcal{T}}_t$

$$\frac{\partial \boldsymbol{B}(\mathcal{T}_t)}{\partial \hat{\mathcal{T}}_t} = \begin{bmatrix} \boldsymbol{\mathcal{F}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{34}$$

Notably, equation (34) demonstrates that–given the assumptions made in the theoretical model–only the domestic propagation of shocks depends on temperature anomalies, captured by the matrix $\boldsymbol{\mathcal{F}}$. In contrast, the global block ($\boldsymbol{H}_2$, $\boldsymbol{B}_{22}$, $\boldsymbol{B}_{12}$) is assumed to be unaffected by temperature anomalies hence enabling an isolation of the impact of temperature variations. This ensures that global shocks are identified independently of temperature anomalies. Consequently, any differences in the country-specific propagation of the global shocks

[9] In the empirical model, we relax the restriction that $\boldsymbol{B}_{22}$ is diagonal (as in the theoretical model) and allow it to be fully unrestricted. Similarly, some rows of $\boldsymbol{B}_{11}(\mathcal{T})$ and $\boldsymbol{B}_{12}$ are identically zero for non-state variables, consistent with the theoretical structure, which is also relaxed in the empirical counterpart.

can then be unequivocally attributed to the temperature anomaly, which is precisely the objective of our analysis.

Two remarks are in order. First, the interaction variable $\hat{\mathcal{T}}_t$ is treated as exogenous with respect to business cycle shocks. Since we employ quarterly temperature anomalies, this assumption is well justified: while temperature anomalies exhibit temporal variation and are ultimately affected by long-run economic activity through carbon emissions, they are unlikely to react to global demand and supply shocks in the short run. Second, equation (32) imposes a first-order Taylor approximation on the temperature effects, in contrast to the general form embodied in equation (3), which allows for higher-order terms, and equation (9) which explicitly assumes non-linearity. We address the potential of higher-order terms in robustness checks in the Appendix.

4.1.3. *Identification of the shocks.* Equation (31) specifies a structural system. However, estimation is limited to its reduced-form representation with the reduced-form error term $\boldsymbol{\varepsilon}_t$ being a linear combination of the structural shocks $\boldsymbol{\varepsilon}_t = \boldsymbol{H}(\mathcal{T})\boldsymbol{u}_t$, with $E(\boldsymbol{u}_t\boldsymbol{u}_t') = \text{diag}\left(\sigma^2_{P_e^*}, \sigma^2_{Y^*}\right)$. Since $\boldsymbol{H}(\mathcal{T})$ is not uniquely determined from reduced-form variances, additional restrictions are necessary.

We identify the shocks by imposing sign restrictions derived from the theoretical model. To ensure that these restrictions are not tied to a specific calibration, we adopt a simulation-based approach and simulate the theoretical model across a wide range of admissible values of the structural parameters. To this purpose, we re-do the exercise outlined in Section 3.7.2 used to establish the sub-panels in the fourth column in Figure 3, however, now applied solely to the impact response of the global demand and supply shock, which allows us to examine the matrix $\boldsymbol{H}(\mathcal{T})$. This yields two key results. First, considering the 1st to 99th percentiles of the simulated entries of $\boldsymbol{H}(\mathcal{T})$, we find that

$$\text{sgn}(\boldsymbol{H}(\mathcal{T} > 0)) = \text{sgn}(\boldsymbol{H}(\mathcal{T} = 0)) \tag{35}$$

implying that the signs of the contemporaneous effects of the global shocks are not influenced by the temperature anomaly (see also the Appendix). Consequently, the temperature anomaly $\mathcal{T}$ can be excluded when determining the sign restrictions. Second, the same analysis provides the following sign restrictions for identifying a contractionary global supply shock (first column of $\boldsymbol{H}$) and an expansionary global demand shock (second column of $\boldsymbol{H}$) for inflation, output and the nominal rate in the home-country and inflation and output at the global level (noting the ordering of the variables as outlined in Section 4.1 and we list only those variables from the theoretical model that are subsequently used in the empirical estimation)

$$\text{sgn}(\boldsymbol{H}) = \begin{array}{r} \\ \vdots \\ \pi_t \\ \hat{y}_t \\ i_t \\ P^*_{e,t} \\ Y^*_t \end{array} \begin{array}{c} \begin{array}{cc} u_{P^*_e,t} & u_{Y^*,t} \end{array} \\ \left[\begin{array}{cc} \vdots & \vdots \\ \geq 0 & \geq 0 \\ \leq 0 & \geq 0 \\ \cdot & \geq 0 \\ \hline \geq 0 & \geq 0 \\ \leq 0 & \geq 0 \end{array}\right] \end{array} \tag{36}$$

where the horizontal line separates home-country (top) from global variables (bottom). According to these restrictions, an expansionary global demand shock is expected to increase inflation and output both domestically and globally, and raise the domestic nominal rate. A contractionary global supply shock raises inflation and lowers output in both dimensions, but its effect on the policy rate remains ambiguous.

For implementation, let $\boldsymbol{\Omega} = E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t')$ denote the reduced-form variance-covariance matrix. By eigenvalue decomposition, we write $\mathbf{V} = \mathbf{E_\Omega}\boldsymbol{\Lambda}_{\boldsymbol{\Omega}}^{1/2}$, with $\mathbf{E_\Omega}$ denoting eigenvectors and $\boldsymbol{\Lambda}_{\boldsymbol{\Omega}}$ eigenvalues. Candidate structural impact matrices are then obtained as $\mathbf{H} = \mathbf{VQ}$, where $\mathbf{Q}$ is a random orthonormal matrix. Following Rubio-Ramírez et al. (2010), $\mathbf{Q}$ is generated by QR decomposition[10] of a Gaussian random matrix. If a candidate $\mathbf{H}$ satisfies the sign restrictions (36), it is retained; otherwise, a new draw of $\mathbf{Q}$ is generated. This procedure is repeated until 5,000 accepted draws are obtained. In the empirical analysis, we report the median estimates of IRFs, along with the 16th and 84th percentiles as credible bands.

4.1.4. *General matters: variables, data and estimation.* We employ gross domestic product (GDP) at constant prices as our measure of output, the GDP

[10]We note that while the estimation and the inference of the reduced-form model rely on a rather non-informative prior for all coefficients, an informative prior emerges in the structural model due to the QR decomposition (Baumeister and Hamilton, 2015).

deflator as our measure of prices, and the three-month government bond rate as the short term interest rate in the vector of home-country variables ($\boldsymbol{y}_{c,t}$), resulting in $n = 3$. The vector of global variables ($\boldsymbol{y}_t^g$) comprises global GDP and a global commodity price index. Output and prices are included in logarithmic four-quarter differences throughout, allowing the resulting series to be interpreted as year-over-year growth rates.[11] The short term interest rate is included in levels.

To capture climatic effects, we utilize the quarterly temperature anomaly ($\hat{\mathcal{T}}_t$) at the country level, defined as deviation from each country's historical average temperature. This approach enables us to account for country-specific climatic variation and effectively control for regional climatic differences. By focusing on temperature anomalies we emphasize the role of relative temperature variations rather than absolute levels. This strategy ensures that our empirical analysis aligns closely with the theoretical framework, which is predicated on the impact of departures from typical climatic conditions, rather than on uniform shifts in temperature across all countries (see also Nath et al., 2024).

Our empirical analysis is based on a quarterly dataset spanning the period 1961Q1 to 2023Q4 and comprising $N = 20$ countries: Australia, Austria, Belgium, Canada, Finland, France, Greece, India, Italy, Japan, Mexico, the Netherlands, Norway, New Zealand, Portugal, South Korea, Spain, Sweden, the United Kingdom, and the United States (see the Appendix for further details).

The selection of countries is determined by the availability of quarterly data for output, prices, and a short term nominal interest rate. We use quarterly data because they better reflect the short run dynamic response of an economy to global demand and supply shocks. In addition, typical supply shocks–such as large swings in commodity or energy prices–could be averaged out at the annual frequency.

[11]This specification is advantageous as it ensures that stationary variables are employed in the VAR model. This facilitates the inclusion of interaction terms as it contributes to avoid the risk of the companion matrix having eigenvalues outside the unit circle. Moreover, Nath et al. (2024) provide an examination of whether temperature changes affect the level or the growth rate of GDP and argue that global warming exerts persistent, but not permanent, effects on economic growth. This finding supports our decision to use growth rates.

The model, consisting of equations (31) and (32), is estimated individually for each country using Bayesian methods.[12] We employ a natural conjugate prior, calibrated to be relatively non-informative, and jointly sample all parameters from the posterior distribution (see Appendix for further details). Moreover, while equation (31) has only one lag, the empirical implementation accommodates an extended lag structure. We incorporate truly exogenous variables into the model, these are, a constant term, a dummy variable to account for the economic turbulence associated with the Covid-19 pandemic, and a disaster-specific variable which serves to capture climate-related disruptions that are not fully spanned by temperature anomalies.[13] We rule out a feedback from local natural disasters to global variables by imposing zero restrictions on the relevant parameters.

4.2. **The general hypothesis: IVAR or VARX.** The IVAR model encapsulates the concept presented in the specification outlined by equation (3) and posits that climate change (proxied by the temperature anomaly) shapes the propagation mechanism of exogenous shocks. This perspective is consistent with our interpretation of the temperature anomaly data and aligns with the relatively minor role of temperature anomalies in accounting for macroeconomic fluctuations (see Ciccarelli and Marotta, 2024, among others). Nonetheless, the empirical specification we propose differs from the alternative considered in the literature, which corresponds to equation (2). The latter gives rise to the following VARX specification

$$\text{VARX:}\quad \boldsymbol{B}_{11}(\mathcal{T}_t)\boldsymbol{y}_{c,t-1} = \boldsymbol{\Gamma}(\bar{\mathcal{T}})\boldsymbol{y}_{c,t-1} + \boldsymbol{\gamma}\hat{\mathcal{T}}_t \tag{37}$$

[12]Our decision to use a VAR framework—specifically, an interacted panel VAR model—rests on two main considerations. First, to ensure comparability and continuity with the literature we critique and build upon—which universally adopts VAR methods—we likewise employ a VAR framework. Second, the (log-)linearised solution of the DSGE model takes the form of a state-dependent VAR, making it the natural model-consistent specification; the joint sign restrictions derived from the theoretical model can be imposed directly on the reduced-form innovation covariance matrix. We acknowledge that Olea et al. (2024) identify conditions under which local projections may outperform VARs. We therefore implement local projections as a complementary robustness exercise in the Appendix, where the results are qualitatively consistent with the baseline findings.

[13]Our natural disaster measure captures country-specific damages (expressed in national currency units relative to nominal GDP) associated with hydrological, meteorological, and climatological events and is drawn from the EM-DAT database, which builds on the classification framework developed by Sapir and Misson (1992). Additional details are provided in the Appendix.

where this equation is now to be used in equation (31) instead of the IVAR specification put forth in equation (33). While the IVAR model nests the VARX model, the key difference concerns the interaction between the temperature anomaly and the vector of endogenous variables in the IVAR model. To statistically differentiate between these specifications, we employ the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), and the Hannan-Quinn Criterion (HQ) and the likelihood-ratio test. Both specifications utilize the same set of home-country $\boldsymbol{y}_{c,t}$ and global variables $\boldsymbol{y}_{g,t}$.[14] The results are presented in Table 2, which lists the preferred specification (IVAR($\mathcal{P}$) or VARX($\mathcal{P}$)) for each information criterion and for each country. Across the countries considered, the information criteria favor the IVAR specification over the VARX specification. Moreover, for all countries, at least one information criterion favors the IVAR model.

According to the HQ information criterion, we choose a lag length of three, as this is the lag length selected by the HQ for most countries. Based on this choice, we employ the likelihood-ratio test (see Hamilton, 1994, Chapter 11.1) to conduct three tests for each country, where we compare the IVAR(3) model with the (i) VARX(3), (ii) VARX(2), and (iii) VARX(4) models. The test results, presented in Table 2, indicate a preference for the IVAR specification, hence substantiating the results of the information criteria. We interpret these findings in favor of our *general hypothesis* and hence of our proposed model specification as outlined by equation (3).

4.3. **The particular hypothesis: the role of climate change.** We now consider IRFs and forecast error variance decompositions (FEVDs). In our context, these statistics are functions of the temperature anomaly $\hat{\mathcal{T}}$. Therefore, the IRFs and FEVDs can be evaluated at different levels of temperature anomalies. This will be used to address our particular hypothesis.

Since we have a set of $c = 1, ..., N$ countries, we consider a panel estimator of the IVAR model for the subsequent analysis. This gives rise to an interacted panel vector autoregressive (IPVAR) model. We use the mean-group estimator of Pesaran and Smith (1995), which was used in an IPVAR setting by Towbin

[14]We calculate their values for each specification for various lags $\mathcal{P}$, considering a maximum of $\mathcal{P}_{\max} = 14$ lags.

TABLE 2. Model specification: IVAR vs. VARX

				Likelihood-ratio test		
				IVAR(3) vs.:		
	AIC	BIC	HQ	VARX(3)$^{(1)}$	VARX(2)$^{(2)}$	VARX(4)$^{(3)}$
Australia	IVAR (12)	VARX (1)	IVAR (1)	IVAR	IVAR	IVAR
Austria	IVAR (7)	IVAR (1)	IVAR (3)	IVAR	IVAR	IVAR
Belgium	IVAR (9)	IVAR (1)	IVAR (2)	IVAR	IVAR	IVAR
Canada	VARX (9)	IVAR (1)	IVAR (1)	IVAR	IVAR	IVAR
Finland	IVAR (11)	IVAR (1)	IVAR (1)	IVAR	IVAR	VARX
France	IVAR (12)	IVAR (1)	IVAR (1)	IVAR	IVAR	IVAR
Greece	VARX (5)	IVAR (1)	IVAR (1)	IVAR	IVAR	IVAR
India	IVAR (9)	IVAR (1)	IVAR (1)	IVAR	IVAR	IVAR
Italy	IVAR (6)	IVAR (1)	IVAR (2)	IVAR	IVAR	IVAR
Japan	IVAR (10)	IVAR (1)	VARX (5)	IVAR	VARX	IVAR
S. Korea	IVAR (12)	IVAR (1)	IVAR (3)	IVAR	IVAR	IVAR
Mexico	IVAR (12)	IVAR (2)	IVAR (3)	IVAR	IVAR	IVAR
Netherlands	IVAR (12)	IVAR (1)	IVAR (1)	VARX	IVAR	IVAR
Norway	VARX (8)	IVAR (1)	VARX (6)	IVAR	IVAR	IVAR
N. Zealand	IVAR (12)	IVAR (1)	IVAR (1)	IVAR	IVAR	IVAR
Portugal	IVAR (12)	IVAR (2)	IVAR (3)	IVAR	IVAR	IVAR
Spain	IVAR (12)	IVAR (1)	IVAR (2)	IVAR	IVAR	IVAR
Sweden	IVAR (12)	IVAR (1)	IVAR (1)	IVAR	VARX	IVAR
UK	IVAR (11)	IVAR (1)	IVAR (1)	VARX	IVAR	IVAR
USA	IVAR (9)	IVAR (1)	IVAR (3)	IVAR	IVAR	IVAR
VARX	3	1	2	2	2	1
IVAR	17	19	18	18	18	19

Note: In case of the information criteria (AIC, BIC, HQ), the values in parentheses indicate the optimal lag length. In case of the likelihood-ratio test, the values in parentheses refer to the specified lag-length. The last two rows indicate the number of times the IVAR or VARX specification was suggested by the information criteria and the likelihood-ratio tests.
(1) The test specification is H_0: VARX(3) vs. H_1: IVAR(3).
(2) The test specification is H_0: VARX(2) vs. H_1: IVAR(3).
(3) The test specification is H_0: VARX(4) vs. H_1: IVAR(3).

and Weber (2013). This panel estimator provides a reasonable control for possible omitted variables, as for instance institutional quality, as highlighted by Hassler et al. (2016), role of labor market institutions, the shares of tradables and nontradables in production, the exchange rate regime and other structural elements. The mean-group estimator of the home-country block in the IPVAR model is given by

$$\overline{\boldsymbol{B}}_j(\hat{\mathcal{T}}) = \sum_{c=1}^{N} \boldsymbol{B}_j^c(\hat{\mathcal{T}})/N \qquad \forall j = 1, ..., \mathcal{P}. \tag{38}$$

which also applies to $\boldsymbol{B}_{12} \to \boldsymbol{B}_{12}^c$, $\boldsymbol{H}_1 \to \boldsymbol{H}_1^c$, $\boldsymbol{\Gamma}_j \to \boldsymbol{\Gamma}_j^c$ and $\boldsymbol{\mathcal{F}}_j \to \boldsymbol{\mathcal{F}}_j^c$ as of equation (32), and we do so across all draws of the posterior distribution of the respective matrices. We compute IRFs and FEVDs using these averaged coefficient matrices and interpret them as an "average home-country" response.

4.3.1. *Empirical implications of temperature anomaly for demand and supply shocks.* The IRFs for the global demand shock are shown in Figure 4 as black solid lines (median of the posterior distribution) and blue dashed lines representing the 68 percent credible bands. The shock is normalized so that global GDP growth increases by one percentage point on impact, which translates into an output growth increase at the (average-)country level of approximately three percentage points. This expansion is accompanied by a rise in inflation which gradually vanishes. These dynamics are significantly influenced by temperature anomalies, $\hat{\mathcal{T}}$. The temperature anomaly variable is set to $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$. These values correspond to the 5th and 95th percentiles of the distribution of the temperature anomaly over time (quarterly frequency) and over countries. When the temperature anomaly is low (sub-panels in the first column), the increase in output is larger than when the temperature anomaly is high (sub-panels in the second column), while the opposite effect is observed for inflation. The differences in the responses are statistically significantly different from zero (sub-panels in the third column). These findings align with the theoretical implications discussed in Section 3.7.2.

While Figure 4 illustrates the path of the IRFs over a long horizon (up to the 16th quarter), it only presents the IRFs for two specific values of the temperature anomaly variable. To better understand the sensitivity of the IRFs to $\hat{\mathcal{T}}$, we extend this analysis by also showing the effects at a particular horizon with respect to various levels of $\hat{\mathcal{T}}$. The results are displayed in Figure 5. The blue line in the first sub-panel (with the gray-shaded area representing the 68 percent credible bands) illustrates the median effect of the global demand shock on the (average-)country inflation at a four-quarter horizon across different levels of temperature anomalies. As evident, the impact of the global demand shock on inflation increases as the temperature anomaly rises, indicating that demand shocks have a larger effect on inflation when the temperature anomaly is high. Conversely, the effect on (average-)country output diminishes as the temperature anomaly increases, as shown by the downward-sloping black solid line in the second sub-panel, which shows the effect of the global demand shock on (average-)country output at the four-quarter horizon

FIGURE 4. Demand shock

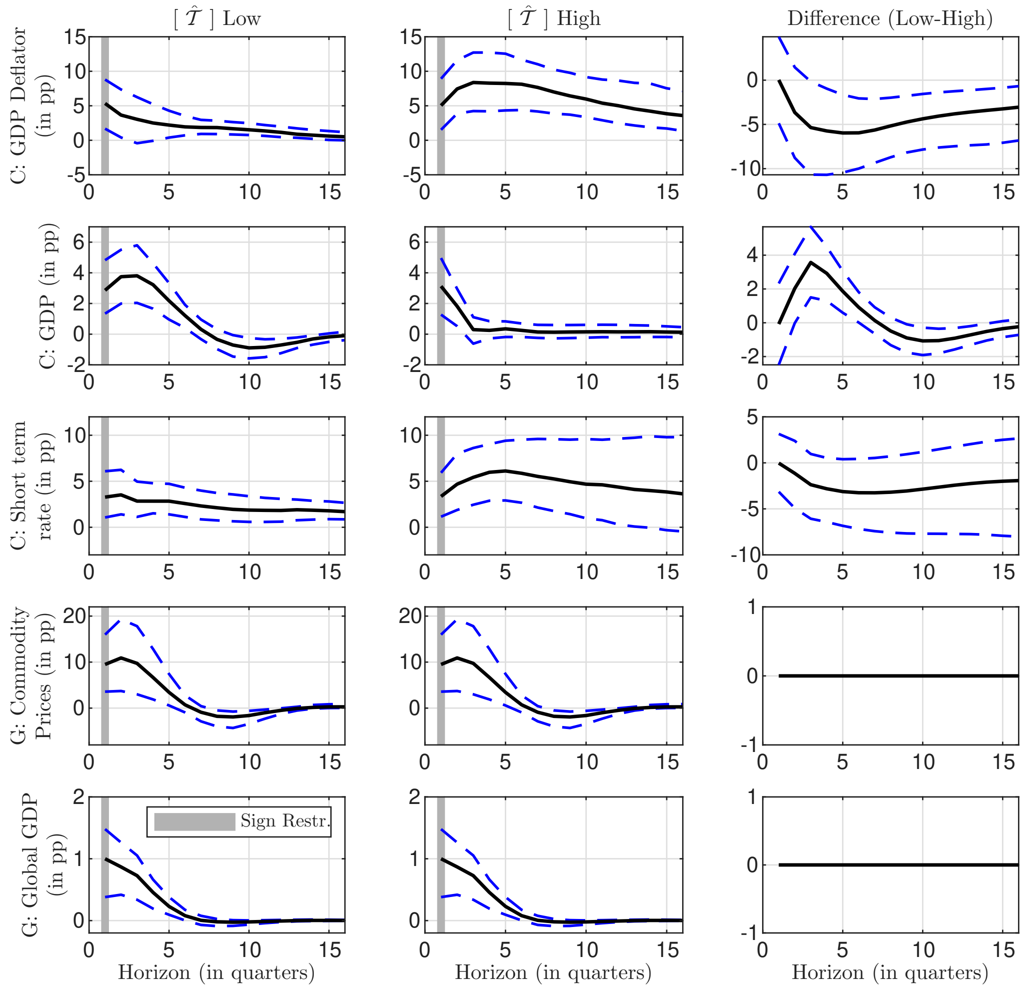

Note: The figure shows the sensitivity of the IRFs of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the temperature anomaly. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$. The solid line represents the posterior median, while the dashed lines denote the 68 percent credible interval.

for varying levels of $\hat{\mathcal{T}}$. The various dotted lines represent horizons other than the four-quarter horizon. In each case, they align with the path of the blue solid line (four-quarter horizon), thereby confirming the results pertaining to shorter and longer horizons.

These empirical findings corroborate the theoretical results, indicating that the temperature anomaly exerts a pronounced influence on the transmission mechanism of demand shocks. In particular, the findings indicate that elevated temperature anomalies increase the slope of the Phillips curve, thereby engendering more pronounced inflation effects and relatively muted output effects

Figure 5. Sensitivity of IRFs on temperature anomalies

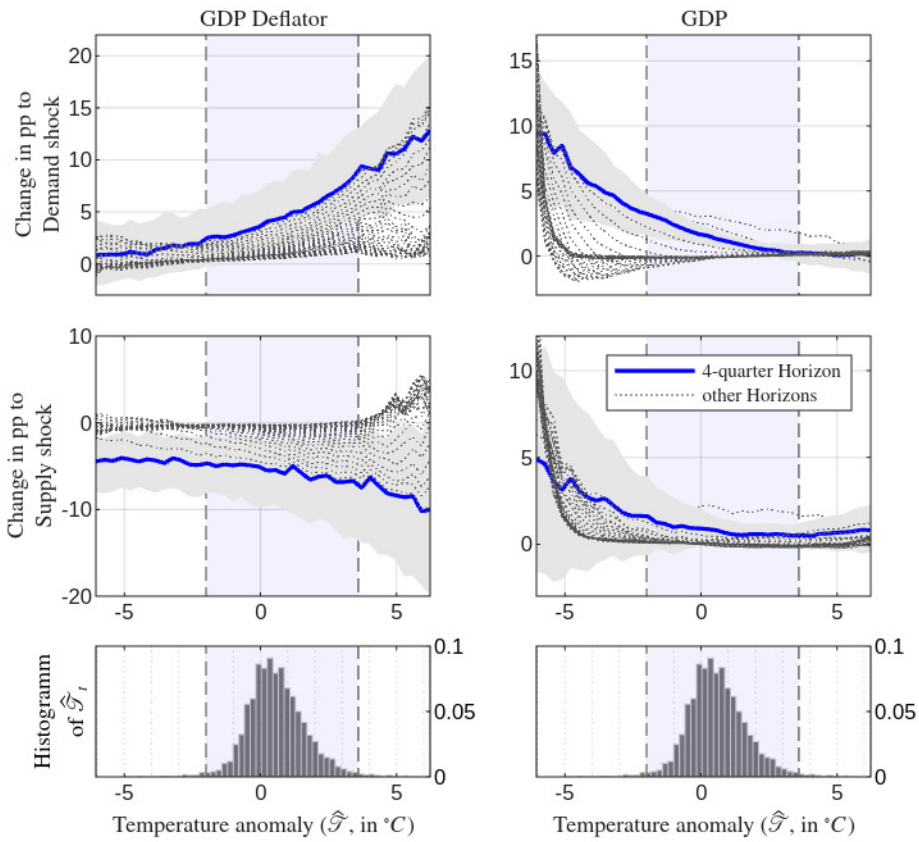

Note: The figure shows how the IRFs (at a particular horizon) vary with the temperature anomalies. The blue lines are the (median) IRFs at the four-quarter horizon (with the 68 percent credible interval given by the gray-shaded area), while the dotted lines represent IRFs at horizons other than the four-quarter horizon. The third sub-panel displays the histogram of the temperature anomaly across time (quarterly frequency) and countries and the vertical lines indicate the 5th and 95th percentiles.

in response to demand shocks. This argument is substantiated by the results of the supply shock, as illustrated in Figure 6. As can be observed, the two shock-assigning variables (global GDP and global commodity prices) exhibit opposing movements. We observe an immediate spillover effect on (average-)country inflation, which declines accordingly. The most notable finding here is the insensitivity of the IRFs to temperature anomalies. As the sub-panels in the third column of Figure 6 illustrate, there is no statistically significant evidence that the IRFs differ between a high and a low value of the temperature anomaly. This result is not specific to the two values of $\hat{\mathcal{T}}$ used in Figure 6. Instead, it is evident across a range of values, as demonstrated in the sub-panels of the second row in Figure 5. The blue lines again quantify the impact of the supply shock on (average-)country inflation and output at

FIGURE 6. Supply shock

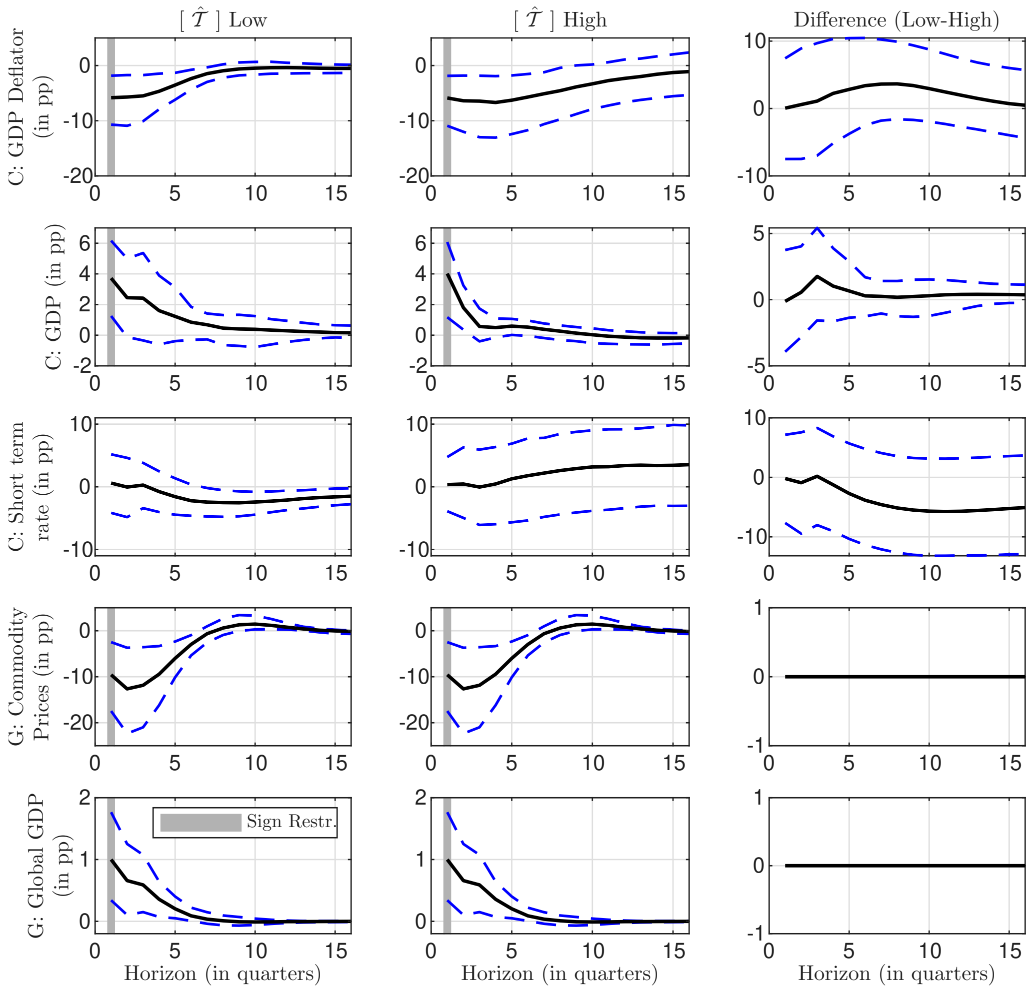

Note: The figure shows the sensitivity of the IRFs of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the temperature anomaly. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$. The solid line represents the posterior median, while the dashed lines denote the 68 percent credible interval.

the four-quarter horizon. This impact is illustrated for various levels of the temperature anomaly. As can be observed, the impact on prices and output remains largely constant (prices fall while output rises), but there is no indication that the observed point estimates differ in a statistically significant manner across low and high values of the temperature anomaly, as the credible bands (gray-shaded area) consistently overlap.

Table 3 summarizes the quantitative effects. It details the impact of a rise in the growth rate of global GDP following an expansionary global demand shock, where the shock is normalized such that the average effect over the first

TABLE 3. Global demand shock: Effects over the one-year horizon

	Country			Global	
	Prices$^{(1)}$ (in pp)	GDP$^{(2)}$ (in pp)	Rate$^{(3)}$ (in pp)	Prices$^{(4)}$ (in pp)	GDP$^{(2)}$ (in pp)
$\mathcal{T} = 0$ °C	5.8 *	3.4 *	5.7 *	13.4 *	1.0 *
$\mathcal{T} = 1$ °C	7.0 *	2.8 *	5.8 *	13.4 *	1.0 *
$\mathcal{T} = 2$ °C	8.0 *	1.9 *	6.3 *	13.4 *	1.0 *
Differences					
$\mathcal{T}_1 - \mathcal{T}_0$	1.2	-0.6	0.1	0.0	0.0
$\mathcal{T}_2 - \mathcal{T}_0$	2.2	-1.5 *	0.7	0.0	0.0

Notes: The asterisks (*) indicate values that are statistically significantly different from zero at the 16 percent level.
(1) The GDP deflator is used in year-over-year relative changes.
(2) GDP is used in year-over-year relative changes.
(3) A short term nominal interest rate is used.
(4) A commodity price index is used in year-over-year relative changes.

four quarters on the global GDP growth rate (year-over-year) is one percentage point. Over the same one-year horizon, global commodity price inflation (year-over-year) increases by approximately 13 percentage points (median).

With a view to the effects at the (average-)country level, when there are no temperature anomalies ($\hat{\mathcal{T}} = 0$), the inflation rate increases by six percentage points, while the GDP growth rate rises by around three percentage points (median) over a one-year horizon. However, in a climatic environment in which the temperature is higher ($\hat{\mathcal{T}} = 1$ °C), the inflation rate is now higher (7 percentage points) and the GDP growth rate is lower (less than three percentage points). Figure 1 shows that a temperature anomaly of +1 °C has been surpassed by 2015. As highlighted in the last two rows in Table 3, the results show that a 1 °C rise in temperature amplifies the inflationary impact of the global demand shock by one percentage point while dampening the increase in the growth rate of GDP by 0.6 percentage points. These effects become more pronounced and also statistically significant in an environment with a higher temperature anomaly (2 °C and beyond).

4.3.2. *Quantifying temperature anomalies for demand and supply shocks.* The previous results demonstrated the extent to which temperature anomalies influence an economy's adjustment path in response to shocks. The findings indicated that, in the case of a demand shock, temperature anomalies play a

FIGURE 7. Forecast error variance decomposition

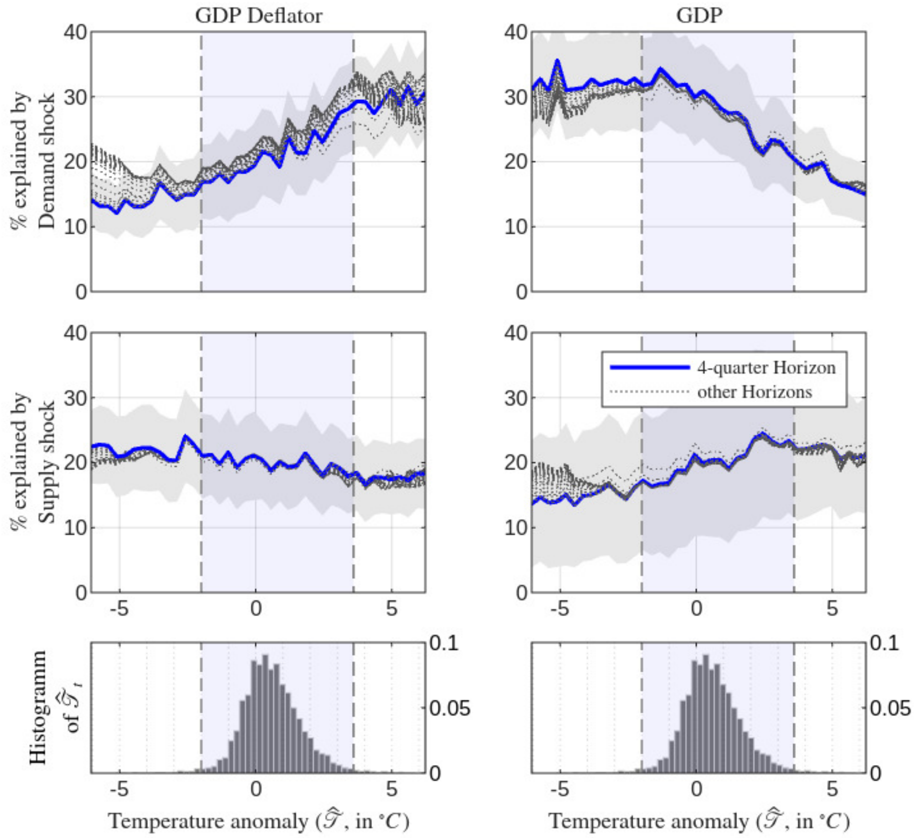

Note: The figure shows how the FEVDs (at a particular horizon) vary with the temperature anomalies. The blue lines are the (median) FEVDs at the four-quarter horizon (with the 68 percent credible interval given by the gray-shaded area), while the dotted lines represent FEVDs at horizons other than the four-quarter horizon. The third sub-panel displays the histogram of the temperature anomaly across time (quarterly frequency) and countries and the vertical lines indicate the 5th and 95th percentiles.

significant role, while not for supply shocks. We now extend this analysis to address two further aspects. First, we examine how much of the variation in output and inflation at the (average-)country level can be attributed to global demand and supply shocks. Second, we explore whether temperature anomalies affect the share of variation in output and inflation explained by these shocks. To address these aspects, we present evidence based on the FEVD.

The results are presented in Figure 7. We begin by discussing the findings related to the supply shock, which are shown in the sub-panels of the second row. The blue line depicts the proportion of inflation and output fluctuations attributable to the supply shock over a one-year horizon as a function of the temperature anomaly. As shown, when the temperature anomaly is low, the supply shock accounts for approximately 20 percent of inflation fluctuations, and this share hardly changes for higher values of temperature anomalies. For output fluctuations over the same horizon, the supply shock explains about 14 percent when the temperature anomaly is negative and around 20 percent when the temperature anomaly is high and positive. These proportions remain relatively stable even when shorter horizons are considered (represented by the gray dotted lines). Overall, the results suggest that temperature anomalies have only a minor impact on the proportion of price and output fluctuations explained by supply shocks. Moreover, these differences are not statistically significant and are small in magnitude.

This contrasts with the results for the demand shock, shown in the upper row of Figure 7. Starting with output, we see that for the one-year horizon (blue line), demand shocks account for a significant proportion of output variation when the temperature anomaly is negative (up to 30 percent). However, this proportion decreases noticeably when the temperature anomaly rises (down to 15 percent). This pattern is also evident for shorter horizons (gray-dotted lines). In other words, the role of demand shocks in explaining output variations diminishes as the temperature anomaly increases.

The opposite trend is observed for inflation. Within a four-quarter horizon (blue line), demand shocks explain a smaller proportion of inflation variation (around 13 percent) when the temperature anomaly is negative, but this proportion increases to above 30 percent as the temperature rises. This pattern suggests that the influence of demand shocks on inflation intensifies as the temperature increases. This pattern prevails also for horizons beyond one year.

Our findings indicate that while the proportion of inflation and output variations attributable to supply shocks remains unaffected by temperature anomalies, the opposite is evident in the context of demand shocks. The temperature anomaly affects the share of variation in output explained by demand shocks by more than ten percentage points. This contrasts with the results of Ciccarelli et al. (2024), who find that climate shocks account for less than one percent of real activity.

4.4. **Robustness and diagnostic checks.** We perform extensive robustness checks. These include an examination of non-linearities, omitted variables, the use of alternative measures of prices and output, an alternative identification approach, and additional interaction variables. We find a high degree of qualitative robustness of our results with respect to these extensions, details of which are provided in the Appendix.

Finally, we note that while the inclusion of additional control variables[15] reduces the risk that the estimated temperature effects are attributable to specific confounders, it cannot establish the formal exogeneity of the temperature anomaly. The analysis rules out a range of competing explanations, but the assumption that temperature anomalies are predetermined with respect to business cycle shocks remains a maintained identifying restriction, and the robustness checks do not constitute proof of strict exogeneity.

5. Conclusion

We argue that climate change—manifested in a gradual increase in global mean temperatures—should be conceptualized as structural change rather than as an exogenous shock. Accordingly, we advance the *general hypothesis* that climate change primarily alters the transmission of traditional macroeconomic shocks, rather than constituting an independent source of business cycle disturbances. Building on this premise, we propose the *particular hypothesis* that climate change steepens the Phillips curve (or aggregate supply curve), thereby amplifying the inflationary effects of demand shocks while dampening the corresponding output response.

We develop a theoretical model to formalize this idea and demonstrate that climate change, proxied by temperature anomalies, increases the slope of the Phillips curve. This results in an amplification of price effects from demand shocks and a muting of the output response. Our empirical analysis lends support to this conclusion. Notably, there is no evidence suggesting that temperature anomalies affect the price and output reactions to supply shocks. These results corroborate the hypothesis that temperature anomalies alter the slope of the Phillips curve, thereby influencing the transmission of demand shocks but not of supply shocks.

We show that even a small increase in mean surface temperatures implies measurable effects on the transmission of demand shocks to output and prices. For example, in a +1 °C global warming scenario, demand shocks induce a 20

[15]We consider patent activity, R&D intensity, demographic variables, carbon pricing, and a linear time trend in the Online Appendix.

percent larger impact on inflation, while the output response is approximately 15 percent smaller.

Consequences for fiscal and monetary policy. These findings have implications for economic policy, particularly with regard to fiscal and monetary measures designed to stabilize the business cycle. In an environment characterized by rising temperatures, expansionary fiscal policies used to combat aggregate demand downturns may prove less effective. As temperature anomalies increase, such policies could lead to significant price fluctuations with minimal impact on real activity. Moreover, our results suggest that monetary policy may be more effective in stabilizing prices and less effective in stabilizing output.

References

Abalo, Kodzovi, Brent Boehlert, Thanh Bui, Andrew Burns, Diego Castillo, Unnada Chewpreecha, Alexander Haider, Stephane Hallegatte, Charl Jooste, Florent McIsaac, Heather Jane Ruberl, Kim Smet, and Ken Strzepek (2025) 'The macroeconomic implications of climate change impacts and adaptation options: A modeling approach.' Policy Research Working Paper Series 11133, The World Bank

Abbritti, Mirko, and Sebastian Weber (2018) 'Reassessing the role of labor market institutions for the business cycle.' *International Journal of Central Banking* 14(1), 1–34

Acevedo, Sebastian, Mico Mrkaic, Natalija Novta, Evgenia Pugacheva, and Petia Topalova (2020) 'The effects of weather shocks on economic activity: What are the channels of impact?' *Journal of Macroeconomics* 65, 103207

Allen, Myles R., David J. Frame, Chris Huntingford, Chris D. Jones, Jason A. Lowe, Malte Meinshausen, and Nicolai Meinshausen (2009) 'Warming caused by cumulative carbon emissions towards the trillionth tonne.' *Nature* 458(7242), 1163–1166

Bai, Jushan, and Pierre Perron (1998) 'Estimating and testing linear models with multiple structural changes.' *Econometrica* 66(1), 47–78

——— (2003) 'Computation and analysis of multiple structural change models.' *Journal of Applied Econometrics* 18(1), 1–22

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti (2026) 'State dependence of monetary policy during global supply chain disruptions.' CESifo Working Paper Series 12451, CESifo

Barrage, Lint, and William Nordhaus (2024) 'Policies, projections, and the social cost of carbon: Results from the DICE-2023 model.' *Proceedings of the National Academy of Sciences* 121(13), e2312030121

Barrios, Salvador, Bazoumana Ouattara, and Eric Strobl (2008) 'The impact of climatic change on agricultural production: Is it different for Africa?' *Food Policy* 33(4), 287–298

Baumeister, Christiane, and James D. Hamilton (2015) 'Sign restrictions, structural vector autoregressions, and useful prior information.' *Econometrica* 83(5), 1963–1999

Bento, Antonio M., Noah Miller, Mehreen Mookerjee, and Edson Severnini (2023) 'A unifying approach to measuring climate change impacts and adaptation.' *Journal of Environmental Economics and Management* 121, 102843

Berg, Kimberly A., Chadwick C. Curtis, and Nelson C. Mark (2024) 'GDP and temperature: Evidence on cross-country response heterogeneity.' *European Economic Review* 169, 104833

Bettarelli, Luca, Davide Furceri, Loredana Pisano, and Pietro Pizzuto (2025) 'Greenflation: Empirical evidence using macro, regional and sectoral data.' *European Economic Review* p. 104983

Bilal, Adrien, and Diego R Känzig (2026) 'The macroeconomic impact of climate change: Global versus local temperature.' *The Quarterly Journal of Economics* 141(2), 889–944

Bilal, Adrien, and James H. Stock (2025) 'A guide to macroeconomics and climate change.' NBER Working Papers 33567, National Bureau of Economic Research, Inc, Mar

Blanchard, Olivier J., and David R. Johnson (2017) *Macroeconomics* (Pearson)

Cantelmo, Alessandro, Nikos Fatouros, Giovanni Melina, and Chris Papageorgiou (2024) 'Monetary policy under natural disaster shocks.' *International Economic Review* 65(3), 1441–1497

Casey, Gregory, Stephie Fried, and Ethan Goode (2023) 'Projecting the impact of rising temperatures: The role of macroeconomic dynamics.' *IMF Economic Review* 71(3), 688–718

Christoffel, Kai, Keith Kuester, and Tobias Linzert (2009) 'The role of labor markets for euro area monetary policy.' *European Economic Review* 53(8), 908–936

Ciccarelli, Matteo, and Fulvia Marotta (2024) 'Demand or Supply? An empirical exploration of the effects of climate change on the macroeconomy.' *Energy Economics* 129(C), S0140988323006618

Ciccarelli, Matteo, Friderike Kuik, and Catalina Martínez Hernández (2024) 'The asymmetric effects of temperature shocks on inflation in the largest euro area countries.' *European Economic Review* p. 104805

Cohen, Gail, Joao Tovar Jalles, Prakash Loungani, and Pietro Pizzuto (2022) 'Trends and cycles in CO2 emissions and incomes: Cross-country evidence on decoupling.' *Journal of Macroeconomics* 71, 103397

Crespo Cuaresma, Jesús, and Christian Glocker (2023) 'Production structure, tradability and fiscal spending multipliers.' *Journal of International Money and Finance* 138, 102921

De Sario, Manuela, Francesca Katherine de'Donato, Michela Bonafede, Alessandro Marinaccio, Miriam Levi, Filippo Ariani, Marco Morabito, and Paola Michelozzi (2023) 'Occupational heat stress, heat-related effects and the related social and economic loss: A scoping literature review.' *Frontiers in Public Health* 11, 1173553

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2012) 'Temperature shocks and economic growth: Evidence from the last half century.' *American Economic Journal: Macroeconomics* 4(3), 66–95

Doda, Baran (2014) 'Evidence on business cycles and CO2 emissions.' *Journal of Macroeconomics* 40, 214–227

Donadelli, Michael, Marcus Jüppner, and Sergio Vergalli (2022) 'Temperature variability and the macroeconomy: A world tour.' *Environmental and resource economics* 83(1), 221–259

Donadelli, Michael, Marcus Jüppner, Max Riedel, and Christian Schlag (2017) 'Temperature shocks and welfare costs.' *Journal of Economic Dynamics and Control* 82(C), 331–355

Drescher, Katharina, and Benedikt Janzen (2025) 'When weather wounds workers: The impact of temperature on workplace accidents.' *Journal of Public Economics* 241, 105258

Fernández-Villaverde, Jesús, Kenneth T Gillingham, and Simon Scheidegger (2024) 'Climate change through the lens of macroeconomic modeling.' *Annual Review of Economics* 17, 125–150

Fesselmeyer, Eric, Haoming Liu, Alberto Salvo, and Rhita P B Simorangkir (2024) 'Heat and observed economic activity in the rich urban tropics.' *The Economic Journal* 134(664), 3445–3460

Galí, Jordi (2015) *Monetary policy, inflation, and the business cycle: An introduction to the New Keynesian framework and its applications,* 2nd ed. (Princeton University Press)

Gasteiger, Emanuel, and Alex Grimaud (2023) 'Price setting frequency and the Phillips curve.' *European Economic Review* 158, 104535

Glocker, Christian, and Philipp Wegmüller (2024) 'Energy price surges and inflation: Fiscal policy to the rescue?' *Journal of International Money and Finance* 149, 103201

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014) 'Optimal taxes on fossil fuel in general equilibrium.' *Econometrica* 82(1), 41–88

Hamilton, James D. (1994) *Time Series Analysis* (Princeton University Press)

Hassler, John, and Per Krusell (2012) 'Economics and climate change: Integrated assessment in a multi-region world.' *Journal of the European Economic Association* 10(5), 974–1000

Hassler, John, Per Krusell, and Jonas Nycander (2016) 'Climate policy.' *Economic Policy* 31(87), 503–558

Held, Isaac M, and Brian J Soden (2006) 'Robust responses of the hydrological cycle to global warming.' *Journal of climate* 19(21), 5686–5699

Henseler, Martin, and Ingmar Schumacher (2019) 'The impact of weather on economic growth and its production factors.' *Climatic Change* 154(3), 417–433

Huber, Florian, Tamás Krisztin, and Michael Pfarrhofer (2023) 'A Bayesian panel vector autoregression to analyze the impact of climate shocks on high-income economies.' *The Annals of Applied Statistics* 17(2), 1543–1573

Iliyasu, Jamilu, and Aliyu Rafindadi Sanusi (2024) 'Climate change's impact on commodity prices: A new challenge for monetary policy.' *Portuguese Economic Journal* 23(2), 187–212

Jo, Ara (2025) 'Substitution between clean and dirty energy with biased technical change.' *International Economic Review* 66(2), 883–902

Kumar, Naveen, and Dibyendu Maiti (2024) 'Long-run macroeconomic impact of climate change on total factor productivity — evidence from emerging economies.' *Structural Change and Economic Dynamics* 68(C), 204–223

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

Letta, Marco, and Richard S. J. Tol (2019) 'Weather, climate and total factor productivity.' *Environmental & Resource Economics* 73(1), 283–305

Linsenmeier, Manuel (2024) 'Seasonal temperature variability and economic cycles.' *Journal of Macroeconomics* 79, 103568

Manne, Alan, Robert Mendelsohn, and Richard Richels (1995) 'MERGE: A model for evaluating regional and global effects of GHG reduction policies.' *Energy Policy* 23(1), 17–34

Marcott, Shaun A, Jeremy D Shakun, Peter U Clark, and Alan C Mix (2013) 'A reconstruction of regional and global temperature for the past 11,300 years.' *Science* 339(6124), 1198–1201

Mas-Colell, Andreu, Michael Whinston, and Jerry Green (1995) *Microeconomic Theory* (Oxford University Press)

Mohnen, Volker A., Walter Goldstein, and Wei-Chyung Wang (1990) 'The scientific challenge of measuring climate change.' *Energy Policy* 18(7), 641–651

Mukherjee, K., and Bazoumana Ouattara (2021) 'Climate and monetary policy: Do temperature shocks lead to inflationary pressures?' *Climatic Change* 167(3), 1–21

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (2024) 'How much will global warming cool global growth?' NBER Working Papers 32761, National Bureau of Economic Research, Inc, July

Nordhaus, William (1991) 'To slow or not to slow: The economics of the greenhouse effect.' *The Economic Journal* 101(407), 920–37

——— (2014) *A question of balance: Weighing the options on global warming policies* (Yale University Press, New Haven, CT)

Olea, José Luis Montiel, Mikkel Plagborg-Moller, Eric Qian, and Christian Wolf (2024) 'Double robustness of local projections and some unpleasant VARithmetic.' NBER Working Papers 32495, National Bureau of Economic Research, Inc

Onofrei, Mihaela, Anca Florentina Vatamanu, and Elena Cigu (2022) 'The relationship between economic growth and CO2 emissions in eu countries: A cointegration analysis.' *Frontiers in Environmental Science* 10, 934885

Ortiz-Bobea, Ariel, Toby R. Ault, Carlos M. Carrillo, Robert G. Chambers, and David B. Lobell (2021) 'Anthropogenic climate change has slowed global agricultural productivity growth.' *Nature Climate Change* 11(4), 306–312

Osman, Matthew B, Jessica E Tierney, Jiang Zhu, Robert Tardif, Gregory J Hakim, Jonathan King, and Christopher J Poulsen (2021) 'Globally resolved surface temperatures since the last glacial maximum.' *Nature* 599(7884), 239–244

Pesaran, M. Hashem, and Ron Smith (1995) 'Estimating long-run relationships from dynamic heterogeneous panels.' *Journal of Econometrics* 68(1), 79–113

Pretis, Felix (2021) 'Exogeneity in climate econometrics.' *Energy Economics* 96, 105122

Ricke, Katharine L., and Ken Caldeira (2014) 'Maximum warming occurs about one decade after a carbon dioxide emission.' *Environmental Research Letters* 9(12), 124002

Rubio-Ramírez, Juan F, Daniel F Waggoner, and Tao Zha (2010) 'Structural vector autoregressions: Theory of identification and algorithms for inference.' *The Review of Economic Studies* 77(2), 665–696

Ruff, Tyler W., and J. David Neelin (2012) 'Long tails in regional surface temperature probability distributions with implications for extremes under global warming.' *Geophysical Research Letters*

Sapir, Debarati G., and Claudine Misson (1992) 'The development of a database on disasters.' *Disasters* 16(1), 74–80

Tol, Richard S. J. (2024) 'A meta-analysis of the total economic impact of climate change.' *Energy Policy* 185, 113922

Towbin, Pascal, and Sebastian Weber (2013) 'Limits of floating exchange rates: The role of foreign currency debt and import structure.' *Journal of Development Economics* 101(C), 179–194

van der Ploeg, Frederick, and Armon Rezai (2026) 'Climate change, climate policy, and the macroeconomy.' CESifo Working Paper Series 12480, CESifo

van der Ploeg, Frederick, and Cees Withagen (2014) 'Growth, renewables, and the optimal carbon tax.' *International Economic Review* 55(1), 283–311

Zhang, Wei, Ning Ding, Yilong Han, Jie He, and Na Zhang (2023) 'The impact of temperature on labor productivity–evidence from temperature-sensitive enterprises.' *Frontiers in Environmental Science* 10, 1039668

Zhao, Xiaobing, Mason Gerety, and Nicolai V. Kuminoff (2018) 'Revisiting the temperature-economic growth relationship using global subnational data.' *Journal of Environmental Management* 223, 537–544

Zolghadr-Asli, Babak, Maedeh Enayati, Hamid Reza Pourghasemi, Mojtaba Naghdyzadegan Jahromi, and John P Tiefenbacher (2021) 'Application of Granger-causality to study the climate change impacts on depletion patterns of inland water bodies.' *Hydrological Sciences Journal* 66(12), 1767–1776

Appendix

CLIMATE CHANGE AND ECONOMIC DYNAMICS: TEMPERATURE-DEPENDENT SHOCK PROPAGATION

This Appendix complements the main text by providing: (i) additional details on the theoretical model, including the full set of equilibrium conditions and a simulation exercise that examines the signs of the matrix of impact effects of structural shocks; and (ii) an alternative, deliberately parsimonious theoretical framework that elucidates the core mechanism from a simpler perspective. It further details the empirical model specification, documents the data (sources, construction, and transformations), and presents a range of robustness checks. In addition, it explores alternative measures of the temperature anomaly and uses the rainfall anomaly as a different indicator of climate change within the empirical framework, reporting and discussing the corresponding results in depth. Additionally, we provide, among others, robustness checks using an alternative panel framework, an alternative shock identification framework, and additional control variables. The Appendix concludes with a Bayesian estimation of several key structural parameters of the theoretical model.

1. Equilibrium equation system, solution procedure and the temperature anomaly

The following lists the equations describing the equilibrium dynamics based on relative prices for which domestic final goods prices ($P_{D,t}$) is used as reference price.

Households.

- $1 = E_t \Lambda_{t,t+1} \frac{R_t}{1+\pi_{t+1}}$
- $1 - \psi_B \frac{s_t}{p_t} B_t^* = E_t \Lambda_{t,t+1} \frac{R^*}{1+\pi_{t+1}} \frac{S_{t+1}}{S_t}$
- $1 = E_t \Lambda_{t,t+1} \left(1 - \delta + \frac{p_{k,t+1}}{p_{t+1}}\right)$
- $\frac{w_t}{p_t} = \chi h_t^{\phi} C_t^{\sigma}$

Energy producers.

- $e_t = \xi_E \left((1-\zeta) e_{f,t}^{\frac{\kappa-1}{\kappa}} + \zeta k_{t-1}^{\frac{\kappa-1}{\kappa}}\right)^{\frac{\kappa}{\kappa-1}}$
- $\frac{k_{t-1}}{e_{f,t}} = \frac{1-\zeta}{\zeta} \left(\frac{p_{k,t}}{s_t P_{e,t}^*}\right)^{-\kappa}$
- $k_t = (1-\delta) k_{t-1} + I_t$

Final goods producers.

- $y_t = \frac{\xi_A(\hat{\mathcal{T}})}{1-\varsigma} h_t$
- $mc_t = \frac{1}{\xi_A(\hat{\mathcal{T}})} \left((1-\varsigma) w_t + \varsigma p_{e,t}\right)$

Prices and inflation.

- $p_t = (1-\vartheta) p_{c,t} + \vartheta p_{e,t}$
- $p_{c,t} = s_t^{1-\varrho}$
- $p_{e,t} = \left((1-\zeta)\left(s_t P_{e,t}^*\right)^{1-\kappa} + \zeta \left(p_{k,t}\right)^{1-\kappa}\right)^{\frac{1}{1-\kappa}}$
- $\pi_{D,t} = \beta E_t \pi_{D,t+1} + \frac{(1-\xi_p\beta)(1-\xi_p)}{\xi_p} \hat{mc}_t$

Market clearing and Balance of Payments.

- $e_t = \vartheta C_t + \frac{\varsigma}{\xi_A(\hat{\mathcal{T}})} y_t$
- $y_t = \varrho p_{c,t} \left[\left(1 - \vartheta + \frac{p_{e,t}}{p_{c,t}} \frac{\vartheta}{\varrho}\right) C_t + I_t\right] + s_t X_t^* - s_t P_{e,t}^* e_{f,t}$
- $s_t B_t^* - R^* s_t B_{t-1}^* = s_t X_t^* - s_t P_{e,t}^* e_{f,t}$

Carbon emissions, temperature anomaly and TFP.

- $M_t^{**} = (1-\varphi) M_{t-1}^{**} + \eta(e_{f,t} + e_{f,t}^*)$ where $e_{f,t}^* = e_f^*(Y_t^*, P_{e,t}^*)$
- $\mathcal{T}(M_t^{**}) = \bar{\mathcal{T}}$
- $\xi_A(\hat{\mathcal{T}}) = e^{-\alpha \hat{\mathcal{T}} + \bar{\xi}_A}$

FIGURE 1. Temperature sensitivity of the $\boldsymbol{H}$ matrix

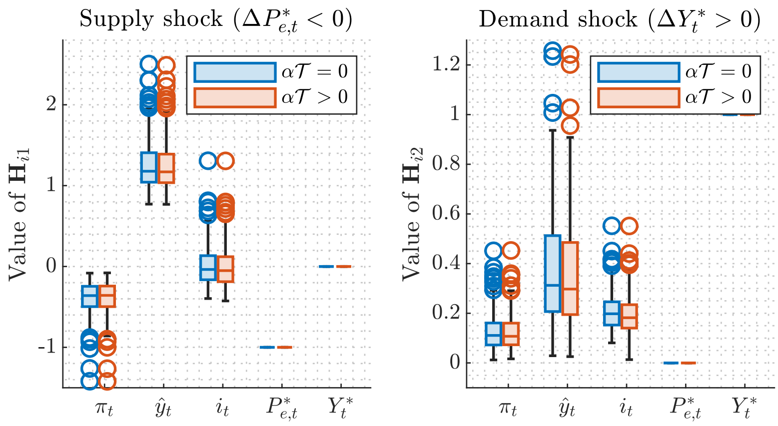

Note: The figure shows the sensitivity of the coefficients of the $\boldsymbol{H}_{n+2\times 2}$ matrix where n refers to the number of endogenous country variables and the two columns to the supply shock (first sub-panel) and demand shock (second sub-panel).

Monetary policy.

- $i_t = \bar{\imath} + \phi_\pi \pi_t + \phi_y \hat{y}_t$ where $i_t = \ln(R_t)$

where in equilibrium we have that $P^H_{e,t} = P^F_{e,t} = P_{e,t}$. The relative prices are given by $s_t = S_t/P_{D,t}$, $p_{k,t} = P_{k,t}/P_{D,t}$, $p_t = P_t/P_{D,t}$, $w_t = W_t/P_{D,t}$, $p_{c,t} = P_{c,t}/P_{D,t}$, and $p_{e,t} = P_{e,t}/P_{D,t}$. Moreover, we have that $P_{F,t} = S_t$, $C_t = (1-\vartheta)^{-1}c_t = \vartheta^{-1}e_{H,t}$, $\hat{Y}^*_t \sim \mathrm{AR}(1)$ and $\hat{P}^*_{e,t} \sim \mathrm{AR}(1)$. Finally, $mc_t = MC_t(\hat{\mathcal{T}})/P_{D,t}$ is real marginal costs, the relation between the CPI and domestic goods prices is given in relative terms by $p_t = \frac{1+\pi_t}{1+\pi_{D,t}} p_{t-1}$, the change in the real exchange rate is given by $s_t/s_{t-1} = \frac{S_t/S_{t-1}}{1+\pi_{D,t}}$, and the export demand function is $X_t = (1/s_t)^{-\theta^*} Y^*$, where P^* is normalized to unity. A hat on a variable denotes its relative deviation from its steady state value.

1.1. **Solution procedure and simulation.** We simulate the model using established techniques, following a sequence that accounts for the dependence of the steady state on the prevailing temperature anomaly. Specifically, for each fixed level of the temperature anomaly $\mathcal{T}$, we first determine the associated steady state of the model—recognizing that the steady state itself varies as a function of $\mathcal{T}$. Subsequently, we obtain the dynamic solution by performing a standard log-linearization of the model around this new steady state (and solving it based on the assumption of rational expectations), which enables the simulation of the economy's response to demand and supply shocks. This procedure is systematically iterated over a grid of temperature anomaly values: for each value, we recompute the steady state and resolve the dynamics for that specific climate environment.

Importantly, this approach interprets each $\mathcal{T}$ as defining a distinct, persistent climate regime, with the corresponding steady state serving as the reference point for the (log-)linearization. In this framework, changing the temperature anomaly corresponds to altering a structural parameter of the model. We then analyze how (transitory) exogenous shocks are propagated and amplified in each new climatic setting.

1.2. **The temperature anomaly and the contemporaneous impact of the global shocks.** Our empirical identification strategy crucially relies on the contemporaneous impact of the global demand and supply shocks on the country-level inflation and output measures. The key question at this stage concerns the extent to which the temperature anomaly shapes this impact qualitatively. To this purpose, we examine the sensitivity of the contemporaneous impact responses of the global demand and supply shock on inflation and output at the country-level. This only concerns the matrix $\boldsymbol{H}$ (as defined in equations 1-3 in Section 2 of the main part). Figure 1 illustrates the sensitivity of the matrix $\boldsymbol{H}$ with respect to the temperature anomaly ($\mathcal{T}$). The matrix $\boldsymbol{H}$ characterizes the contemporaneous responses of the endogenous variables in the theoretical model to two fundamental structural shocks: a global demand shock and a global supply shock. As previously described, all structural parameters are assigned independent uniform distributions, with upper and lower bounds specified in the Table presenting the calibration of the model.

To assess the impact of temperature anomalies, we conduct 2,000 simulations of the theoretical model under two scenarios: (i) $\mathcal{T} > 0$, and (ii) $\mathcal{T} = 0$. This approach enables us to isolate the effect of changes in $\mathcal{T}$ on the distribution of responses, while preserving flexibility in model calibration. In the figure, the blue and orange box plots display the distributions of the corresponding entries of $\boldsymbol{H}$, focusing on selected reduced-form coefficients for inflation, output, the nominal interest rate, and the two foreign variables (global output and global prices).

Two key insights emerge from this analysis. First, we observe that $\text{sgn}(\boldsymbol{H}(\mathcal{T} > 0)) = \text{sgn}(\boldsymbol{H}(\mathcal{T} = 0))$; that is, the sign of each element in $\boldsymbol{H}$ remains invariant to the temperature anomaly. Second, this procedure allows for the identification of robust sign restrictions on the structural response coefficients, which are subsequently employed for the empirical identification strategy (which is outlined in detail in Section 4 in the main part).

2. Temperature anomaly, the AS curve and demand shocks – a simple(r) theoretical methodology

The following section provides an alternative theoretical framework for addressing our research question. The aim here is to express and analyze the research question within the simplest possible theoretical model environment. This is intentional in order to underscore the simplicity of the overarching research idea more explicitly. To this end, we employ the standard aggregate demand-–aggregate supply (AS–AD) model, as commonly presented in undergraduate macroeconomic textbooks (see Sachs and Larraín B., 1993; Blanchard and Johnson, 2017, among others).

Consistent with the empirical findings of Kumar and Maiti (2024), we again assume that total factor productivity ξ_A is negatively affected by climate change, and $\mathcal{T}(M_t^{**}) = \bar{\mathcal{T}}$. Specifically, we model ξ_A as a function of $\mathcal{T}$, that is, $\xi_A = \xi_A(\mathcal{T})$, with $\xi_A' < 0$. For the remaining components of the model, we adopt the framework presented in Chapter 7 of Blanchard and Johnson (2017). This introduces the following four key equations for the purpose of establishing the AS curve: (1) the aggregate production function $Y = \xi_A(\mathcal{T})N$, (2) the price-setting equation, which extends the marginal cost relation by a markup (μ), $P = (1+\mu)\frac{W}{\xi_A(\mathcal{T})}$, (3) the wage-setting equation, $W = P^e f(u, z)$, with $f_u < 0$ and $f_{uu} > 0$, and (4) the definition of the unemployment rate, $u = 1 - \frac{N}{L}$. In these equations, Y represents aggregate output, N denotes employment, L is the labor force, P and P^e are the actual and expected price levels, respectively, W is the nominal wage, and z is a catchall variable encompassing factors such as employment protection, unemployment benefits, and similar institutional elements.

These equations can be put together to obtain the following expression for the aggregate supply (AS) curve, which identifies the relation between output

FIGURE 2. A view from the AS-AD model

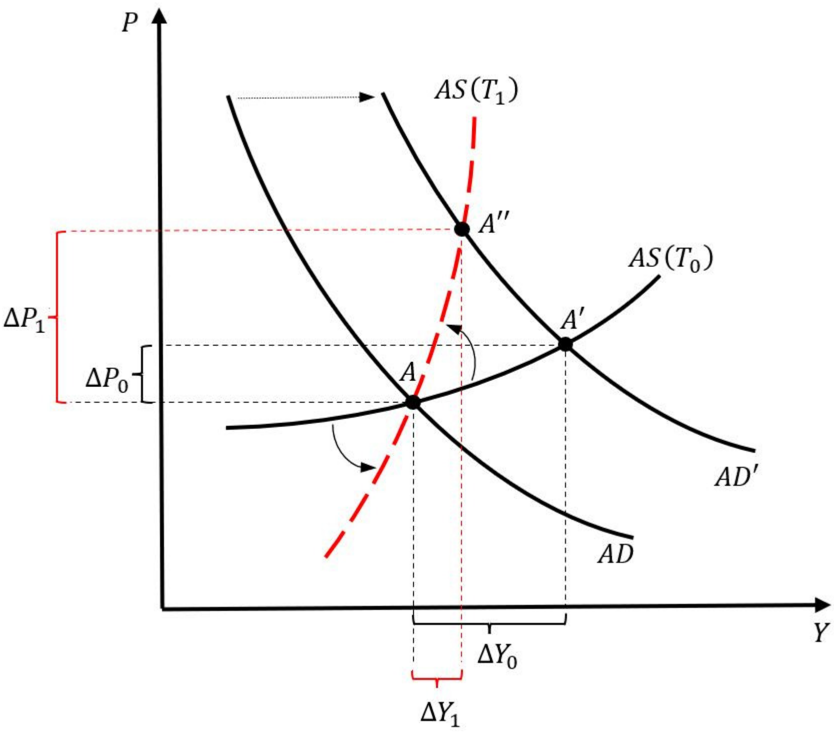

and prices from a supply side view

$$P = (1+\mu)P^e f\left(1 - \frac{Y}{\xi_A(\mathcal{T})L}, z\right) \tag{1}$$

From this we can identify the following for the slope of the AS curve

$$\left.\frac{dP}{dY}\right|_{\text{AS}} = -\frac{(1+\mu)P^e}{\xi_A(\mathcal{T})^2 L} \cdot f_u > 0 \tag{2}$$

Furthermore, we have that

$$\frac{d}{d\mathcal{T}}\left(\left.\frac{dP}{dY}\right|_{\text{AS}}\right) = \frac{d^2P}{dY\,d\mathcal{T}} = (1+\mu)P^e\xi_A'\frac{2f_u\xi_A(\mathcal{T})L - f_{uu}Y}{\xi_A(\mathcal{T})^4L^2} > 0 \tag{3}$$

Consequently, the slope of the AS curve increases with the temperature anomaly $\mathcal{T}$. This result mirrors the one previously presented in the theoretical (DSGE) model.

Figure 2 illustrates the immediate effects of an increasing temperature anomaly $\mathcal{T}$ on the implications of demand shocks. Consider the initial equilibrium at point A, determined by the intersection of the original AD curve and the AS($\mathcal{T}_0$) curve, which represents the AS curve under a low temperature anomaly environment ($\mathcal{T}_0$). Following an expansionary aggregate demand shock, the economy moves to a new equilibrium at point A', where output increases by ΔY_0 and prices rise by ΔP_0.

When the temperature anomaly rises from $\mathcal{T}_0$ to $\mathcal{T}_1$, the AS curve pivots to the left, now depicted by the red-dashed line, AS($\mathcal{T}_1$). The equilibrium with the original AD curve occurs now at point A, and following the same expansionary demand shock, the economy reaches a final equilibrium at point A''. In this scenario, output and prices still increase, but by smaller (output) and larger (prices) amounts. Crucially, we observe that $\Delta Y_1 < \Delta Y_0$ and $\Delta P_1 > \Delta P_0$. In other words, under increasing temperature anomalies, the expansionary demand shock leads to a more pronounced rise in prices, while dampening the increase in output.

3. Bayesian estimation and prior densities

Our benchmark (interacted Bayesian) VAR model (at the country-level) can be re-arranged to the following expression:

$$Y = Xb + E, \tag{4}$$

where X now includes all regressors of equation (31) of the main part (that is, lagged endogenous and exogenous variables, and also the interaction variables[1]), and E has a variance-covariance matrix Σ. We use a rather diffuse version of the conjugate prior densities. To this purpose, we utilize the Normal-Wishart prior density for b and Σ^{-1}: $p(b, \Sigma^{-1}) = p(b)p(\Sigma^{-1})$, where $b \sim N\left(\underline{b}, \underline{V}\right)$ and $\Sigma^{-1} \sim W\left(\underline{H}, \underline{v}\right)$. The rather agnostic setup of the prior densities is then obtained by using $\underline{v} = 0$ and $\underline{H}^{-1} = I$, where I is the identity matrix of conformable size. For the slope parameters we set $\underline{b} = 0$ and $\underline{V} = 100I$. Given this particular specification for the prior densities, we obtain the following conditional posterior densities for $p(b|Y, \Sigma^{-1})$ and $p(\Sigma^{-1}|Y, b)$:

$$b|Y, \Sigma^{-1} \sim N\left(\bar{b}, \bar{V}\right) \qquad \Sigma^{-1}|Y, b \sim W\left(\bar{H}, \bar{v}\right), \tag{5}$$

with $\bar{V} = \left(\underline{V}^{-1} + \sum_{t=1}^{T} X'\Sigma^{-1}X\right)^{-1}$, $\bar{b} = \bar{V}\left(\underline{V}^{-1}\underline{b} + \sum_{t=1}^{T} X'\Sigma^{-1}Y\right)$, $\bar{v} = T - \underline{v}$ and $\bar{H} = \left(\underline{H}^{-1} + \sum_{t=1}^{T} X'\Sigma^{-1}(Y - Xb)(Y - Xb)'\right)^{-1}$.

We employ a Gibbs sampler to draw from the multivariate Normal $p(b|Y, \Sigma^{-1})$ and the Wishart $p(\Sigma^{-1}|Y, b)$ distribution.

[1] Since the interaction variable does not enter the endogenous variables contemporaneously, but influences them only through lagged values of the endogenous state vector, the country-level VAR can be equivalently expressed as a VARX (vector autoregressive model with exogenous regressors, X), with the interaction terms included in the vector of exogenous variables.

We sample 6,000 draws from the posterior distribution. After discarding the first 1,000, we are left with 5,000 draws for each parameter. As is common in the literature for Bayesian estimation of VARs, we use rejection sampling to impose stability on the BVAR coefficients and only keep stable draws. Our results are qualitatively not affected by this choice of the sampling.

4. The data: general overview, sources and transformations

The empirical model is estimated using quarterly time series data covering seven key variables: a measure of global output (GDP), global prices, and country-specific data on output (GDP), prices (GDP deflator), short-term nominal interest rates, an indicator for the severeness of country-specific natural disasters, and temperature anomalies.

In the empirical specification, we include a set of truly exogenous regressors, namely a constant and a country-specific dummy variable designed to capture the turbulence associated with the COVID-19 pandemic. This dummy is constructed separately for each country in the sample and takes three values: -1 in quarters in which a lockdown policy was imposed, 1 in quarters in which lockdown measures were lifted (i.e., relaxed), and 0 otherwise. Although the initial introduction of lockdowns coincides across many countries in early 2020, subsequent policy actions are markedly heterogeneous: the timing, intensity, and duration of both the imposition and relaxation of lockdown measures vary substantially across countries. Empirically, the slope coefficient on this COVID-19 dummy is highly statistically different from zero in every country. Most importantly, including this control yields reduced-form residuals that are consistent with normality in all country-specific VAR models according to the Jarque–Bera test; by contrast, omitting the COVID-19 dummy generally leads to a rejection of residual normality.

The primary source for macroeconomic variables is the *OECD Main Economic Indicators*, which offers extensive and harmonized coverage across countries. The choice of sample and countries is largely determined by the availability of sufficiently long quarterly data series. In cases where country-level GDP and GDP deflator data are incomplete prior to 1961, we reconstruct or *back-cast* missing series following established practices. Specifically, we utilize related series, such as employment and unemployment rates, industrial production, and retail sales (for GDP); and the consumer price index (CPI) and producer price index (PPI) (for the GDP deflator). Backcasting procedures have been implemented for Austria and New Zealand, to achieve consistent long-term national accounts data.

4.1. **Temperature anomalies: FAOSTAT.** Temperature anomaly data are obtained from the FAOSTAT Temperature Change domain, provided by the United Nations Food and Agriculture Organization (FAO) following this link (retrieved in June, 2024). FAOSTAT disseminates detailed statistics of land surface air temperature change by country, derived from the NASA Goddard Institute for Space Studies (NASA–GISS) Global Surface Temperature Change data (GISTEMP). The FAOSTAT dataset offers monthly, seasonal, and annual temperature anomalies for 198 countries and 39 territories, spanning the period 1961–2024. Temperature anomalies are computed as deviations (°C) from the 1951–1980 baseline climatological mean:

- **Reference period:** The baseline period for calculating anomalies is the average land temperature between 1951 and 1980.
- **Anomaly calculation:** For each country, month, and season, anomalies represent deviations from the corresponding period's mean during the reference years. The dataset reports temperature anomalies—the deviation of surface air temperature from the average in the baseline period, expressed in degrees Celsius (°C). In particular, country temperature anomalies from FAOSTAT are computed by subtracting the 1951-80 average at station level (more on this below) from the current temperature. Data are available at monthly (seasonally adjusted or unadjusted), and annual frequency. We aggregate the monthly data to a quarterly frequency by computing the arithmetic mean of the observations across the three corresponding months of each quarter.
- **Source:** The original source—the NASA Goddard Institute for Space Studies (NASA–GISS) GISTEMP dataset—contains temperature records from as early as 1880, but FAOSTAT aggregates and publishes country-level statistics from 1961. The data are based upon thousands of meteorological stations globally (see Food and of the United Nations, 2023; Emediegwu and Ubabukoh, 2023). Data are calculated from land-based

meteorological stations and aggregated using the FAO Global Administrative Unit Layers (GAUL) to provide country-level statistics. Aggregates are area-weighted for regional reporting using FAO Land Use dataset (see Food and of the United Nations, 2023).

- **Meteorological Station Coverage:** The dataset is based on the Global Historical Climatology Network (GHCN v4), which consists of continuously updated data from over 26,000 meteorological stations worldwide (see Food and Agriculture Organization of the United Nations, 2025).
- **Countries/Territories Covered:** As of 2025, statistics are available for 198 countries and 39 territories, making it one of the most comprehensive global land temperature anomaly datasets.

4.2. **Temperature anomalies: Berkeley Earth database.** As an alternative, we consider the Berkeley Earth database, which offers regional and country-level land-surface temperature data using the Berkeley Earth averaging method. Berkeley Earth publishes country-level estimates of warming rates since 1960, reported as °C per century, for 185+ countries and territories. The key statistic is "Warming since 1960" (°C/century), representing the linear trend in annual mean temperature anomaly over the period from 1960 to the most recent year available. Each value is given with an uncertainty margin, representing the statistical confidence (typically 1 standard deviation) of the estimated anomalies. Berkeley Earth uses the full Global Historical Climatology Network (GHCN v4), World Meteorological Organization data, national agency records, and crowd-sourced archives. Its database contains over 50,000 land-based meteorological stations globally, which are used in a spatial averaging algorithm. This approach is peer-reviewed.

The key difference between the FAOSTAT data and those from the Berkeley Earth database concern the measuring points (Meteorological Stations): The FAOSTAT relies on over 26,000 land-based meteorological stations from GHCN v4 (via NASA–GISS); the Berkeley Earth utilizes over 50,000 stations, combining GHCN v4, WMO, national, and private sources.

4.3. **Comparing the two sources for the temperature anomalies.** For both sources, the focus on temperature anomalies—rather than absolute temperature levels—mitigates station-specific biases and long-term local changes unrelated to climate trends. This is especially relevant for macroeconomic analyses, as temperature anomalies represent sustained deviations from typical local climate conditions, which have pronounced economic implications (see Dell et al., 2014, for instance).

Although both the FAOSTAT and Berkeley Earth databases provide temperature anomaly data, conceptual differences in their methodologies lead to variations in the reported statistics. Empirical comparisons reveal a high degree of correlation between the temperature anomalies produced by these two sources, with observed country-level correlation coefficients ranging from approximately 0.71 to 0.97. This strong correlation underscores the substantial similarity between the datasets in capturing broad-scale temperature variations and trends. Nevertheless, the observed deviations in correlation reflect the inherent differences in anomaly computation, temporal resolution, spatial coverage, and data processing protocols employed by each institution. Thus, while FAOSTAT and Berkeley Earth present complementary perspectives on home-country temperature anomalies, the distinction in their conceptual frameworks and analytical approaches naturally results in both close agreement and meaningful divergence within their respective anomaly statistics.

4.4. **Rainfall anomaly as alternative climate change measure.** Our baseline results rely on the FAOSTAT-temperature anomaly as the key country-level proxy for climate change. While the temperature anomaly is widely regarded as an appropriate measure in this context (Hassler et al., 2016; Hassler and Krusell, 2012), a range of alternative climate indicators has been used in the literature, often tailored to the empirical model and the specific channel under investigation (see also the discussion in the Introduction and, inter alia, Huber et al. (2023); Deschênes and Greenstone (2011)). To further substantiate the role of climate conditions in shaping the transmission of demand shocks—and, more specifically, to assess the robustness of our findings to the choice of climate proxy—we replace the temperature anomaly with a country-level rainfall anomaly, as motivated by Giorgi et al. (2019); Palagi et al. (2022) among others as climate change measure.

We obtain rainfall data from the Copernicus Climate Data Store (retrieved January, 9th 2026 following this link). The underlying rainfall observations are recorded in millimeters (mm)—equivalently, liters per square meter (L/m^2),

Figure 3. Country rainfall anomalies

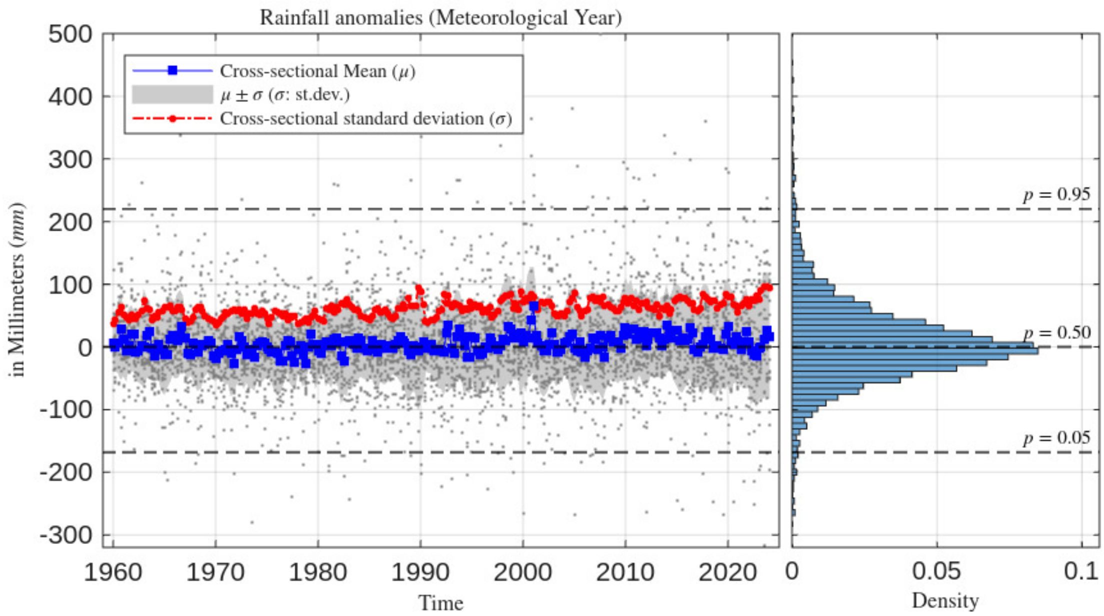

Note: The figure shows the (quarterly) rainfall anomaly (i.e. deviation) by country with respect to a baseline period, corresponding to the period 1960–1980. The series cover the period 1960–2023 for 20 countries. The 5th, 50th and 95th percentiles are derived from the joint distribution, i.e. over all countries and the whole period.

where 1 mm corresponds to 1 L/m^2—and are available at a monthly frequency for each country, constructed by aggregating measurements from multiple weather stations within each country. Because station coverage is incomplete in some months, we first compute, for each country and month, the cross-station mean of available observations to obtain a single national monthly rainfall measure. We then convert the monthly data to quarterly frequency by summing the three monthly values within each quarter. Repeating this procedure for all countries yields a quarterly rainfall time series for each country over the period 1960–2024.

To construct rainfall anomalies, we follow the same approach used by the United Nations for the temperature anomaly series. Specifically, we define the base period as 1960–1980 and compute, for each country and each quarter, the baseline rainfall level as the average rainfall observed in that quarter over the base period. The rainfall anomaly is then defined as the deviation of realized rainfall in a given country–quarter from the corresponding quarter-specific baseline level. By construction, the resulting anomaly series is directly comparable to the temperature anomaly in that it removes deterministic seasonal variation and can therefore be interpreted as a seasonally adjusted measure of unusual rainfall conditions.

Figure 3 plots the rainfall anomalies by country (left panel) and their pooled distribution across all countries and quarters/years (right panel). In contrast to temperature anomalies—which exhibit a clear structural shift around 1977 in our sample—rainfall anomalies display no comparable break. Instead, the series remain centered close to zero throughout the sample, both on average (blue dotted line) and, by and large, at the country level. If anything, the cross-sectional dispersion increases slightly toward the end of the sample; however, this change is modest and not quantitatively and statistically meaningful.

We subsequently substitute the temperature anomaly variable in our empirical model with the rainfall anomaly variable and carry out the same exercises as in the main part. The results are provided and discussed in detail in Section 5.2 below.

4.5. **The natural disaster variable: Concept, source, and implications.** Our baseline specification is an interacted VAR that includes, among other covariates, a measure of natural disasters that enters the system as a strictly exogenous regressor. The disaster variable captures realized damages associated with hydrological, meteorological, and climatological events. It is constructed from the EM-DAT database (Sapir and Misson, 1992); for documentation of the underlying classification system, see EM-DAT's classification guide.

Treating realized disasters as strictly exogenous is well justified in our setting. Given the quarterly frequency of the data, disaster realizations within quarter t are not plausibly driven by contemporaneous macroeconomic innovations or policy surprises occurring within the same quarter. In other words, while macroeconomic conditions may shape exposure and long-run vulnerability, creating a feedback between output growth and climate change over the medium term (Bilal and Känzig, 2026), the within-quarter realization of hydrological, meteorological, and climatological events is predetermined with respect to the system's endogenous variables at horizon t. This identification choice aligns with the broader approach in the literature that models weather- and climate-related shocks as external drivers of the macro-financial system rather than as endogenous outcomes; see, for example, Ciccarelli et al. (2024).

Incorporating the disaster measure remains important because it captures realized, non-temperature channels of climate-related disruption (e.g., physical destruction and acute supply interruptions) that are not fully spanned by temperature anomalies, thereby reducing omitted-variable bias when estimating how the macroeconomic transmission in the (interacted) VAR model varies with the temperature state.

In the remainder of this subsection, we first describe in detail what the disaster variable measures and how it is operationalized in the empirical analysis. We then discuss the economic interpretation of this measure and clarify its role for identification and inference in the interacted VAR.

4.5.1. *What the disaster variable measures.* EM-DAT was established in 1988 and is grounded in an anthropocentric conception of disasters and emergencies, emphasizing events that generate substantial and unexpected harm to human populations. Consistent with its documentation, EM-DAT defines a disaster as "a situation or event which overwhelms local capacity, necessitating a request to the national or international level for external assistance; an unforeseen and often sudden event that causes great damage, destruction, and human suffering" (see EM-DAT general definitions and Sapir and Misson (1992)). The database's core classification principle follows a hazard-based logic, that is, disasters are categorized by the event triggering the adverse outcome.

Under the current methodology (see the detailed description in this document), EM-DAT organizes natural hazards into six main groups: geophysical, hydrological, meteorological, climatological, biological, and extra-terrestrial (with additional subtypes defined within each group). In our empirical application, we restrict attention to disasters in the hydrological, meteorological, and climatological groups.

The EM-DAT data model records disasters and maps their impacts at the country level. Since a single event may affect multiple countries, EM-DAT stores country-specific observations for each affected country; consequently, a given disaster can generate multiple country-level impact records, potentially originating from more than one reporting source. For each of the three hazard groups used in our analysis, we construct two quantitative disaster intensity measures: (i) the affected population relative to the total population of the country, and (ii) reported damage costs relative to the country's nominal GDP.

Table 1. Quantitative disaster measures: sign and statistical significance

	Damage to GDP			Affected population to total population		
	π_t	$\hat{y}_t$	i_t	π_t	$\hat{y}_t$	i_t
Australia	$+^{**}$	$-$	$-$	$+^{*}$	$-^{*}$	$+$
Austria	$-$	$+$	$-$	$-$	$-^{*}$	$-$
Belgium	$-$	$-$	$-^{*}$	$+$	$-^{**}$	$-$
Canada	$-$	$-^{**}$	$-$	$+$	$-^{**}$	$+$
Finland	$+^{***}$	$-$	$-^{***}$	$-$	$-$	$-$
France	$-$	$+^{*}$	$-$	$-$	$-^{*}$	$+$
Greece	$+$	$+$	$-$	$+$	$+$	$-$
India	$+^{**}$	$-^{*}$	$-$	$-$	$-^{*}$	$-^{**}$
Italy	$-$	$-^{*}$	$-$	$+$	$-^{*}$	$-^{*}$
Japan	$+^{*}$	$-^{*}$	$-$	$+$	$-$	$-$
Korea	$+$	$-$	$-$	$+$	$+$	$-$
Mexico	$-$	$-^{*}$	$-$	$-$	$-$	$-$
Netherlands	$+$	$+$	$-$	$+^{*}$	$-$	$-$
New Zealand	$+^{**}$	$-^{**}$	$+$	$+^{**}$	$-$	$+$
Norway	$+$	$-$	$-^{*}$	$-$	$-$	$-^{*}$
Portugal	$-$	$-^{*}$	$-$	$+$	$-^{*}$	$-^{*}$
Spain	$-$	$+$	$+$	$-$	$-$	$-$
Sweden	$+^{**}$	$-$	$+$	$+^{*}$	$-^{*}$	$-^{*}$
United Kingdom	$-$	$+$	$+$	$-$	$+$	$+$
United States	$+^{*}$	$-$	$+$	$+^{*}$	$-^{*}$	$+$

Note: The table lists the results of the estimates for the two disaster measures. The "Damage to GDP" measure is used in the baseline model; the "affected population to total population" measure (i.e. Affected population ratio) is an alternative measure and used as a robustness check. The entries indicate the sign of the estimated coefficients and the stars indicate statistical significance at the one percent (***), five percent (**) and 16 percent level (*).

We then obtain a composite disaster indicator at the country–time level by aggregating across the three hazard groups, simply summing the respective ratios for each period.

4.5.2. *Implications and role in the empirical model.* As described in the main part of the paper, our baseline specification includes a disaster measure as a control variable indicating reported damages from natural disasters (in national currency units) relative to nominal GDP. Table 1 reports the sign (+/−) of corresponding parameter estimates for each country and for each of the three endogenous variables in the VAR system. To illustrate, for Australia the disaster variable is associated with higher inflation and lower output (which is in line with the findings provided in Noy (2009); Botzen et al. (2019)) and lower nominal interest rates; however, only the inflation coefficient is statistically significant (at the 5% level). More generally, across countries, whenever the disaster coefficient is statistically different from zero (at least at the 16% level), its sign is consistent with the expected qualitative pattern: upward pressure on inflation and downward pressure on real activity and interest rates (which is in line with the findings provided in Noy (2009); Botzen et al. (2019) for the effect on output and with Parker (2018); Cevik and Jalles (2024) for the effect on the inflation rate). The main exception is France, where the output coefficient is positive and statistically significant, a result that may reflect short-run reconstruction dynamics (e.g., investment aimed at restoring destroyed capital stock).

Turning to the mean-group panel estimates of the VAR model, we find that only in case of inflation the disaster variable exhibits a statistically significant effect (at the 5% level), whereas the effects on output and the nominal interest rate are not statistically distinguishable from zero. Accordingly, our baseline conclusions are not driven by the inclusion of the disaster control: impulse responses to a global demand shock are visually indistinguishable whether the disaster variable is included in the VAR model or omitted. Nevertheless, given the prevalence of significant inflation effects at the country level and for the mean-group estimate as well, we retain this control in the baseline model. In contrast to Bilal and Känzig (2026), we do not find relevant negative effects on output resulting from more frequent and more intensive natural disasters.

As an additional robustness check, we replace the damage-based disaster measure with an alternative disaster indicator capturing the affected population relative to total population. The country-level estimates for this measure are also reported in Table 1 (right-hand side). The qualitative pattern is broadly similar: disasters are typically associated with higher inflation and lower output and nominal interest rates, particularly in cases where the coefficients are statistically significant. However, a comparison across the two measures indicates that statistically significant coefficients occur more frequently when using the damage-to-GDP ratio than when using the affected-population ratio. Consistent with this, at the mean-group level the affected-population measure is not statistically significant in any equation of the system, suggesting limited explanatory power in our setting. Consequently, the demand-shock IRFs from a model including the affected-population control are again visually indistinguishable from those obtained from a specification without any disaster control. For brevity, we do not report the corresponding IRFs, but Table 1 summarizes the relevant inference and thus provides a direct indication of the role of alternative disaster measures in the VAR system.

5. Robustness and diagnostic checks

The following Section considers various robustness assessments. We examine the role of an alternative measure of the temperature anomaly, of higher order non-linearities, of omitted variables, of alternative measures for prices and output, of an alternative identification approach, of a different country selection, and of additional interaction variables.

5.1. **Using an alternative measure for the temperature anomalies.** The baseline results rely on country-level temperature anomaly data from FAOSTAT. In this section, we re-examine these results using an alternative temperature anomaly measure derived from the Berkeley Earth database. The Berkeley Earth data are originally available at a monthly frequency, which we convert to quarterly series by computing three-month averages. We then incorporate these alternative temperature anomaly data into the IPVAR model, re-estimate the model, and compute the corresponding impulse response functions (IRFs) to global demand and supply shocks. The results are presented in Figures 4 and 5.

Overall, the results are qualitatively consistent with the baseline findings, although some quantitative differences emerge. In particular, an expansionary global demand shock is found to trigger a statistically significant and persistent increase in country-level inflation (prices) and output, with the magnitude and persistence of the response strongly influenced by the prevailing temperature anomaly. Specifically, high temperature anomalies are associated with a pronounced and prolonged rise in prices but a muted output response, whereas low temperature anomalies produce the opposite pattern. These differences in the IRFs for prices and output are statistically significant, as shown in the third-column panels of Figure 4. Thus, the baseline conclusions are corroborated when employing the alternative Berkeley Earth measure of temperature anomalies.

For the expansionary global supply shock, we again find a statistically significant divergence between prices and output at the country level. However, in this case the IRFs do not appear to be significantly affected by temperature

FIGURE 4. Demand shock using Berkeley Earth temperature anomaly

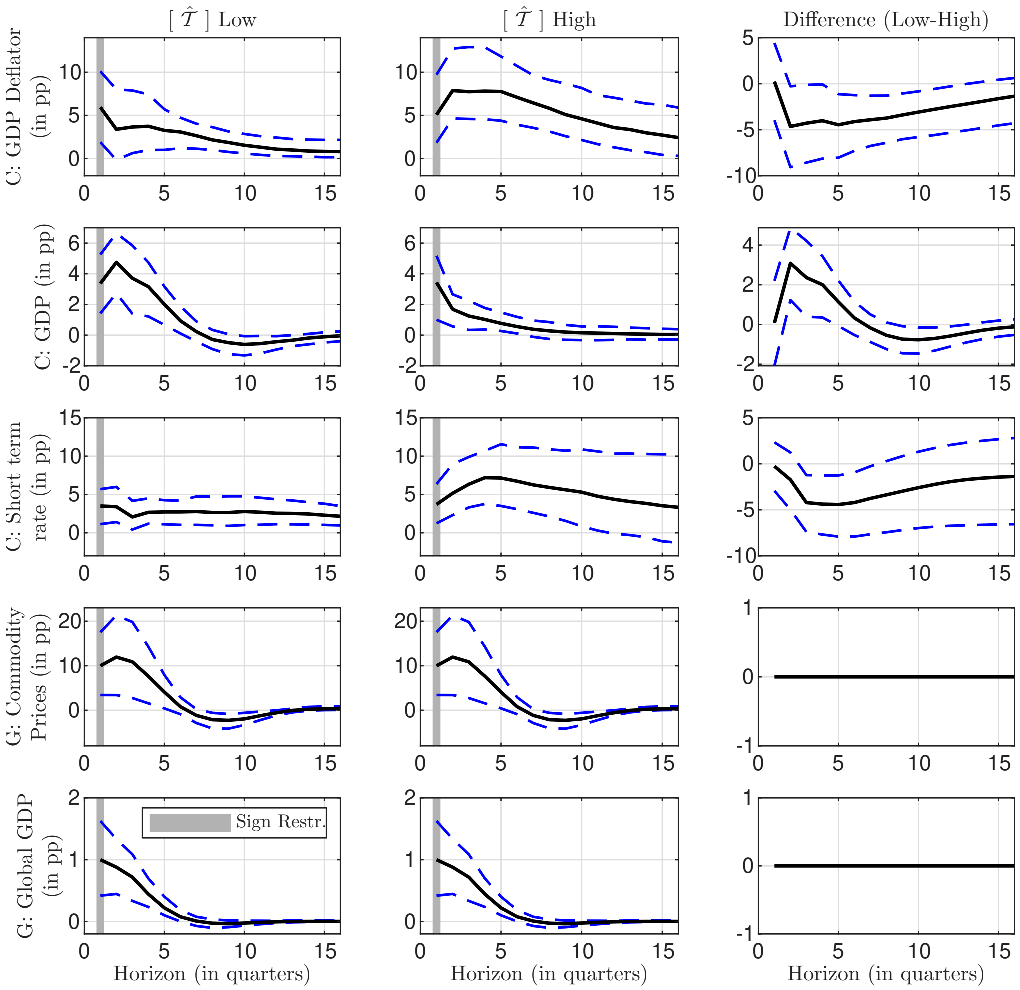

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the temperature anomaly taken from the Berkeley Earth database. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$.

anomalies: the differences between IRFs under high and low anomalies are statistically indistinguishable from zero (see third-column panels in Figure 5). This outcome is consistent with the baseline results obtained using FAOSTAT data.

More generally, the degree of statistical significance in the differences between IRFs for high and low temperature anomaly regimes is greater when using the FAOSTAT temperature anomaly data than when using the Berkeley Earth data.

Figure 5. Supply shock using Berkeley Earth temperature anomaly

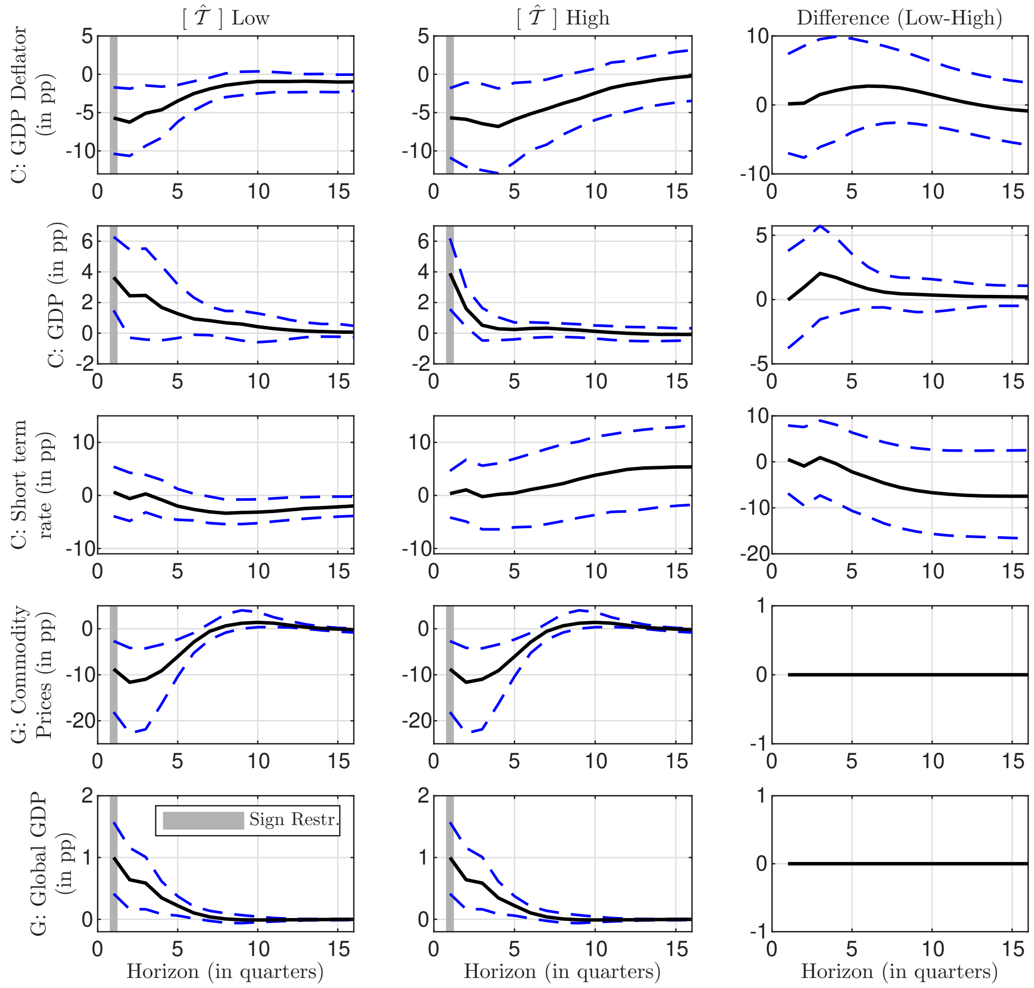

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the temperature anomaly taken from the Berkeley Earth database. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$.

5.2. **Using rainfall anomalies instead of temperature anomalies.** We now report and discuss results obtained when replacing the temperature anomaly with the rainfall anomaly as the interaction variable in the interacted Bayesian panel VAR (IPVAR) introduced in the main part of the paper. We implement exactly the same empirical exercise as in the baseline specification; the only modification is the choice of the interaction variable, which is now the country-level rainfall anomaly. As before, we use the estimated IPVAR to simulate global demand and supply shocks and, crucially, we assess whether the resulting impulse response functions (IRFs) vary systematically with the level of the interaction variable.

FIGURE 6. Demand shock (based on rainfall anomaly instead of temperature anomaly as climate change measure)

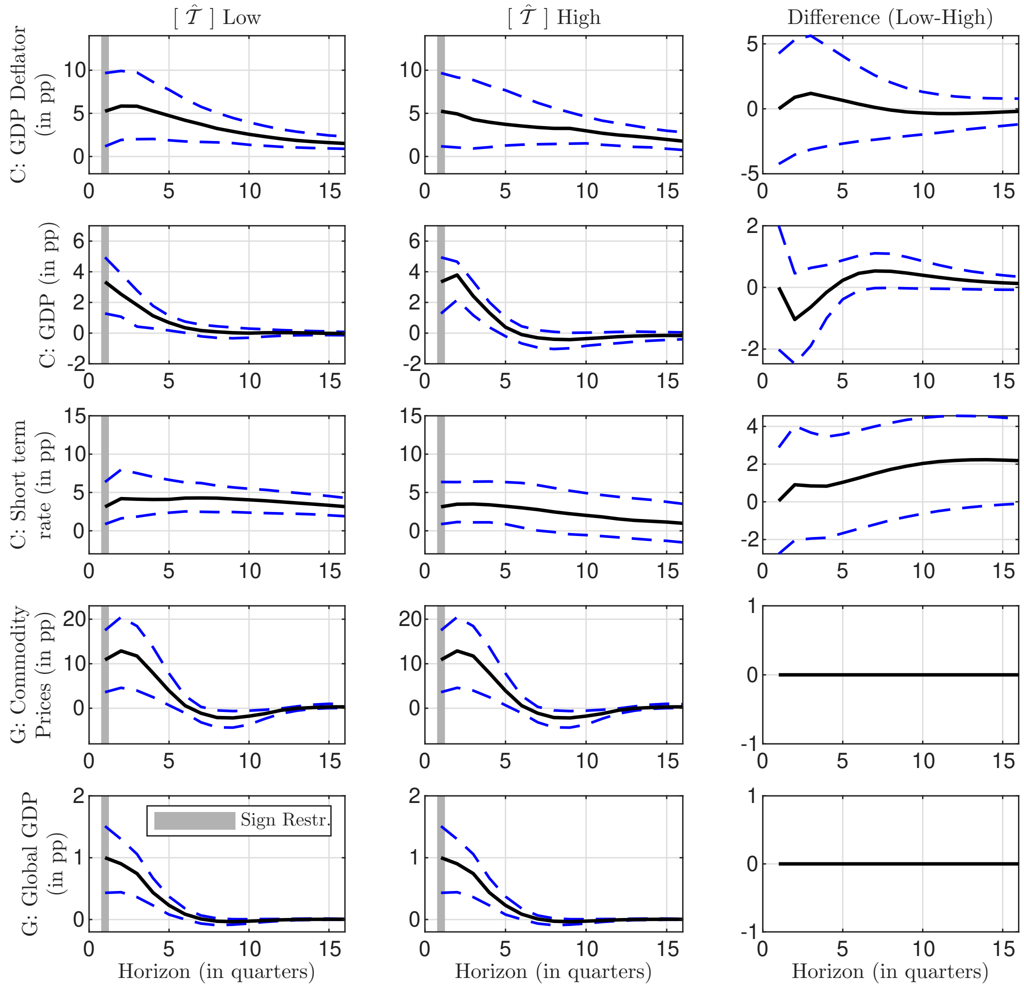

Note: The figure shows the sensitivity of the IRFs of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the rainfall anomaly. The values for the upper and lower bound of rainfall anomaly variable are given by the 95th and 5th percentile of the distribution as shown in the right-hand side panel in Figure 3. The solid line represents the posterior median, while the dashed lines denote the 68 percent credible interval.

Figure 6 presents the results for the global demand shock and figure 7 for the global supply shock. The shock is identified in the same way as in the baseline analysis, and the figures report the responses of global output and global inflation, as well as the average country-level responses of inflation, output, and the nominal interest rate. IRFs are shown conditional on low versus high values of the rainfall anomaly: the low state corresponds to the 5th percentile of the pooled rainfall anomaly distribution (across all countries and the full sample period), whereas the high state corresponds to the 95th

FIGURE 7. Supply shock (based on rainfall anomaly instead of temperature anomaly as climate change measure)

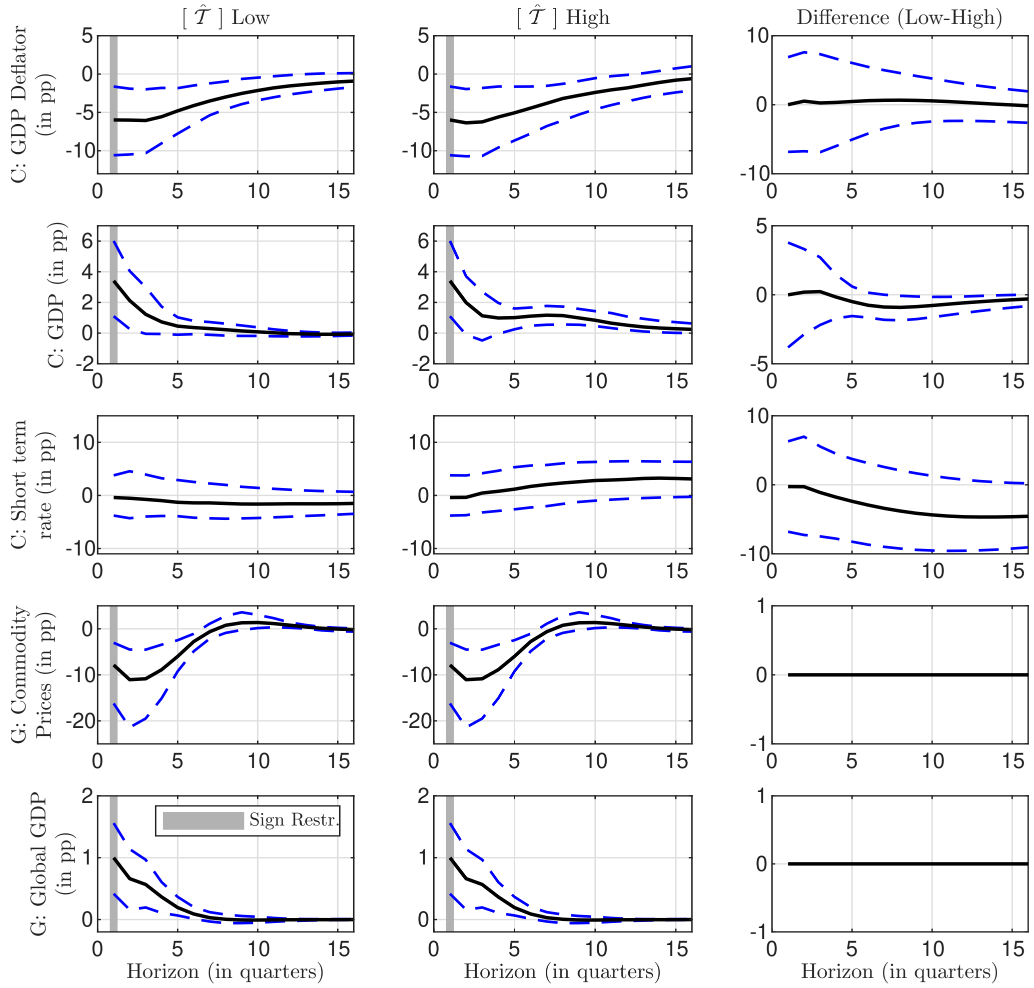

Note: The figure shows the sensitivity of the IRFs of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the rainfall anomaly. The values for the upper and lower bound of rainfall anomaly variable are given by the 95th and 5th percentile of the distribution as shown in the right-hand side panel in Figure 3. The solid line represents the posterior median, while the dashed lines denote the 68 percent credible interval.

percentile. Both percentile thresholds are also marked in the right-hand panel of Figure 3.

The estimated IRFs indicate that an expansionary global demand shock leads to a pronounced increase in inflation and output at the average country level. These responses are statistically significant and display substantial persistence. Importantly, this qualitative pattern is essentially identical under both low and high rainfall anomaly realizations. Consistent with this, the state-dependent differences in the IRFs (third-column panels in Figure 6) are not statistically distinguishable from zero: the corresponding credible intervals overlap throughout the horizon in all cases. Overall, these results suggest that rainfall anomalies do not materially affect the transmission of global demand shocks in our framework. This finding contrasts with the baseline specification based on temperature anomalies and thereby strengthens the interpretation that temperature-based measures capture the relevant state dependence associated with climate conditions, in line with the arguments in Manne et al. (1995); Held and Soden (2006); Hassler and Krusell (2012); Bilal and Stock (2025).

Figure 7 corroborates the baseline results obtained using the temperature anomaly, now employing the rainfall anomaly as the interaction variable. While output and inflation at the (average) country level respond to the global supply shock in a manner that is statistically significantly different from zero, these responses exhibit no systematic dependence on the rainfall anomaly. This pattern closely parallels the baseline findings for the temperature anomaly. Accordingly, the rainfall anomaly does not meaningfully shape the transmission of the global supply shock: the implied state dependence remains negligible, and the differences in IRFs between low and high rainfall-anomaly realizations are not statistically distinguishable from zero. In this sense, the rainfall-based specification mirrors the baseline conclusion that supply-shock transmission is not materially affected by the interaction variable. The full set of supply-shock results for the rainfall-anomaly specification is available upon request.

5.3. **The role of higher order Taylor approximations .** The baseline results were derived from a model that ignored higher order terms in the impact of climate change on the economic system. However, Burke et al. (2015); Chavas et al. (2016); Casey et al. (2023), among others, emphasize the importance of accounting for higher order terms (or non-linearities in general) when evaluating such effects. In light of their findings, we now extend the analysis to incorporate higher order terms. To this end, this modifies the equation for the IVAR model as follows:

$$\text{IVAR (non-linear):} \quad \boldsymbol{B}_{11}^{c}(\mathcal{T}_t) = \boldsymbol{\Gamma}^{c}(\bar{\mathcal{T}}) + \sum_{p=1}^{\mathcal{R}} \boldsymbol{\mathcal{F}}_{p}^{c} \hat{\mathcal{T}}_t^{p} \tag{6}$$

Figure 8. Considering non-linearities

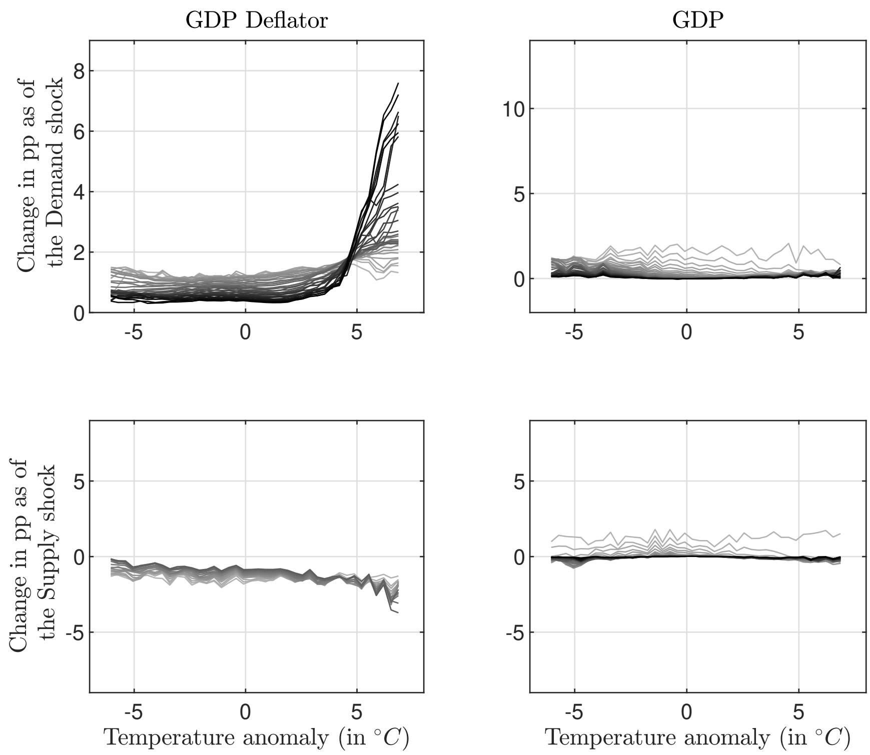

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent) and output (GDP, yoy change in percent) to a global demand (upper sub-panels) and a global supply shock (lower sub-panels) with respect to the temperature anomaly. The bright gray lines represent short-horizon IRFs, with progressively darker lines indicating longer horizons, culminating in black lines that capture the 16-quarter horizon.

In principle, any polynomial could be applied here, with the only limitation being the degrees of freedom in the estimation. Given that the temperature anomaly includes both positive and negative values, we adopt a cubic specification and set $\mathcal{P} = 3$, which conforms with a cubic Taylor expansion around the average level of the temperature anomaly ($\hat{\mathcal{T}}$). We then substitute equation (6) into the baseline VAR and re-estimate the model using Bayesian techniques.

The results are illustrated in Figure 8. We present impulse response functions (IRFs) in the same graphical format as in the main part, which illustrates home-country price and output responses to global demand and supply shocks under varying degrees of temperature anomalies. As before, the bright gray lines represent short-horizon IRFs, with progressively darker lines indicating longer horizons, culminating in black lines that capture the 16-quarter horizon.

Figure 9. Dynamic stability: linear vs. cubic

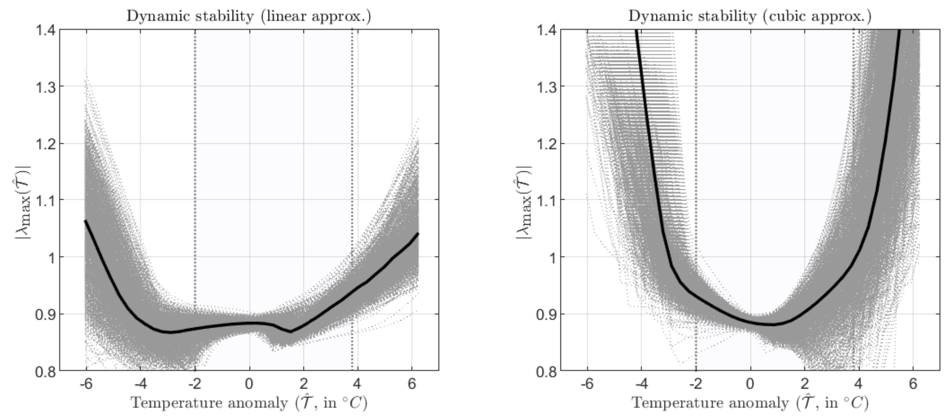

Note: The figure shows the posterior distribution of the absolute value of the maximum eigenvalue of the companion matrix of the IPVAR model as a function of the temperature anomaly.

The second-row sub-panels depict the IRFs of prices and output in response to the supply shock. Prices decline while output rises, yet the IRFs remain relatively flat across different temperature anomalies, indicating that the anomaly does not significantly influence price or output responses to a supply shock. Thus this finding, discussed earlier, is reinforced when non-linearities are considered.

The first-row sub-panels in Figure 8 display the IRFs corresponding to the global demand shock. Both price and output reactions are positive across all horizons, as indicated by the range from bright gray to black lines. Contrary to the baseline results, the output response now appears rather flat with respect to temperature anomalies, suggesting a diminished sensitivity to these anomalies. However, the price response becomes markedly more sensitive to temperature anomalies, particularly for higher positive values. This sensitivity is observed across short and long horizons, with the magnitude of the price response in the non-linear setting significantly exceeding that of the linear model. We interpret this as further support for our baseline findings and as evidence of non-linear effects of temperature anomalies—implying that climate change exerts non-linear influences on the economic system.

As a further check, we examine in detail the stability of the IPVAR model. This is motivated by the observation that the inclusion of non-linearities into the IPVAR model quickly poses challenges regarding the stability of the dynamic system. The stability of the IPVAR model is determined by the eigenvalues of the companion matrix (see Hamilton, 1994, Chapter 10). Since the matrix lag polynomial of the IPVAR model depends on the temperature anomaly, the eigenvalues of the companion matrix hence directly depend on the temperature anomaly.

We illustrate the sensitivity of the eigenvalues with respect to the temperature anomaly in Figure 9. The figure displays the posterior density of the absolute value of the maximum eigenvalue of the companion matrix of the IPVAR model as a function of the temperature anomaly. This is shown for both the baseline model (left-hand side sub-panel), which assumes a linear relationship between the temperature anomaly and the matrix lag polynomial, and for a non-linear extension (right-hand side sub-panel), which considers a cubic relationship, as outlined in equation (6). The thick black solid line represents the median across the draws in each case. Both sub-panels indicate the values of the temperature anomaly corresponding to the 5th and 95th percentiles of the respective histogram (or density distribution) of the temperature anomaly across countries and time.

The left-hand side sub-panel highlights that eigenvalues (in absolute value terms) larger than unity emerge primarily at the extremes of the temperature anomaly distribution, i.e., at values that occur infrequently. For most of the variation in the temperature anomaly, the linear IPVAR model remains stable, with all the draws of the maximum eigenvalue below unity. Even for values beyond the 5th and 95th percentiles of the respective histogram, both stable and unstable draws for the maximum eigenvalue occur, with the median slowly but steadily moving towards instability of the IPVAR model.

In stark contrast, the non-linear (i.e., cubic) system exhibits a much quicker transition towards instability. The non-linear system remains dynamically stable within a temperature anomaly range from -1.5 to 2, with all draws of the (absolute value) of the maximum eigenvalue below unity. However, outside this narrow temperature interval, the system quickly drifts toward instability, with a strikingly rapid and pronounced pace.

In the cubic model, the size of the maximum eigenvalue increases at a notably faster rate with the temperature anomaly than in the linear model. This

FIGURE 10. Demand shock: Different measures for the prices indices

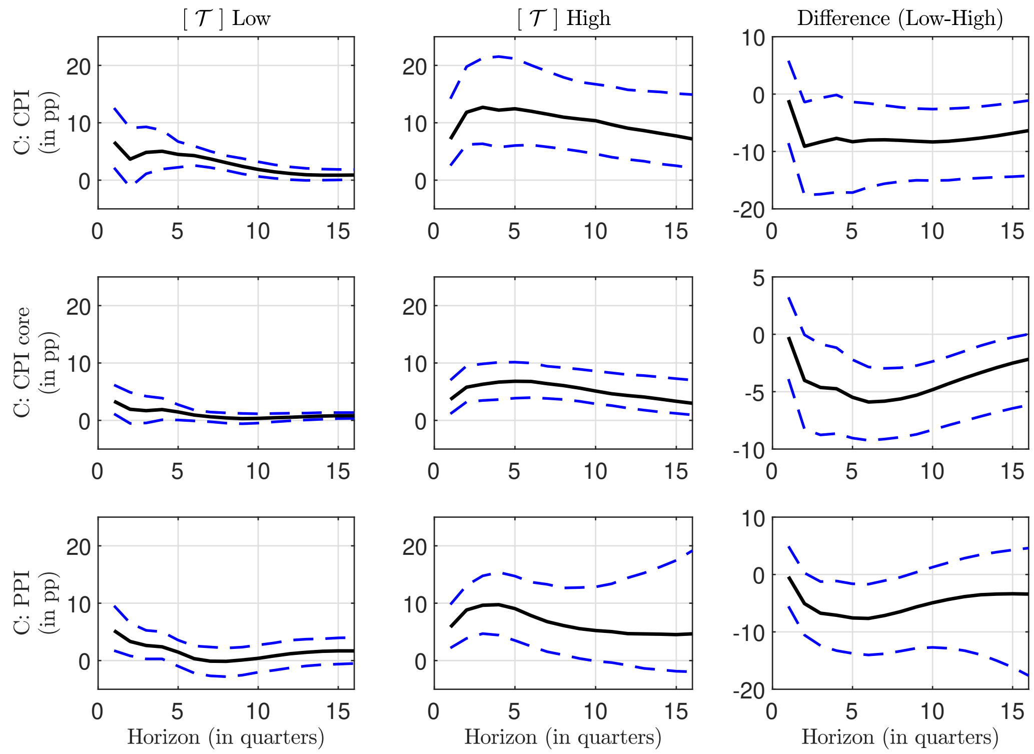

Note: The figure shows the sensitivity of the impulse response functions of various inflation measures to a global demand shock with respect to the temperature anomaly.

highlights the difficulty in assessing the effects of the temperature anomaly within the IPVAR model. On the one hand, this rapid transition to instability in the non-linear model makes the system harder to interpret, but on the other hand, it suggests that the linear model may underestimate the effects of the temperature anomaly. While in principle, higher-order interaction terms are possible, it has been documented by Mittnik (1990) and Granger (1998) that in this case the model (formally referred to a multivariate Generalized Vector Autoregressive, GAR, models) may become unstable when including squared or higher-powered interaction terms. In practice, we encounter precisely these issues when considering cubic interactions, and Caggiano et al. (2017) similarly warn against model instability in variants featuring multiple interaction terms.

5.4. **Different measures for the price indices.** This section extends the baseline results by assessing their robustness with respect to different price measures. While the baseline model employs the GDP deflator as the primary price measure, we now substitute this variable with three alternatives: (i) the Consumer Price Index (CPI), (ii) the core CPI, and (iii) the Producer

FIGURE 11. Supply shock: Different measures for the prices indices

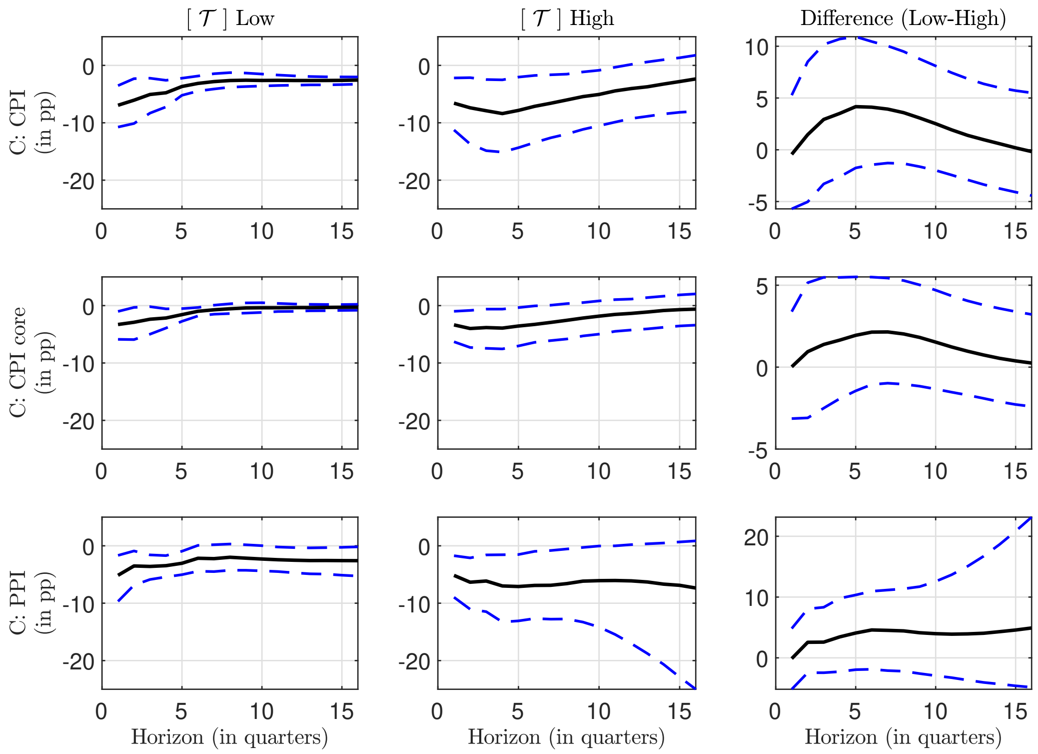

Note: The figure shows the sensitivity of the impulse response functions of various inflation measures to a global supply shock with respect to the temperature anomaly.

Price Index (PPI). In each case, we replace the GDP deflator with one of these alternative price measures in the baseline model and replicate the baseline analysis. Similar to the treatment of the GDP deflator, each alternative price measure (used at a quarterly frequency) is entered as a year-over-year percentage change. Our focus remains on the global demand shock and its implications for price responses, with the results displayed in Figure 10.

The sub-panels in the first row of Figure 10 show the responses of the CPI to a global demand shock, differentiated by a low (left-hand side) and a high (middle sub-panel) temperature anomaly, with the difference between these two states depicted in the right-hand sub-panel. In response to the expansionary demand shock, the CPI rises in both temperature anomaly states. Notably, the increase is more pronounced when the temperature anomaly is high. The right-hand sub-panel indicates that this difference is statistically significant, corroborating the baseline findings.

The second row of sub-panels presents the results for the core CPI. As with the CPI, prices increase following the expansionary global demand shock, with the magnitude of the increase being greater when the temperature anomaly is high. We observe a similar pattern for the PPI, where price responses are also more substantial under high temperature anomalies. These results provide further evidence supporting the robustness of the baseline findings concerning price reactions to expansionary global demand shocks. More importantly, they highlight the significant role of temperature anomalies in shaping these price responses, as their impact remains statistically significant across all price measures.

Figure 11 reports the corresponding analysis for the global supply shock. As in the previous exercise, we replace the baseline price measure (the GDP deflator) with each of the three alternative price indicators in turn, and re-estimate the model sequentially for each. The figure presents the resulting impulse response functions of the price measures to an (expansionary) supply shock, conditional on different levels of the temperature anomaly. In all cases, the immediate and short-term responses of domestic prices are negative, both on impact and for several subsequent quarters. This pattern holds irrespective of whether the temperature anomaly is low or high. Crucially, the differences in the price responses across the two temperature regimes are not statistically distinguishable from zero, thereby reinforcing the conclusions drawn from the baseline specification.

5.5. **Different measures for real activity.** This section extends the baseline analysis by evaluating the robustness of the results with respect to different measures of real economic activity. While the baseline results employ GDP as the primary indicator of real activity, we now replace this variable with two alternatives: (i) the industrial production index (IP) and (ii) the number of unemployment people. Both series are sourced from the OECD Main Economic Indicators database. In each case, we substitute GDP with one of these real activity measures in the baseline model. Consistent with the treatment of GDP, each alternative measure (used at a quarterly frequency) is entered as a year-over-year percentage change. The same exercise as conducted for the baseline results is replicated here, with the focus on the demand shock and its implications for the various real activity measures. The results are presented in Figure 12.

FIGURE 12. Demand shock: Different measures for output

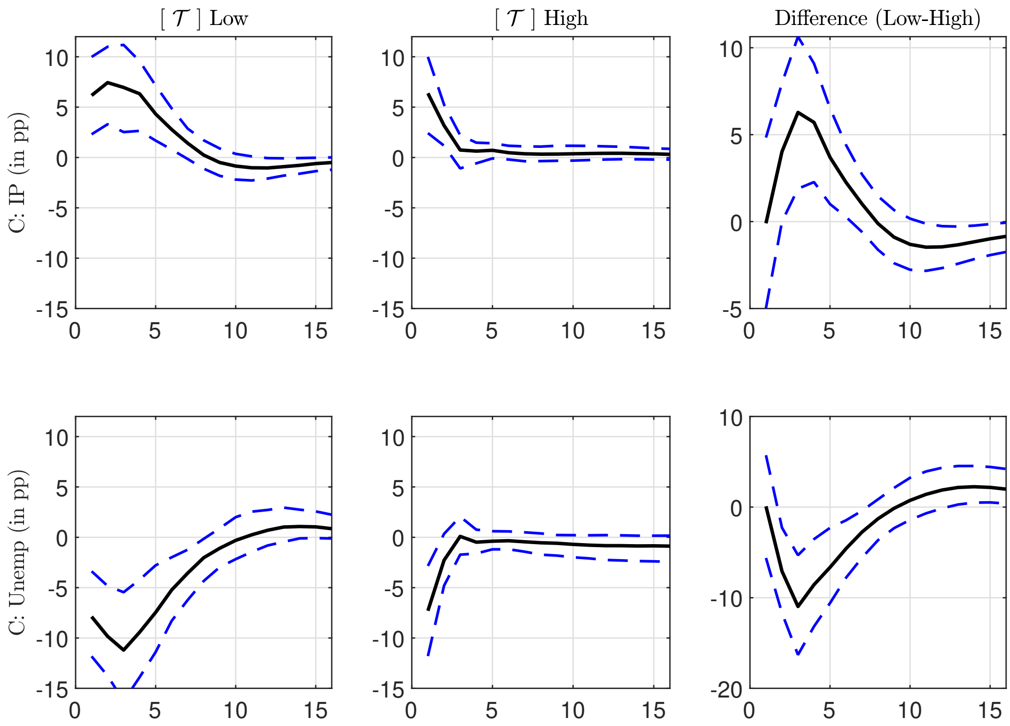

Note: The figure shows the sensitivity of the impulse response functions of various output measures to a global demand shock with respect to the temperature anomaly.

The sub-panels in the first row of Figure 12 illustrate the responses of IP to the global demand shock, with results disaggregated by low (left-hand side) and high (middle sub-panel) values of the temperature anomaly. The difference between these two states is depicted in the right-hand sub-panel. As shown, the expansionary demand shock leads to an increase in IP. Notably, the increase is more pronounced when the temperature anomaly is low. The right-hand sub-panel highlights that this difference is statistically significant, thereby confirming the baseline results, though now with IP as the measure of real activity rather than GDP.

The second row of sub-panels displays the results for unemployment. Once again, we observe a significant response in real activity. In this case, unemployment decreases in response to the expansionary global demand shock. Importantly, the decline in unemployment is larger when the temperature anomaly is low.

These results provide further validation of the baseline findings concerning the real activity response to expansionary global demand shocks. More importantly, they underscore the significant role of the temperature anomaly in

FIGURE 13. Supply shock: Different measures for output

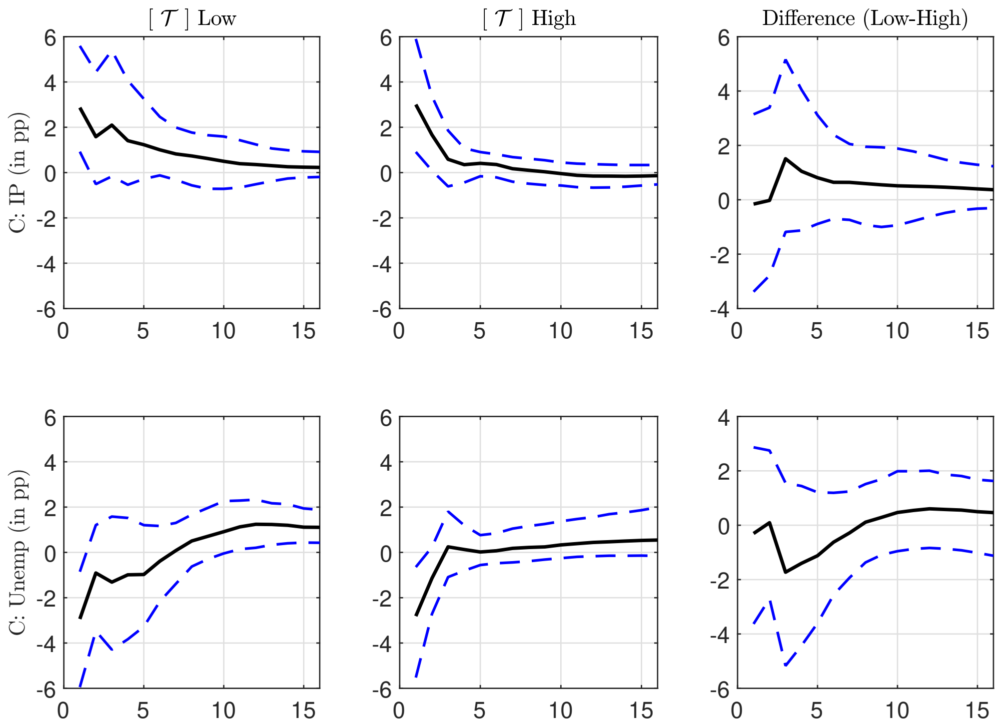

Note: The figure shows the sensitivity of the impulse response functions of various output measures to a global supply shock with respect to the temperature anomaly.

shaping these responses, as its impact remains statistically significant in both alternative real activity measures.

Figure 13 presents the corresponding analysis for the global supply shock using alternative output measures. In analogy to the previous exercise, we replace the baseline output proxy (GDP) with each of the two additional indicators in turn and repeat the estimation sequentially. The figure reports the resulting impulse response functions of the output measures to an (expansionary) global supply shock, conditional on different levels of the temperature anomaly. For industrial production, domestic output reacts positively on impact and remains elevated for several subsequent quarters. In contrast, for unemployment, the response is negative on impact and continues to be so in the quarters that follow. These results apply under both low and high levels of the temperature anomaly. Most importantly, the differences across the two climate regimes are not statistically significant, thereby substantiating the baseline findings.

5.6. **Omitted variables.** Separate Ljung-Box tests on each residual time series cannot reject the null hypothesis that they follow processes which are uncorrelated over time. However, it is still possible that omitted variables matter for the results. To check whether the two identified global shocks are correlated with other (omitted) variables, we first follow Glocker and Towbin (2015) and compute correlations of the estimated structural disturbances with variables that a large class of general equilibrium models suggests as being jointly generated by various shocks.

We compute correlations of up to six leads and lags between the global demand and supply shock and the growth rate of the domestic stock market indices, employment, consumption, the Brent oil price[2], global industrial production, the price of natural gas (according to Dutch TTF), the domestic long-term government bond rates and a global stock market index.

The cross-correlations indicate that none of the omitted variables correlates significantly with the structural shocks. The statistical importance of the cross-correlations has been judged by means of the upper and lower limits of an asymptotic 95 percent confidence tunnel for the null hypothesis of no cross-correlation.

5.7. **Alternative identification approach.** In this section, we explore an alternative identification approach. While this approach still relies on sign restrictions too, it uses an approach as proposed by Glocker and Piribauer (2025). We start from the IPVAR model, and hence the average cross-country coefficient matrices, to identify the model. Let its structural form now be given by

$$\mathbf{A}\boldsymbol{y}_t = \sum_{j=1}^{\mathcal{P}} \mathbf{B}_j(\hat{\mathcal{T}})\boldsymbol{y}_{t-j} + \boldsymbol{u}_t \tag{7}$$

where $\boldsymbol{u}_t$ having a diagonal variance-covariance matrix. Identification requires imposing restrictions on the matrix $\mathbf{A}$. Following Glocker and Piribauer (2025), we impose sign restrictions on each row of $\mathbf{A}$ in accordance with the theoretical model. As regards the global part ($\boldsymbol{g}_t$: global commodity price index, global GDP), we consider an aggregate supply and an aggregate demand relation. As regards the country part ($\boldsymbol{y}_{c,t}$: GDP deflator, GDP, short-term

[2]We use the cyclical component of the oil price obtained after applying the Christiano-Fitzgerald filter on the logarithm of the oil price.

interest rate), this suggests to use (i) the Phillips curve, (ii) the consumption-Euler equation in conjunction with the aggregate resource constraint, and (iii) the Taylor rule. This motivates the following sign restrictions on the elements of matrix $\mathbf{A}$ along the rows

$$\text{sgn}(\mathbf{A}) = \begin{bmatrix} + & - & \cdot & \cdot & \cdot \\ + & + & \cdot & \cdot & \cdot \\ \cdot & \cdot & + & - & \cdot \\ \cdot & \cdot & + & \cdot & + \\ \cdot & \cdot & - & - & + \end{bmatrix} \tag{8}$$

where a dot $(\cdot)$ indicates an unrestricted element, and the symbols "+" and "−" denote the sign restrictions imposed on the corresponding elements. The first two rows identify the global supply and demand relations, the third to fifth rows the (country-level) Phillips relation, consumption-Euler equation and the Taylor rule.

We use the eigenvector-eigenvalue decomposition of the variance-covariance matrix $\mathbf{\Omega}$ and define $\mathbf{V} = \mathbf{E}_{\mathbf{\Omega}}\mathbf{\Lambda}_{\mathbf{\Omega}}^{1/2}$, where $\mathbf{E}_{\mathbf{\Omega}}$ is the (orthogonal) matrix of eigenvectors and $\mathbf{\Lambda}_{\mathbf{\Omega}}$ is the (diagonal) matrix of eigenvalues. In view of the properties of this decomposition in case of symmetric matrices ($\mathbf{E}'_{\mathbf{\Omega}} = \mathbf{E}_{\mathbf{\Omega}}^{-1}$, since $\mathbf{E}_{\mathbf{\Omega}}$ is an orthogonal matrix), we have that $\mathbf{\Omega} = \mathbf{V}\mathbf{V}'$. We then multiply $\mathbf{V}$ with an orthonormal matrix $\mathbf{Q}$ to obtain $\mathbf{H} = \mathbf{V}\mathbf{Q}$, from which we can compute a candidate for $\mathbf{A} = \mathbf{H}^{-1}$. The matrix $\mathbf{Q}$ is a random orthonormal matrix of the same dimension as $\mathbf{\Omega}$. We follow Rubio-Ramírez et al. (2010) and compute $\mathbf{Q}$ by drawing an independent standard normal matrix $\mathbf{X}$ (of the same size as $\mathbf{\Omega}$), and apply the QR decomposition[3] $\mathbf{X} = \mathbf{Q}\mathbf{R}$ to get $\mathbf{Q}$.

Regarding the sign restrictions, if the corresponding entries in a candidate matrix $\mathbf{A}$ are consistent with the sign restrictions outlined in (8), we retain the candidate draw and proceed to the next draw until a total of 5,000 accepted draws is obtained. Otherwise, a new $\mathbf{Q}$ matrix is drawn until the sign restrictions are satisfied. We report the median and the 16th and 84th percentile of the distribution as confidence bands.

[3] We note that while the estimation and the inference of the reduced-form model rely on a rather non-informative prior for all coefficients, an informative prior emerges in the structural model due to the QR decomposition (Baumeister and Hamilton, 2015).

Figure 14. Global demand shock – alternative identification

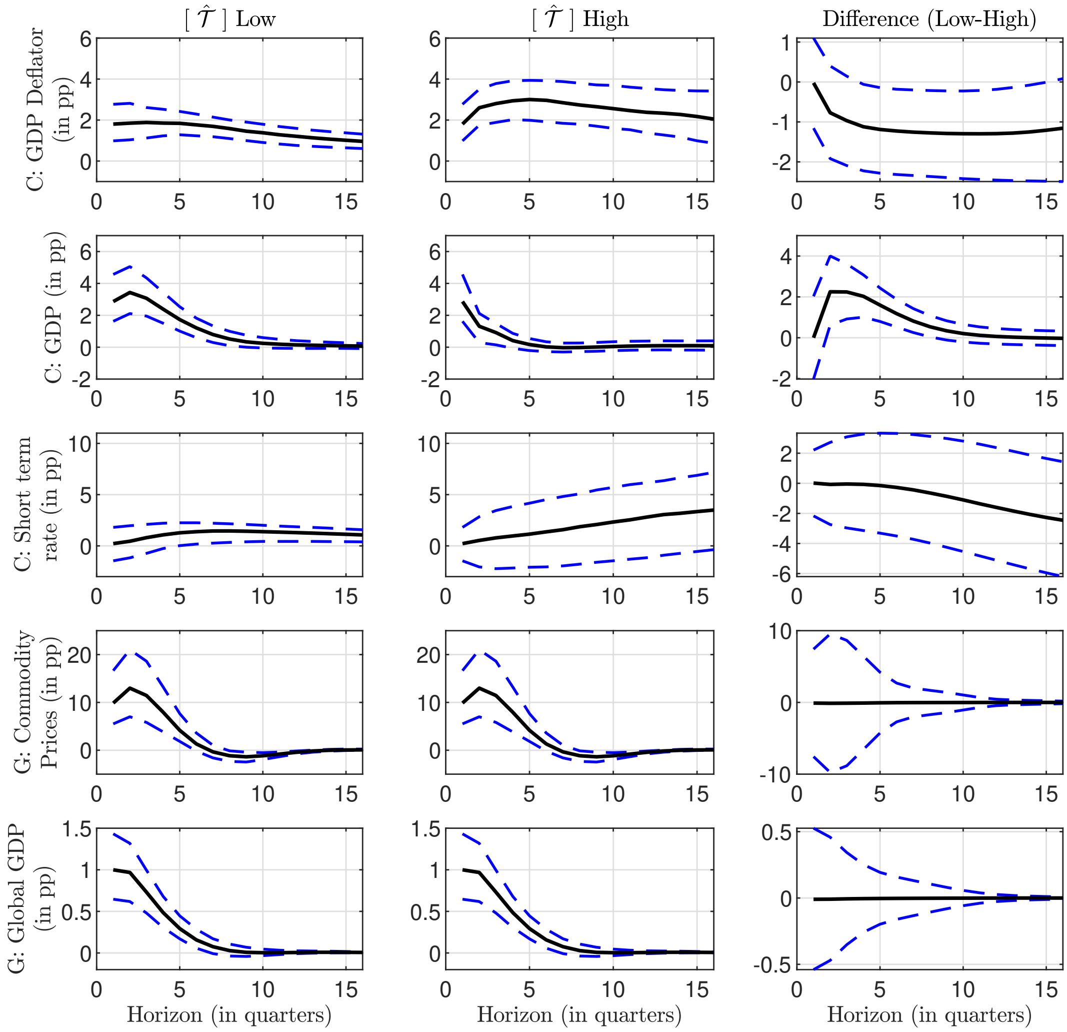

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the temperature anomaly.

We now revisit the analysis of global supply and demand shocks, as represented by the first and second elements of $\mathbf{u}_t$, respectively, to examine their effects at the country level. More importantly, we assess the extent to which the transmission mechanism of these shocks is influenced by different levels of the temperature anomaly $\hat{\mathcal{T}}$. The results for both shocks are presented in Figures 14 and 15.

Starting with the global demand shock, Figure 14 illustrates that the global demand shock raises global GDP and global commodity prices, which is accompanied by increases in domestic prices, GDP, and the interest rate. These responses are statistically significant and exhibit a high degree of inertia. Of particular interest is the observation that the impulse responses at the home-country level are strongly shaped by the temperature anomaly. The sub-panels in the left-hand column display the responses when the temperature anomaly is low, those in the middle show responses at a high temperature anomaly level, and the right-hand column illustrates the differences between the two. As shown, the price response is significantly larger when the temperature anomaly is high, whereas the output response is notably weaker in this case. This pattern replicates the baseline results from the main analysis, albeit with a different method of implementing the sign restrictions. Notably, as highlighted in the sub-panels of the third column, the differences in impulse responses across low and high temperature anomaly levels are statistically significant (with 68% error bands). This underscores that the transmission mechanism of demand shocks is substantially influenced by the temperature anomaly, and thus by climate change.

Turning to the global supply shock, Figure 15 presents the results. We consider an expansionary global supply shock that causes a significant decline in global commodity prices, accompanied by a rise in global GDP. At the home-country level, prices decline, and GDP increases. While the price responses at the home-country level are statistically significant, the output responses are only borderline significant. More importantly, we find no evidence suggesting that the temperature anomaly affects the transmission of supply shocks. In other words, the temperature anomaly does not influence the impulse response functions for supply shocks, which corroborates the findings from the main analysis. This result further reinforces the notion that temperature anomalies, and by extension climate change, affect the slope of the Phillips curve, thereby altering the effects of demand shocks while leaving the effects of supply shocks unchanged.

5.8. **Alternative panel methodology: Pooled Bayesian Panel VAR Model.** In this section, we examine the robustness of our baseline findings by considering an alternative econometric approach: a pooled Bayesian panel VAR model. This specification contrasts with the mean-group estimator employed in our main analysis and provides an alternative perspective on the role of temperature anomalies in shaping macroeconomic dynamics. In the pooled approach, all home-country observations are stacked, and a common set of

FIGURE 15. Global supply shock – alternative identification

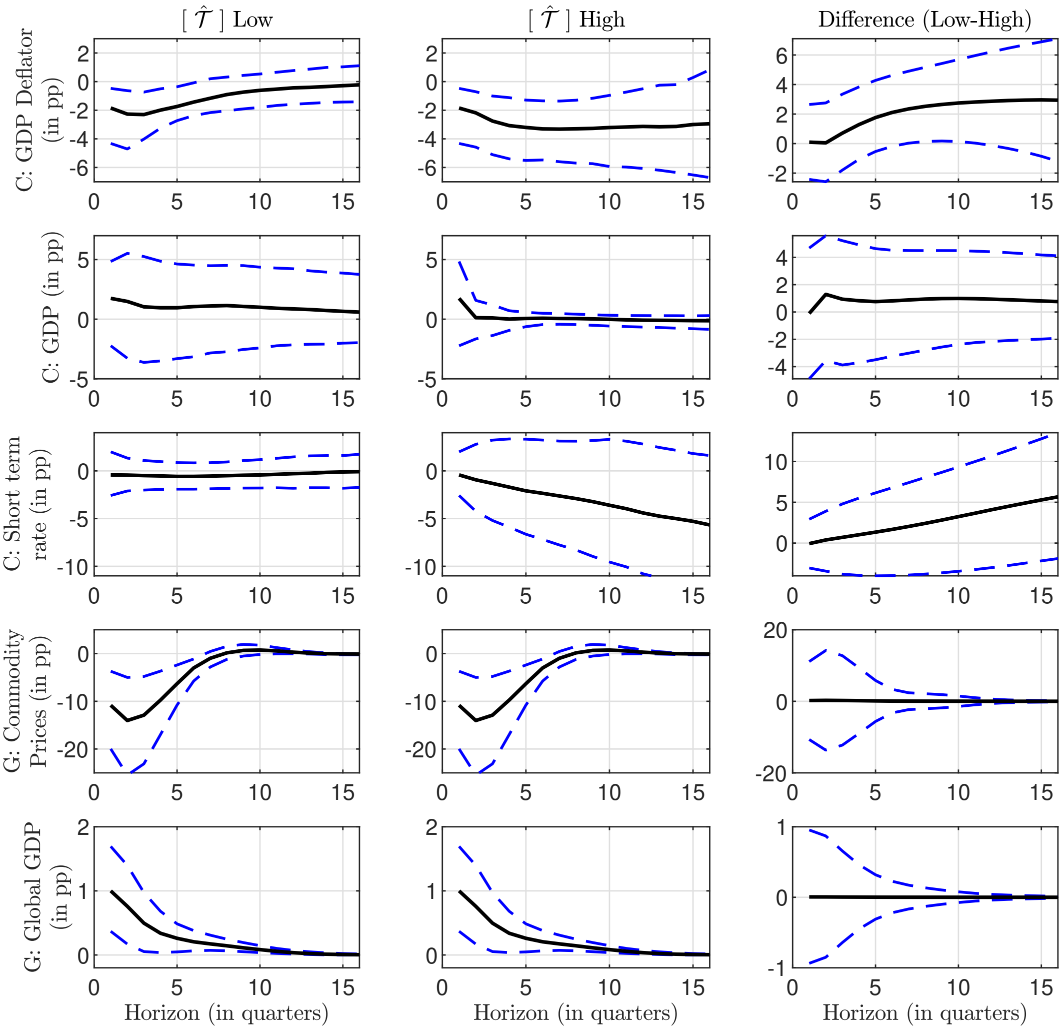

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the temperature anomaly.

autoregressive coefficients is estimated for the entire panel, thus imposing parameter homogeneity across countries. This contrasts with the mean-group approach, which estimates home-country parameter matrices and computes the impulse response functions (IRFs) as the average response across countries.

5.8.1. *Model Specification.* The pooled Bayesian panel VAR model maintains the same structural framework as equation (31) in the main text, but imposes coefficient homogeneity across countries. The model specification is given by:

$$
(9) \qquad \begin{bmatrix} \underset{(n\times 1)}{\boldsymbol{y}_{c,t}} \\ \underset{(2\times 1)}{\boldsymbol{y}_{g,t}} \end{bmatrix} = \begin{bmatrix} \underset{(n\times n)}{\boldsymbol{B}_{11}(\mathcal{T})} & \underset{(n\times 2)}{\boldsymbol{B}_{12}} \\ \underset{(2\times n)}{\boldsymbol{0}} & \underset{(2\times 2)}{\boldsymbol{B}_{22}} \end{bmatrix} \begin{bmatrix} \underset{(n\times 1)}{\boldsymbol{y}_{c,t-1}} \\ \underset{(2\times 1)}{\boldsymbol{y}_{g,t-1}} \end{bmatrix} + \begin{bmatrix} \underset{(n\times 2)}{\boldsymbol{H}_1(\mathcal{T})} \\ \underset{(2\times 2)}{\boldsymbol{H}_2} \end{bmatrix} \underset{(2\times 1)}{\boldsymbol{u}_t}
$$

where the subscript c denotes again country $c = 1, ..., N$, and the coefficient matrices do not vary across countries by construction. We use the same set of exogenous variables as in the main text. The temperature interaction structure follows the same first-order Taylor approximation as in the baseline model:

$$
(10) \qquad B_{11}(\hat{\mathcal{T}}_{c,t}) = \Gamma(\bar{T}_c) + F\hat{\mathcal{T}}_{c,t}
$$

where Γ and F are $n \times n$ matrices of coefficients of the pooled model, and $\hat{\mathcal{T}}_{c,t}$ represents the home-country temperature anomaly for country c at time t.

The key difference from the baseline specification lies in the treatment of cross-country heterogeneity. While the mean-group estimator allows for country-specific coefficients $B_j^c(\hat{T})$ and subsequently averages them as shown in equation (38) of the main part, the pooled model directly estimates common coefficients across all countries.

5.8.2. *Methodological Considerations.* The pooled approach offers both advantages and disadvantages relative to the mean-group estimator. On the positive side, pooling increases the effective sample size by utilizing information from all countries simultaneously, potentially leading to more precise parameter estimates and narrower confidence intervals. This efficiency gain is particularly valuable when individual country time series are relatively short or when cross-country parameter heterogeneity is limited.

However, the pooled specification comes with important limitations. Most notably, it imposes the restrictive assumption of coefficient homogeneity across countries, which may not be empirically justified given differences in economic structures, institutional frameworks, and climatic conditions. If true heterogeneity exists, pooled estimates may suffer from aggregation bias and provide misleading inference about the average effects of temperature anomalies (see Canova and Ciccarelli, 2013, for further details).

The mean-group estimator used in our baseline analysis addresses this concern by allowing for country-specific coefficients while still providing interpretable average effects. This approach is more robust to cross-country heterogeneity and follows the recommendation of Pesaran and Smith (1995) for dynamic panel models where slope coefficients may vary across units.

5.8.3. *Interpretation of Impulse Response Functions.* The interpretation of impulse response functions differs between the two approaches. In our baseline model using the mean-group estimator, the IRFs represent the "average country response" – that is, the arithmetic mean of home-country responses to global shocks. This interpretation explicitly acknowledges cross-country heterogeneity while providing a summary measure of typical economic responses.

In contrast, the pooled model generates IRFs that represent the response of a "representative country" with characteristics given by the pooled coefficients. These IRFs reflect the response that would be observed if all countries behaved according to the estimated pooled parameters. While this interpretation is more restrictive, it provides a direct estimate of the common response pattern, assuming such commonality exists.

The temperature interaction effects in the pooled model have a similar interpretation to the baseline case: the IRFs can be evaluated at different levels of temperature anomalies to assess how climate conditions influence the transmission of global shocks. However, these effects now represent the common temperature sensitivity across all countries rather than the average of home-country sensitivities.

5.8.4. *Empirical Results.* Using the pooled Bayesian panel VAR model, we conduct the same analysis as performed with the baseline mean-group specification. We compute impulse response functions for global demand and supply shocks, allowing for heterogeneity across distinct levels of temperature anomalies. The identification strategy remains identical, employing sign restrictions derived from our theoretical model as outlined in equation (36) in the main part.

Our findings strongly corroborate the results obtained with the baseline model. Temperature anomalies continue to exhibit no significant influence on the transmission mechanism of global supply shocks (see Figure 17). The impulse responses to supply shocks remain largely invariant to temperature

FIGURE 16. Pooled Panel BVAR: Demand shock

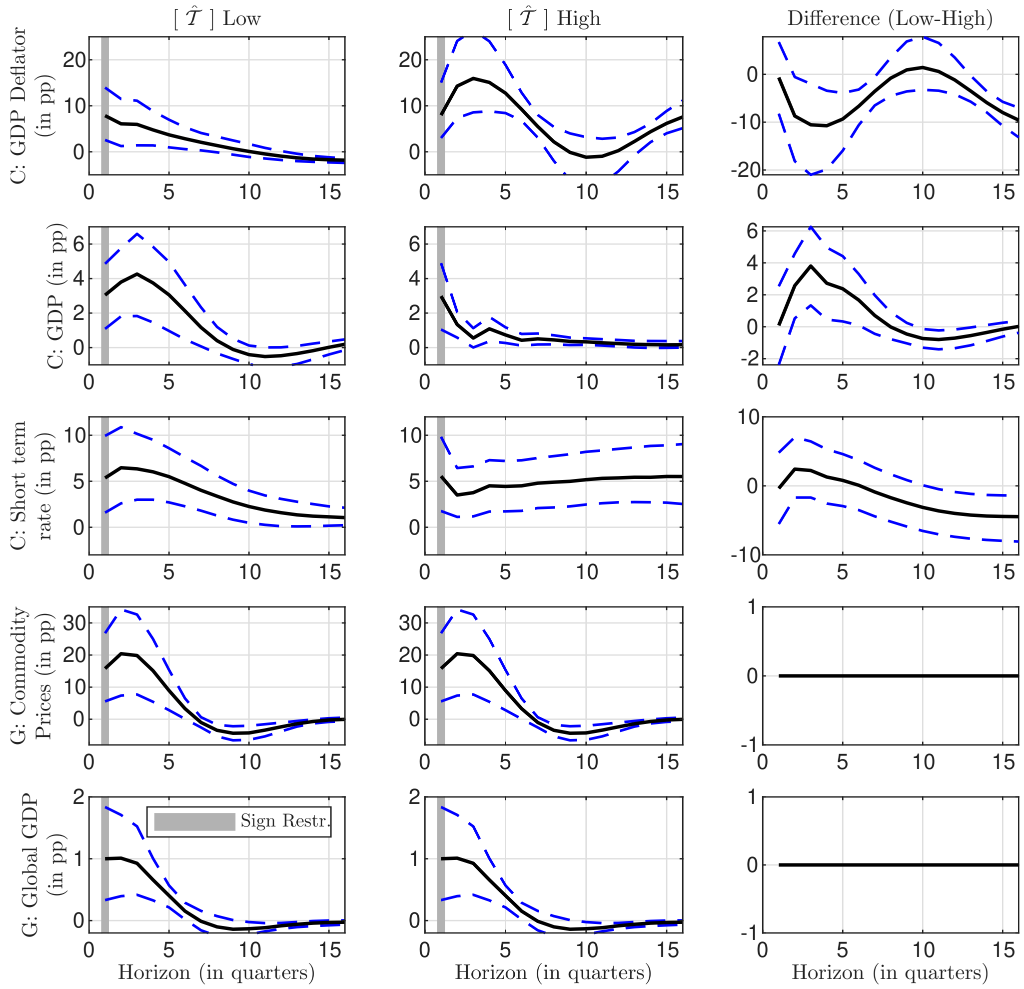

Note: The figure provides results based on an alternative empirical model: Pooled Bayesian Panel VAR Model with an interaction term. It shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the temperature anomaly. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$.

conditions, with confidence intervals that consistently encompass zero differences across temperature states.

For global demand shocks (see Figure 16), the pooled model confirms our core findings. Higher levels of temperature anomalies amplify the inflationary effects of demand shocks while simultaneously dampening the output

FIGURE 17. Pooled Panel BVAR: Supply shock

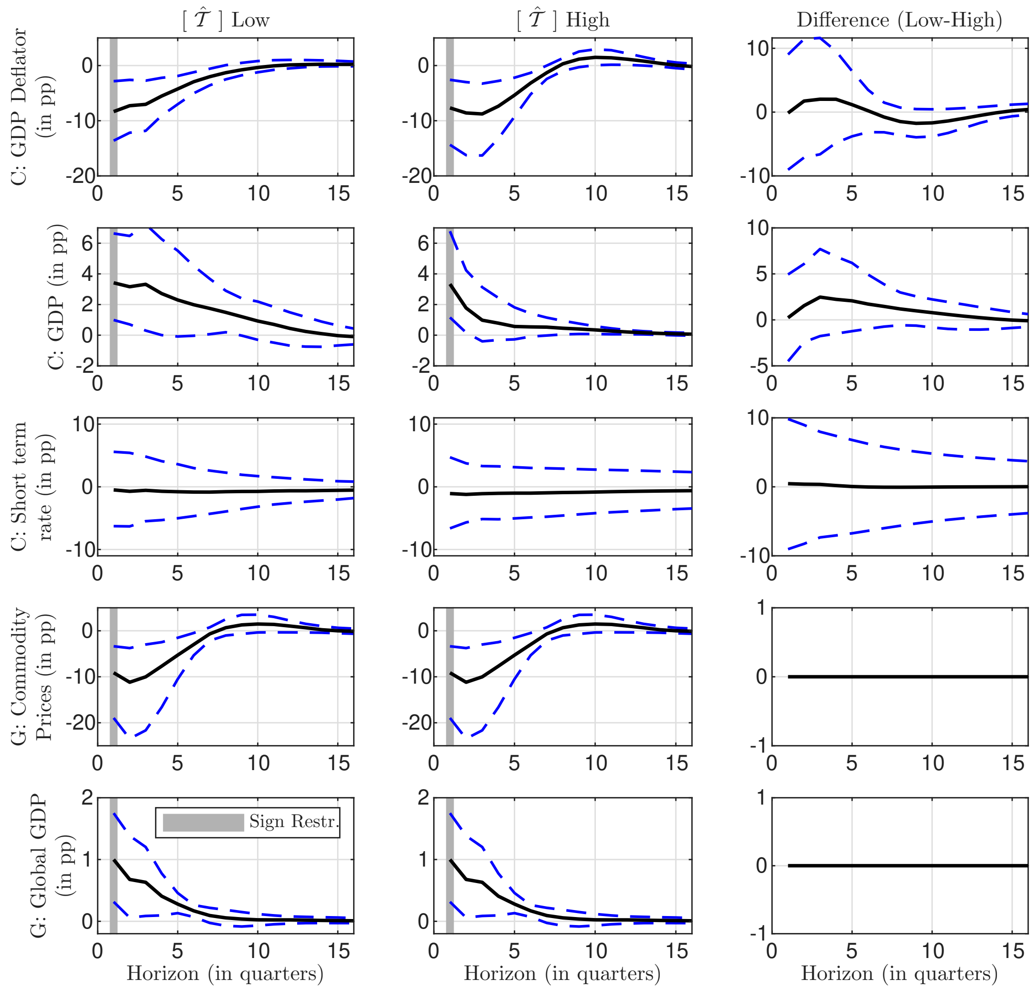

Note: The figure provides results based on an alternative empirical model: Pooled Bayesian Panel VAR Model with an interaction term. It shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the temperature anomaly. The temperature anomaly variable is defined by: $\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$.

responses. The magnitude and statistical significance of these temperature-dependent effects align closely with those reported in our main analysis. Specifically, elevated temperature anomalies steepen the effective Phillips curve relationship, transforming demand expansions into more pronounced stagflationary episodes.

The quantitative effects estimated under the pooled specification are comparable to those shown in Table 3 of the main text. A global demand shock normalized to raise global GDP by one percentage point continues to generate larger price responses and smaller output responses when temperature anomalies are elevated. These findings reinforce our interpretation that climate change fundamentally alters the propagation mechanism of demand shocks rather than acting as an independent source of business cycle fluctuations.

5.8.5. *Discussion.* We consider the consistency of results across these two distinct econometric approaches as a confirmation of the robustness of our baseline findings. The agreement between the mean-group and pooled specifications suggests that our conclusions are not artifacts of a particular estimation methodology but reflect genuine economic relationships in the data.

The fact that both approaches yield qualitatively and quantitatively similar results provides additional confidence in our central hypothesis that temperature anomalies alter the slope of the Phillips curve, thereby influencing the transmission of demand shocks but leaving supply shock propagation largely unaffected. This consistency is particularly reassuring given the different assumptions underlying each approach regarding cross-country heterogeneity.

Moreover, the robustness across estimation methods strengthens our conclusions regarding the implications of climate change for macroeconomic stabilization. The finding that rising temperatures may reduce the effectiveness of demand-side fiscal policies while increasing their inflationary consequences appears to be a robust empirical regularity that emerges regardless of the specific panel VAR methodology employed.

5.9. **Local projections as a complementary empirical approach.** As an alternative to the interacted panel BVAR used in the main part, we now reexamine the responses to the global demand and supply shocks by means of the local projection (LP) method of Jordà (2005). The choice of the interacted panel BVAR as our baseline empirical model is derived from the structure of the theoretical framework: the (log-)linearised solution of the DSGE model in Section 2 of the main part takes the form of a (state-dependent) reduced-form VAR, with the structural matrices $\boldsymbol{A}(\hat{\mathcal{T}})$ and $\boldsymbol{H}(\hat{\mathcal{T}})$ entering the impact and lag dynamics in a fully consistent way (see also Sections 3-4 of the main part). The empirical VAR is therefore the natural model-consistent specification to confront with the data, and the joint sign restrictions on $\boldsymbol{A}$ derived from the theoretical model can be imposed directly on the reduced-form innovation covariance matrix. Local projections provide a flexible alternative that is robust to certain forms of dynamic mis-specification (Plagborg-Møller and Wolf, 2021; Olea et al., 2024), but the model-based mapping from theory to estimation is more direct in the VAR. We therefore use the LP approach as a complementary check rather than as the primary empirical specification, and verify that the qualitative pattern of the temperature-dependent transmission of demand and supply shocks is preserved under this alternative.

5.9.1. *Implementation: panel structure, interaction term and shock identification.* The LP exercise re-uses the data and the variable selection of the baseline panel BVAR. The country block $\boldsymbol{y}_{c,t}$ contains the year-on-year percentage change in the GDP deflator, the year-on-year percentage change in real GDP, and the short-term nominal interest rate. The global block $\boldsymbol{g}_t$ contains the year-on-year percentage in the global commodity price index and in global GDP. The interaction variable is the country temperature anomaly $\hat{\mathcal{T}}_{c,t}$ measured from the FAO series, demeaned at the country level so that variation in the interaction is purely within-country. We further include the same "strange-observation" dummy as in the baseline specification to absorb the period of unusual global activity around the Covid-19 pandemic.

Three features of the implementation are worth highlighting because they are non-standard in the typical LP application:

(i) Identification of the global shocks. Local projections require an exogenous shock series on the right-hand side. We obtain it from a small auxiliary VAR estimated on the global block alone. Sign restrictions are imposed on the impact matrix in the same spirit as in equation (8): the demand shock raises both global commodity prices and global GDP on impact, while the (expansionary) supply shock raises global GDP and lowers global commodity prices. Following Rubio-Ramírez et al. (2010), we draw orthonormal $\boldsymbol{Q}$ matrices via the QR decomposition of independent standard normal matrices, retaining 2,000 sign-admissible rotations. To obtain a single time series of structural innovations to be used as regressors in the LP step, we collapse the set of admissible impact matrices to its entry-wise median (a "median-target" rotation in the spirit of Fry and Pagan, 2011), which yields the structural-shock series $\hat{\boldsymbol{\varepsilon}}_t = (\hat{\varepsilon}_t^d, \hat{\varepsilon}_t^s)'$. The sign of the supply-shock series is flipped so the LP IRFs trace out the expansionary supply shock used in Figure 6 of the main part.

(ii) Panel structure. For every country $c = 1, \ldots, N$, every variable v in $\boldsymbol{y}_{c,t}$, and every horizon $h = 0, \ldots, H$, we estimate

$$
\begin{aligned}
y^{v}_{c,t+h} = {} & \alpha^{v,h}_{c} + \tau^{v,h}_{c}\, t + \beta^{v,h}_{d}\hat{\varepsilon}^{d}_{t} + \gamma^{v,h}_{d}\hat{\varepsilon}^{d}_{t}\widetilde{\mathcal{T}}_{c,t} \\
& + \beta^{v,h}_{s}\hat{\varepsilon}^{s}_{t} + \gamma^{v,h}_{s}\hat{\varepsilon}^{s}_{t}\widetilde{\mathcal{T}}_{c,t} + \delta^{v,h}_{c}\widetilde{\mathcal{T}}_{c,t} \\
& + \sum_{k=1}^{\mathcal{P}} \boldsymbol{\Phi}^{v,h}_{c,k}\boldsymbol{w}_{c,t-k} + \eta^{v,h}_{c} D_t + u^{v,h}_{c,t+h},
\end{aligned}
\tag{11}
$$

where $\widetilde{\mathcal{T}}_{c,t} = \hat{\mathcal{T}}_{c,t} - \bar{\mathcal{T}}_{c}$ is the country-demeaned temperature anomaly, $\boldsymbol{w}_{c,t} = (\boldsymbol{y}'_{c,t}, \boldsymbol{g}'_{t})'$ collects the country and global controls, D_t is the dummy variable (for the Covid-19 period), and $\mathcal{P} = 3$ matches the lag length of the baseline panel BVAR. The constant $\alpha^{v,h}_{c}$ is country-specific and the linear time trend absorbs the secular drift in temperature anomalies and global trade integration that would otherwise risk loading on $\hat{\varepsilon}_t\widetilde{\mathcal{T}}_{c,t}$. We aggregate the country-level coefficients to the panel level using the mean-group estimator of Pesaran and Smith (1995), that is, by averaging $\hat{\beta}^{v,h}_{c}$ and $\hat{\gamma}^{v,h}_{c}$ across c. The pointwise standard error of the mean-group estimator is the cross-country standard deviation divided by $\sqrt{N}$, from which we construct 68 percent pointwise bands. This direct cross-country averaging mirrors the Pesaran–Smith approach used in our baseline interacted panel BVAR and ensures that the LP results and the BVAR results refer to the same population object: the average country response.

(iii) Evaluation of the IRFs at low and high temperature anomalies. Because $\widetilde{\mathcal{T}}_{c,t}$ enters (11) multiplicatively with the structural shock, the impulse response at any temperature value τ for country c and variable v is given by

$$
\mathrm{IRF}_{c,v,h}(\tau) = \hat{\beta}^{v,h}_{c} + \hat{\gamma}^{v,h}_{c}(\tau - \bar{\mathcal{T}}_{c}).
$$

We follow the IPVAR convention and report IRFs at a low and a high value of the temperature anomaly. Because the IPVAR figures use a regular grid spanning the empirical support of $\hat{\mathcal{T}}_{c,t}$, while the LP estimates of $\hat{\gamma}$ are particularly noisy in the thin tails of the cross-country IAV distribution, we evaluate the LP IRFs at the 10th and 90th percentile of the pooled IAV (rather than at the extreme grid endpoints used in Figures 4 and 6 of the main part). This choice trades a little extrapolation distance for substantially tighter bands without altering the qualitative pattern. Finally, IRFs are normalised so that the impact effect of each shock on global GDP equals one percentage point, exactly as in the IPVAR figures of the main part.

(iv) Width of the confidence bands. A practical observation about Figures 18 and 19 below is that the 68 percent pointwise bands look tighter than what is typically reported in single-country LP applications. The reason lies in how the bands are constructed. They are the standard error of the mean-group estimator in the sense of Pesaran and Smith (1995): at every horizon h, we take the cross-country standard deviation of the country-level LP coefficients $\hat{\beta}_c^{v,h}$ (and analogously for $\hat{\gamma}_c^{v,h}$) and divide it by $\sqrt{N}$, with $N = 20$. This standard error captures the uncertainty about the *average* country response under cross-country coefficient heterogeneity. It does *not*, however, include the within-country sampling uncertainty of the country-level LP regressions, which in single-country LPs is typically the dominant source of width and is normally reported through a Newey–West / HAC correction that accounts for the moving-average structure of the LP residuals at horizon h. We deliberately follow the panel convention rather than the single-country one because we want the LP figures to be directly comparable to the IPVAR figures of the main part: in the IPVAR, bands are also constructed on the mean-group object (averaging country-level posterior IRFs), so reporting cross-country mean-group bands here keeps the units of uncertainty aligned across the two methods. Augmenting the LP standard errors with within-country HAC uncertainty would mechanically widen the bands, but it would not affect the qualitative conclusions drawn from the figures.

5.9.2. *Empirical results.* Figures 18 and 19 report the LP-based impulse responses for the demand and the (expansionary) supply shock, in the same layout used for the IPVAR results in the main part. Inspection of the figures shows that the LP estimates reproduce the qualitative content of our baseline findings.

For the global demand shock (Figure 18), the country-level GDP deflator, real GDP, and the short-term interest rate all rise on impact, with hump-shaped dynamics that closely resemble the corresponding IRFs from the IPVAR baseline (Figure 4 of the main part). More importantly, the responses

Figure 18. Local projections: global demand shock

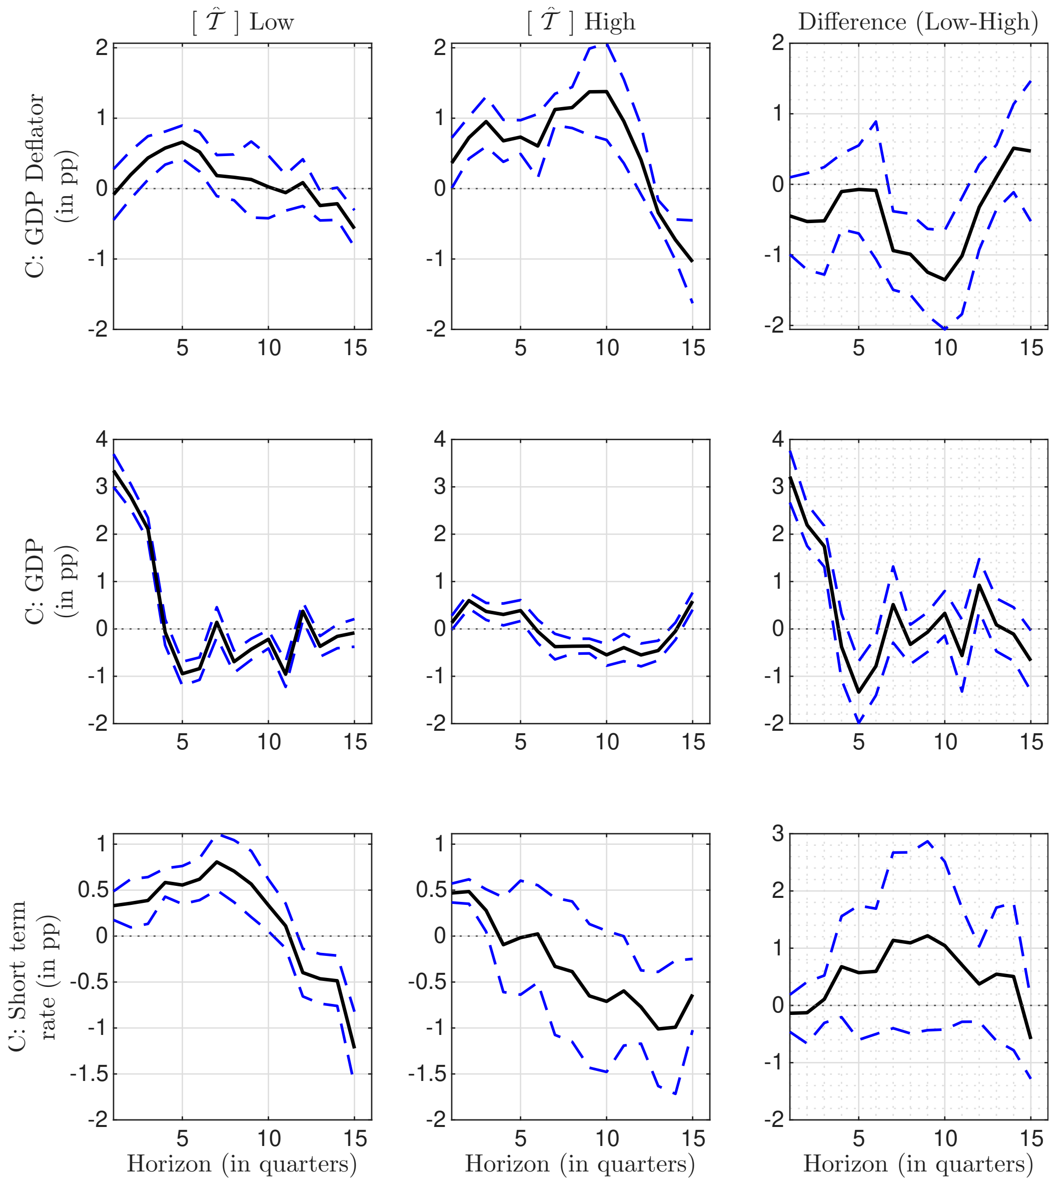

Note: Mean-group LP impulse responses to a global demand shock identified from sign restrictions on a small auxiliary VAR for global commodity prices and global GDP. The country block contains the GDP deflator, GDP, and the short-term nominal interest rate. IRFs are evaluated at low (10th percentile) and high (90th percentile) values of the country-demeaned Berkeley temperature anomaly and normalised to a one-percentage-point impact effect on global GDP. Solid lines are the mean-group point estimates; dashed lines are 68 percent pointwise bands.

depend on the level of the temperature anomaly. Inflation is statistically distinguishable across the two temperature states, with the corresponding "Difference (Low–High)" panel showing a sign reversal that is consistent with the BVAR finding of an amplified inflationary response when temperatures are

FIGURE 19. Local projections: global supply shock (expansionary)

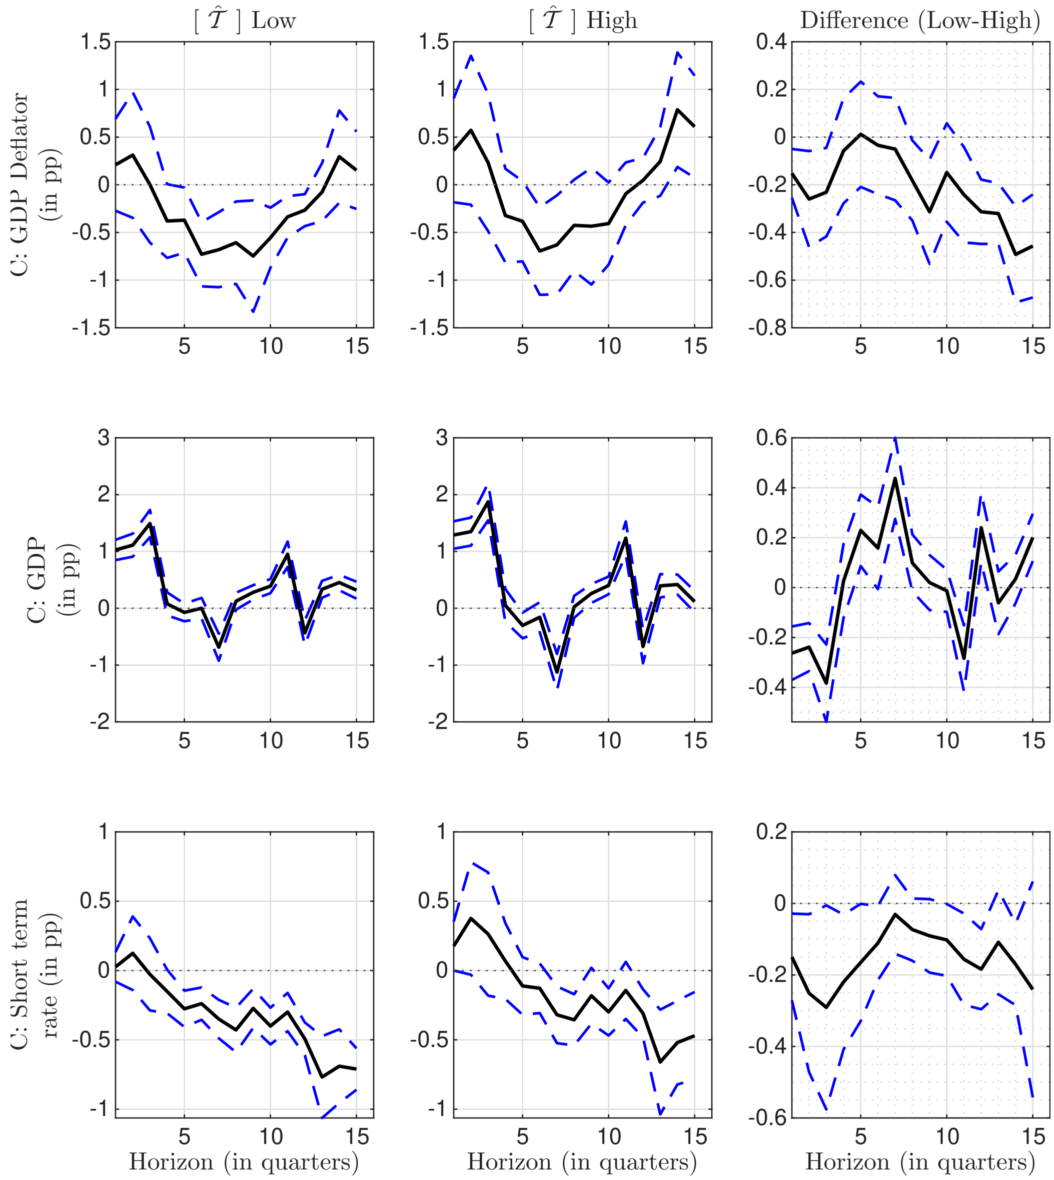

Note: Mean-group LP impulse responses to an expansionary global supply shock identified from sign restrictions on a small auxiliary VAR for global commodity prices and global GDP. Sign convention is set so that an expansionary supply shock lowers global commodity prices and raises global GDP. The country block contains the GDP deflator, GDP, and the short-term nominal interest rate. IRFs are evaluated at low (10th percentile) and high (90th percentile) values of the country-demeaned Berkeley temperature anomaly and normalised to a one-percentage-point impact effect on global GDP. Solid lines are the mean-group point estimates; dashed lines are 68 percent pointwise bands.

elevated. The output response, in turn, is comparatively muted in the high-temperature state, again mirroring the IPVAR pattern of a dampened real activity response when the climate-related interaction term is large.

For the global supply shock (Figure 19), the responses are once again qualitatively in line with the IPVAR baseline (Figure 6 of the main part). An expansionary supply shock lowers domestic prices and raises domestic output, while the short-term interest rate tends to decline. The temperature dependence of the supply-shock IRFs is much weaker than that of the demand-shock IRFs: the "Difference (Low–High)" panels are noisier and the temperature-induced shifts are not robustly estimated. This mirrors the central message of the main part of the paper, namely that climate change leaves the propagation of supply shocks essentially unchanged while it materially alters the propagation of demand shocks.

Taken together, the LP results provide an additional, methodologically distinct piece of evidence in support of our baseline conclusion: temperature anomalies modify the transmission of global demand shocks in line with the predictions of the theoretical model, while the transmission of global supply shocks remains largely insensitive to the temperature anomaly. The fact that this pattern emerges from a panel local projection specification — which differs from the interacted panel BVAR in its identification timing, its lack of cross-equation restrictions, a certain lack of smoothing the IRFs, and its single-equation orientation — reinforces our reading that the temperature-dependent steepening of the Phillips curve relationship documented in the main part is not an artefact of the VAR-based methodology.

5.9.3. *Pooled-panel local projections.* The mean-group LP exercise above is the LP analog of the baseline (mean-group) IPVAR estimated in the main part. As a parallel to the pooled Bayesian panel VAR robustness check reported in Section "Alternative panel methodology" above, we now also implement a *pooled-panel* version of the local projection. Imposing coefficient homogeneity across countries trades the heterogeneity-robust mean-group estimator of Pesaran and Smith (1995) for the larger effective sample size of the stacked panel. The two panel-LP estimators thus span the same methodological dichotomy — between heterogeneity and pooled efficiency — as the two panel-BVAR estimators considered earlier, and provide a fully symmetric LP–BVAR robustness picture.

The pooled LP specification re-uses the data, the global-shock identification, and the construction of the country-demeaned interaction term $\widetilde{\mathcal{T}}_{c,t}$ described above. The single difference is in the panel structure: rather than estimating equation (11) country-by-country and aggregating via the mean-group estimator, we stack all country observations and run, for every variable $v \in \{\text{GDP deflator, GDP, short-term rate}\}$ and every horizon h, a single pooled OLS regression

$$\begin{aligned} y^{v}_{c,t+h} = {} & \alpha^{v,h}_{c} + \tau^{v,h}_{c}\, t + \beta^{v,h}_{d}\hat{\varepsilon}^{d}_{t} + \gamma^{v,h}_{d}\hat{\varepsilon}^{d}_{t}\,\widetilde{\mathcal{T}}_{c,t} \\ & + \beta^{v,h}_{s}\hat{\varepsilon}^{s}_{t} + \gamma^{v,h}_{s}\hat{\varepsilon}^{s}_{t}\,\widetilde{\mathcal{T}}_{c,t} + \delta^{v,h}\widetilde{\mathcal{T}}_{c,t} \\ & + \sum_{k=1}^{\mathcal{P}} \mathbf{\Phi}^{v,h}_{k}\,\boldsymbol{w}_{c,t-k} + \eta^{v,h} D_{t} + u^{v,h}_{c,t+h}, \end{aligned} \tag{12}$$

in which the slope coefficients $\beta^{v,h}_{d}$, $\gamma^{v,h}_{d}$, $\beta^{v,h}_{s}$, $\gamma^{v,h}_{s}$, $\delta^{v,h}$, $\mathbf{\Phi}^{v,h}_{k}$, $\eta^{v,h}$ are common to all countries, while the unobserved cross-country level differences are absorbed by country fixed effects $\alpha^{v,h}_{c}$ and country-specific linear time trends $\tau^{v,h}_{c}$. The coefficients on the interaction terms $\hat{\varepsilon}^{d}_{t}\widetilde{\mathcal{T}}_{c,t}$ and $\hat{\varepsilon}^{s}_{t}\widetilde{\mathcal{T}}_{c,t}$ are identified by exploiting within-country variation in the interaction term together with the cross-country common variation in the respective structural shock. The IRFs at low and high temperature anomaly are again evaluated at the 10th and 90th percentile of the pooled IAV and normalised to a one-percentage-point impact effect on global GDP, exactly as in the mean-group LP.

Standard errors of the pooled estimator are obtained as one-way cluster-robust covariances with countries as clusters (Cameron-Gelbach-Miller correction, see, e.g., the panel-LP discussion in Plagborg-Møller and Wolf, 2021). Variance of the IRF at temperature value τ, $\text{Var}(\hat{\beta}^{v,h}_{d} + \hat{\gamma}^{v,h}_{d}(\tau - \bar{\mathcal{T}}))$, is computed using the joint cluster-robust covariance of $(\hat{\beta}^{v,h}_{d}, \hat{\gamma}^{v,h}_{d})$, and analogously for the supply shock. As in the pooled BVAR, the larger effective sample size of the stacked panel produces tighter pointwise bands than those obtained for the mean-group estimator.

Figures 20 and 21 report the pooled LP IRFs for the global demand and the (expansionary) supply shock. The qualitative content of the figures is the same as that of the mean-group LP and of the IPVAR baseline.

For the global demand shock (Figure 20), the response of the country-level GDP deflator builds up over the projection horizon, with a clearly steeper hump in the high temperature state and a more contained response in the low state. The corresponding “Difference (Low–High)” panel is statistically significant for inflation over a wide range of horizons, in line with the IPVAR

FIGURE 20. Pooled-panel local projections: global demand shock

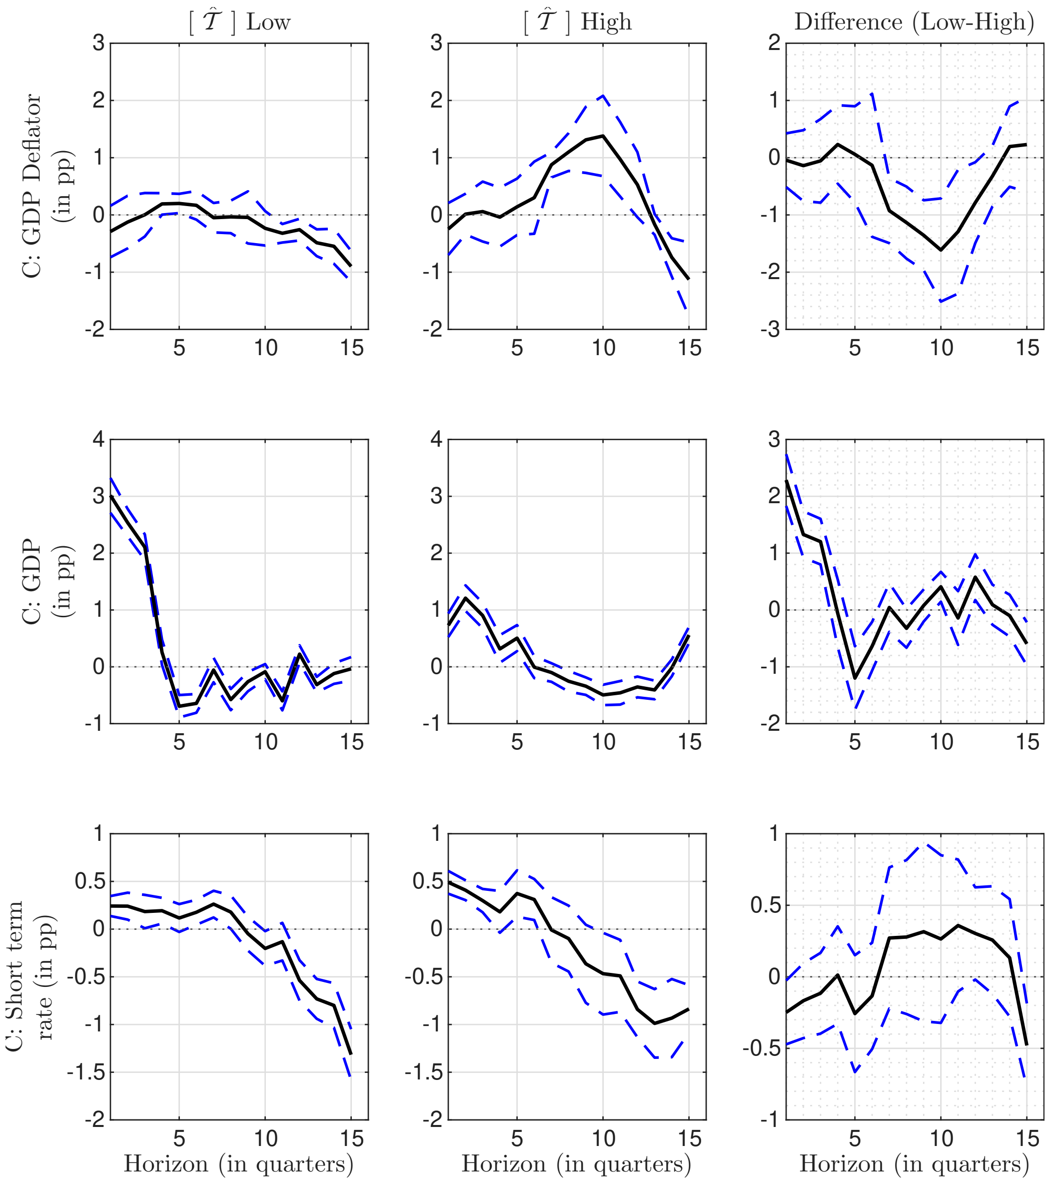

Note: Pooled-panel LP impulse responses to a global demand shock identified from sign restrictions on a small auxiliary VAR for global commodity prices and global GDP. The pooled LP imposes coefficient homogeneity across countries; country fixed effects and country-specific linear time trends absorb cross-country level differences and trends. The country block contains the GDP deflator, GDP, and the short-term nominal interest rate. IRFs are evaluated at low (10th percentile) and high (90th percentile) values of the country-demeaned Berkeley temperature anomaly and normalised to a one-percentage-point impact effect on global GDP. Solid lines are the point estimates; dashed lines are 68 percent pointwise bands based on one-way cluster-robust standard errors (clusters = countries).

finding that elevated temperature anomalies amplify the inflationary impact of demand shocks. The real GDP response also exhibits temperature dependence: the impact effect on output is markedly stronger in the low-temperature state than in the high-temperature state, while the medium-horizon difference

Figure 21. Pooled-panel local projections: global supply shock (expansionary)

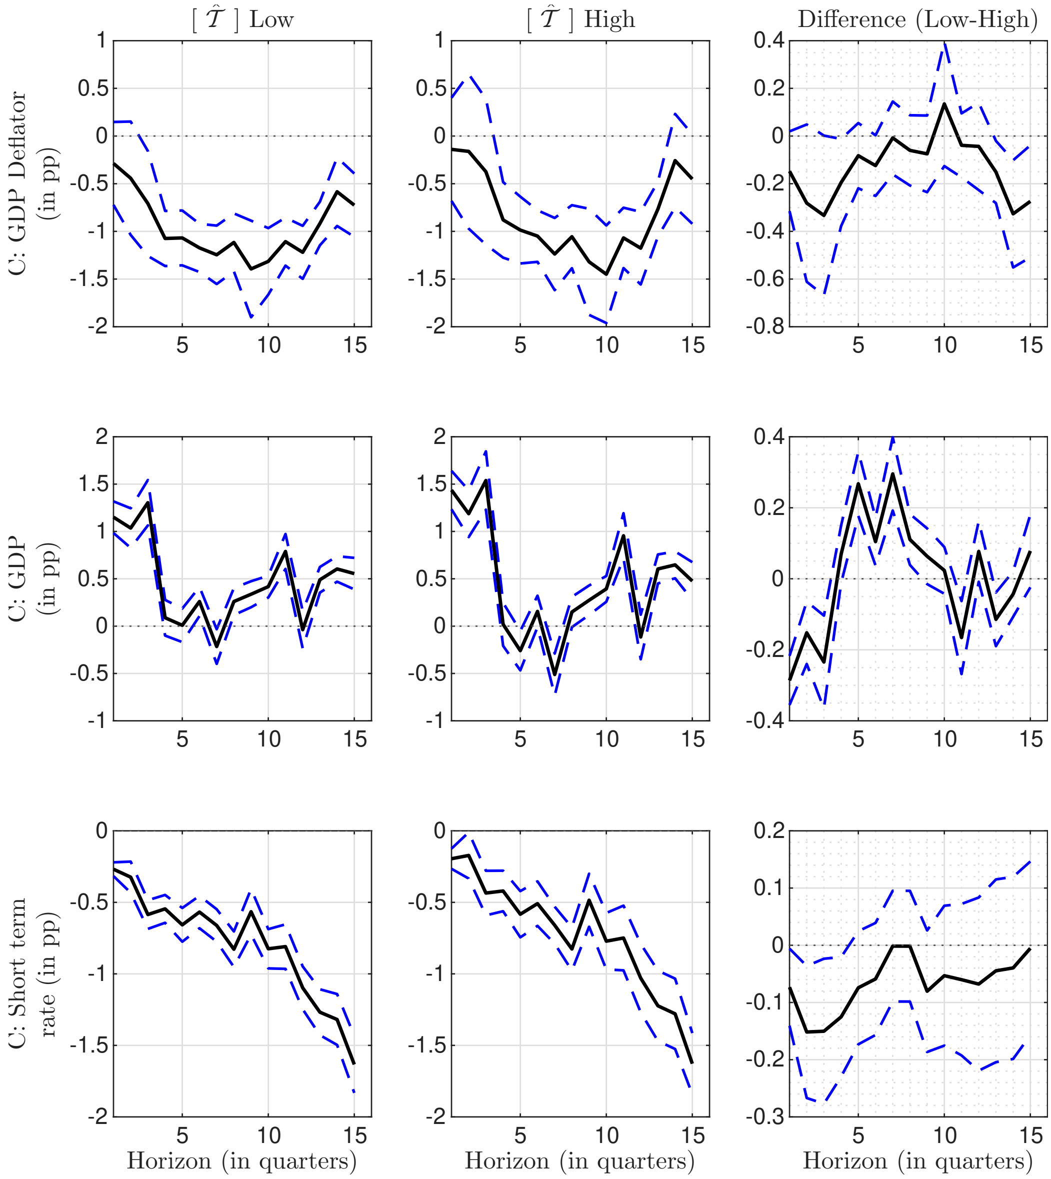

Note: Pooled-panel LP impulse responses to an expansionary global supply shock identified from sign restrictions on a small auxiliary VAR for global commodity prices and global GDP. Sign convention is set so that an expansionary supply shock lowers global commodity prices and raises global GDP. The pooled LP imposes coefficient homogeneity across countries; country fixed effects and country-specific linear time trends absorb cross-country level differences and trends. IRFs are evaluated at low (10th percentile) and high (90th percentile) values of the country-demeaned Berkeley temperature anomaly and normalised to a one-percentage-point impact effect on global GDP. Solid lines are the point estimates; dashed lines are 68 percent pointwise bands based on one-way cluster-robust standard errors (clusters = countries).

becomes statistically less pronounced — a pattern that is qualitatively consistent with the dampening of the cumulative real activity response when the climate-related interaction term is large. The interest-rate response, finally, mirrors the inflation response: monetary policy reacts more strongly when temperatures are high though the response is statistically insignificant.

For the expansionary global supply shock (Figure 21), the responses are again in line with the IPVAR baseline of the main part: domestic prices fall, domestic output rises, and the short-term interest rate declines. The responses to the supply shock display modest temperature dependence in the price block and only weak temperature dependence in the output and interest-rate blocks. As in the mean-group LP, the temperature-induced shifts are markedly smaller than those documented for the demand shock, in line with our central reading that the temperature anomaly primarily reshapes the slope of the Phillips curve and therefore matters for demand-driven adjustments while leaving supply-driven adjustments largely intact.

The narrowing of the confidence bands relative to the mean-group LP is the natural consequence of pooling: the pooled estimator exploits the full $N \times T$ panel rather than averaging N country-level estimates. As discussed in the analogous BVAR section above, this efficiency gain comes at the cost of imposing slope homogeneity — a constraint that the cross-country LP coefficients in the mean-group exercise above suggest is approximately, but not exactly, satisfied. The fact that the pooled LP, the mean-group LP, the pooled BVAR, and the mean-group IPVAR yield qualitatively the same picture is therefore reassuring: the temperature-dependent transmission of demand shocks documented in the main part survives both the BVAR-vs-LP and the mean-group-vs-pooled robustness exercises.

5.10. **Reduced cross-sectional dimension.** The baseline results presented earlier rely on a country set that includes countries with significant heterogeneity across various dimensions. This heterogeneity spans not only the economic dimension—distinguishing between industrial economies and emerging market economies—but also the geographical dimension, which leads to climatic differences among the countries. For example, we consider Northern countries such as Finland and Sweden, as well as tropical/sub-tropical countries such as India and Mexico.

It is natural to assume that climate change affects countries differently, not only due to differences in their level of economic development but also because of their geographic location (Nath et al., 2024). Given these differences, we

FIGURE 22. Demand shock: selected country-set

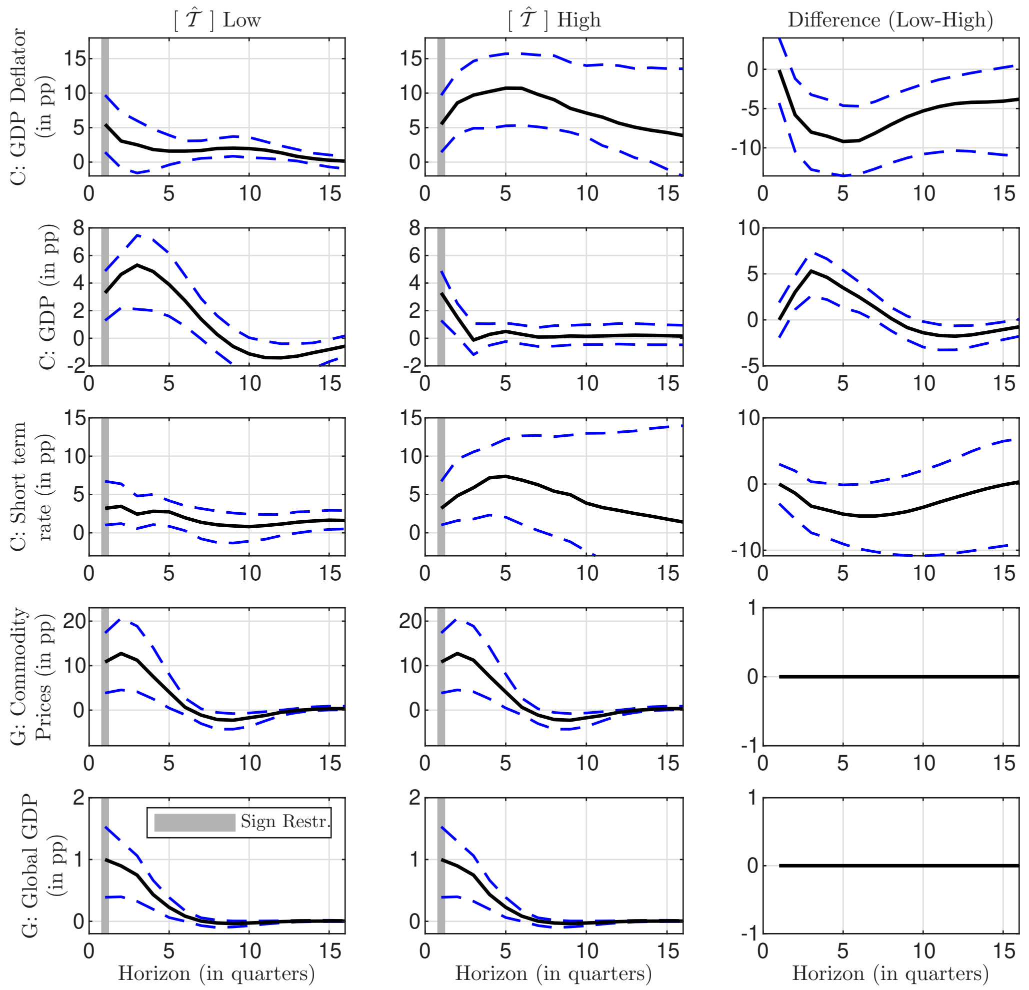

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global demand shock with respect to the temperature anomaly. The figure is based on a selected country set excluding Finland, the Netherlands, Japan, Sweden and the UK.

now examine the stability of our baseline results with respect to a selected sub-set of countries.

As previously highlighted, the likelihood-ratio test (presented in the main part) predominantly favors the IVAR specification for most countries. However, there are a few countries for which the VARX specification constitutes a valid alternative, according to the likelihood-ratio test. These countries include Finland, the Netherlands, Japan, Sweden, and the UK. To test the robustness of our findings, we re-estimate the IPVAR model while excluding these five countries. We then carry out the same analysis as before, examining

FIGURE 23. Supply shock: selected country-set

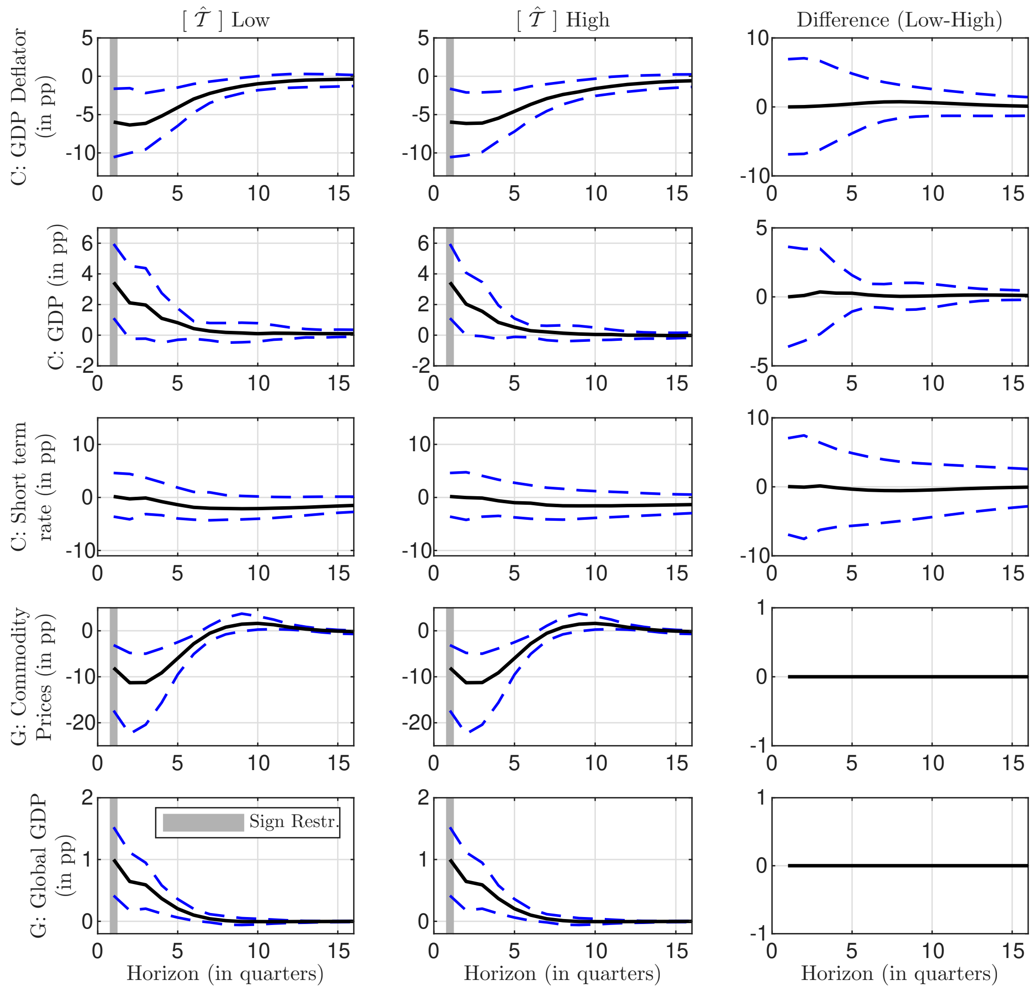

Note: The figure shows the sensitivity of the impulse response functions of inflation (GDP deflator, yoy change in percent), output (GDP, yoy change in percent) and the short term interest rate to a global supply shock with respect to the temperature anomaly. The figure is based on a selected country set excluding Finland, the Netherlands, Japan, Sweden and the UK.

the sensitivity of the impulse response functions (IRFs) at the country level to a demand shock with respect to the temperature anomaly.

The results of this robustness check are presented in Figure 22 fore the global demand shock and in Figure 23 for the global supply shock. Our focus in this analysis is on the IRFs of output and inflation, specifically their sensitivity to the temperature anomaly. As shown in the figure, the (expansionary) demand shock still leads to increases in both output and inflation. Importantly, the differences in the responses of these variables to the temperature anomaly remain pronounced. However, a key observation is that the difference in the IRFs between low and high values of the temperature anomaly is now noticeably larger compared to the baseline results.

For instance, the maximum difference in inflation between low and high temperature anomalies is around four percentage points when using the reduced country coverage, while the difference is only about two percentage points when the full country set is used. A similar pattern is observed for output: with the reduced country coverage, the maximum difference in output responses is around six percentage points, compared to approximately four percentage points when using the full set of countries.

These findings underscore that the choice of countries in the sample does matter to some extent, a point that has been emphasized in the literature by Byrne and Vitenu-Sackey (2024). Nonetheless, we consider this robustness check to provide further confirmation of the validity of our baseline results.

5.11. **Additional interaction variables.** The baseline results rely on a specification of the IPVAR model which uses only one interaction variable: the temperature anomaly $\mathcal{T}_t$. The key equation of the baseline model which captures the non-linearity that arises from the interaction variable reads

$$\text{IVAR:} \quad \boldsymbol{B}_{11}^{c}(\mathcal{T}_t)\boldsymbol{y}_{c,t-1} = \left(\boldsymbol{\Gamma}^{c}(\bar{\mathcal{T}}) + \boldsymbol{\mathcal{F}}^{c}\hat{\mathcal{T}}_t\right)\boldsymbol{y}_{c,t-1} \tag{13}$$

In the theoretical framework, the transmission of the temperature anomaly operates through total factor productivity (TFP): higher temperature anomalies depress TFP and thereby propagate to macroeconomic outcomes. In in line with this, the current empirical specification employs a single interaction variable. This, however, raises the possibility that effects attributable to omitted interaction channels are inadvertently loaded onto the temperature anomaly. This risk arises precisely because the exclusive use of one interaction term, while methodologically parsimonious and conducive to a compact, interpretable model, necessarily abstracts from competing or complementary sources of heterogeneity. The benefit of parsimony thus comes at the cost of potential omitted-variable bias in the interaction structure, or a mis-attribution of effects from other relevant interaction variables to the temperature anomaly. Motivated by these considerations, we now broaden the analysis to incorporate additional interaction elements and assess the robustness of the temperature-anomaly mechanism accordingly. To this purpose, we extend equation (13) as follows

$$\text{(14)} \qquad \text{IVAR:} \quad \boldsymbol{B}_{11}^{c}(\mathcal{T}_t, \zeta_t)\boldsymbol{y}_{c,t-1} = \left(\boldsymbol{\Gamma}^{c}(\bar{\mathcal{T}}, \bar{\zeta}) + \boldsymbol{\mathcal{F}}^{c}\hat{\mathcal{T}}_t + \boldsymbol{\mathcal{Z}}^{c}\hat{\zeta}_t\right)\boldsymbol{y}_{c,t-1}$$

where ζ_t is an additional interaction variable (subsequently explained in more detail), and $\hat{\zeta}_t$ is its relative deviation from its steady state $\bar{\zeta}$. We will subsequently consider empirical variables for ζ_t that characterize structural elements related to TFP at the country level ($\zeta_t \to \zeta_t^c$). In this context, Goldin et al. (2024) identify several key contributors to a productivity slowdown, including mis-measurement, a decline in the contribution of capital per worker, lower spillovers from the growth of intangible capital, and a reduction in allocative efficiency. Sectoral reallocation, a lower contribution of human capital, and a slowdown in innovation also play a role in certain countries. However, also variables beyond a specific TFP relation are possible in shaping the inflationary effect in response to a demand (and supply) shock. Importantly, throughout all subsequent analyses, the specification includes only one additional interaction alongside the temperature anomaly, thereby restricting the model to exactly two interaction variables to preserve parsimony. This constraint is deliberately imposed to balance methodological clarity with the computational limitations of our available hardware.

5.11.1. *Elements shaping TFP beyond the temperature anomaly.* We examine several variables that could influence TFP in a similar way as the temperature anomaly, focusing particularly on those related to innovation and human capital. We consider data from the OECD on patents, measured by applications filed under the Patent Cooperation Treaty (PCT), as an indicator for innovation. Additionally, we include spending on research and development, specifically basic research expenditure as a percentage of GDP, which also serves as an innovation indicator. As indicators for changes in human capital, we use the dependency ratio (the number of pensioners relative to the working-age population, 15-64) and a measure of educational attainment, specifically the percentage of the population with upper secondary or post-secondary non-tertiary education. Using OECD data implies that India and Mexico drop out of the sample.

FIGURE 24. Inflation/price response to global demand shock: additional interaction variables

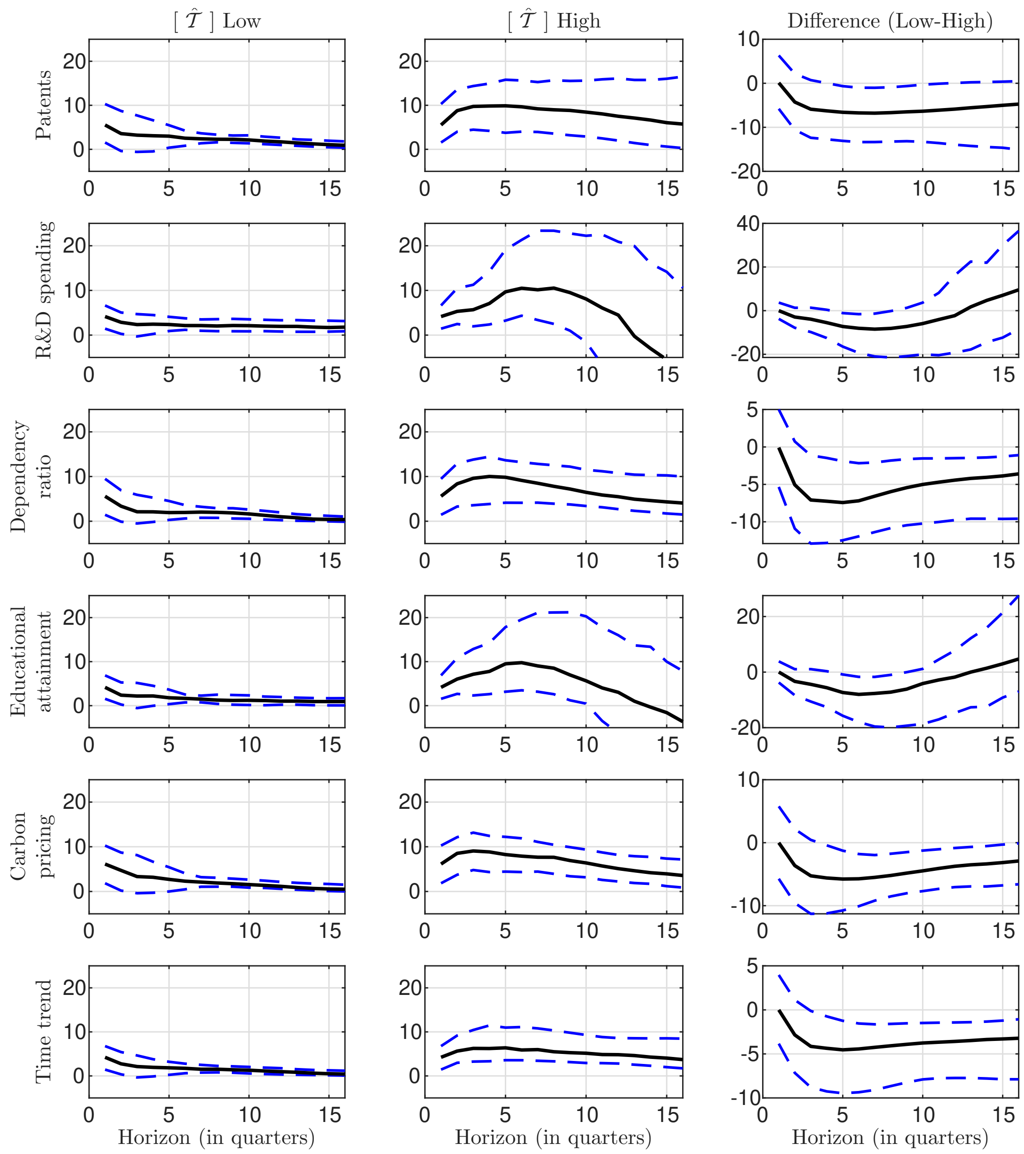

Note: The figure shows the impulse response functions of inflation (GDP deflator) to a global demand shock. The IRFs are shown for low and high values of the temperature anomaly. Each row shows the IRFs of the extended baseline model, where the extension refers to the variable mentioned on the left of each row being used as additional interaction variable in the IPVAR model.

We incorporate these potentially confounding variables into the model by sequentially adding each of the four variables as an additional interaction variable. To examine their effects, we re-estimate the model, apply the identification for global demand and supply shocks and analyze the sensitivity of the model to the temperature anomaly. If any of these additional variables is

FIGURE 25. GDP response to global demand shock: additional interaction variables

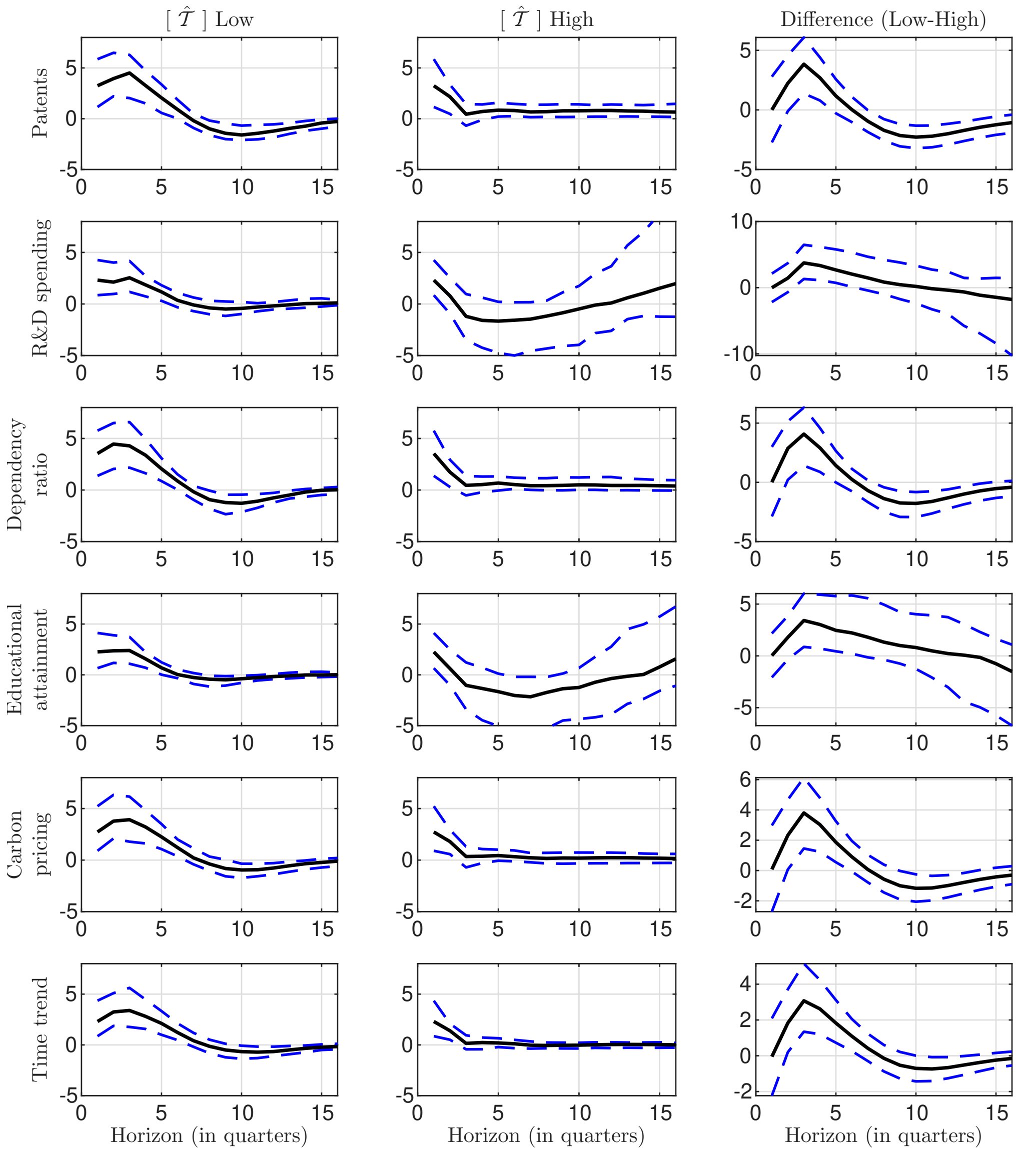

Note: The figure shows the impulse response functions of GDP to a global demand shock. The IRFs are shown for low and high values of the temperature anomaly. Each row shows the IRFs of the extended baseline model, where the extension refers to the variable mentioned on the left of each row being used as additional interaction variable in the IPVAR model.

relevant, we expect changes in the sensitivity of the IRFs of GDP, the GDP deflator, and the interest rate at the country level.

The results are displayed in Figures 24 and 25 for the demand shock, and in Figures 26 and 27 for the supply shock. For illustration, Figure 24 reports the IRFs of the GDP deflator to a global demand shock. As in the baseline, IRFs are shown at different levels of the temperature anomaly (the primary

Figure 26. Inflation/price response to global supply shock: additional interaction variables

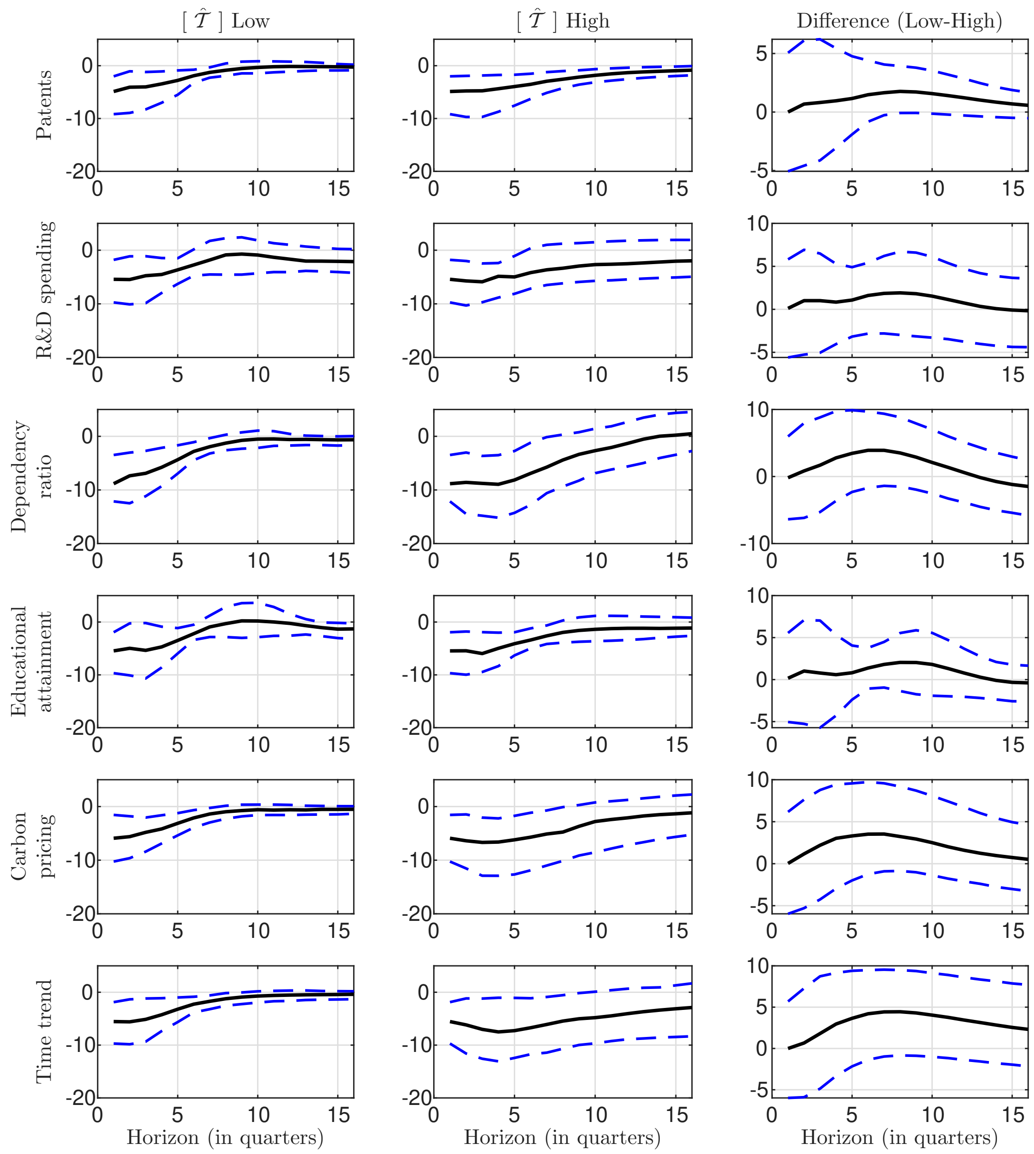

Note: The figure shows the impulse response functions of inflation (GDP deflator) to a global supply shock. The IRFs are shown for low and high values of the temperature anomaly. Each row shows the IRFs of the extended baseline model, where the extension refers to the variable mentioned on the left of each row being used as additional interaction variable in the IPVAR model.

interaction variable). Each row, however, introduces a different *additional* interaction variable, as labeled on the left margin. For example, the first row replicates the baseline specification and augments it with patents as the additional interaction term. The layout of Figures 25–27 mirrors that of Figure 24, but applies to GDP and to the supply shock, respectively. In all cases, we compute time-and-country means for each additional interaction variable and

FIGURE 27. GDP response to global supply shock: additional interaction variables

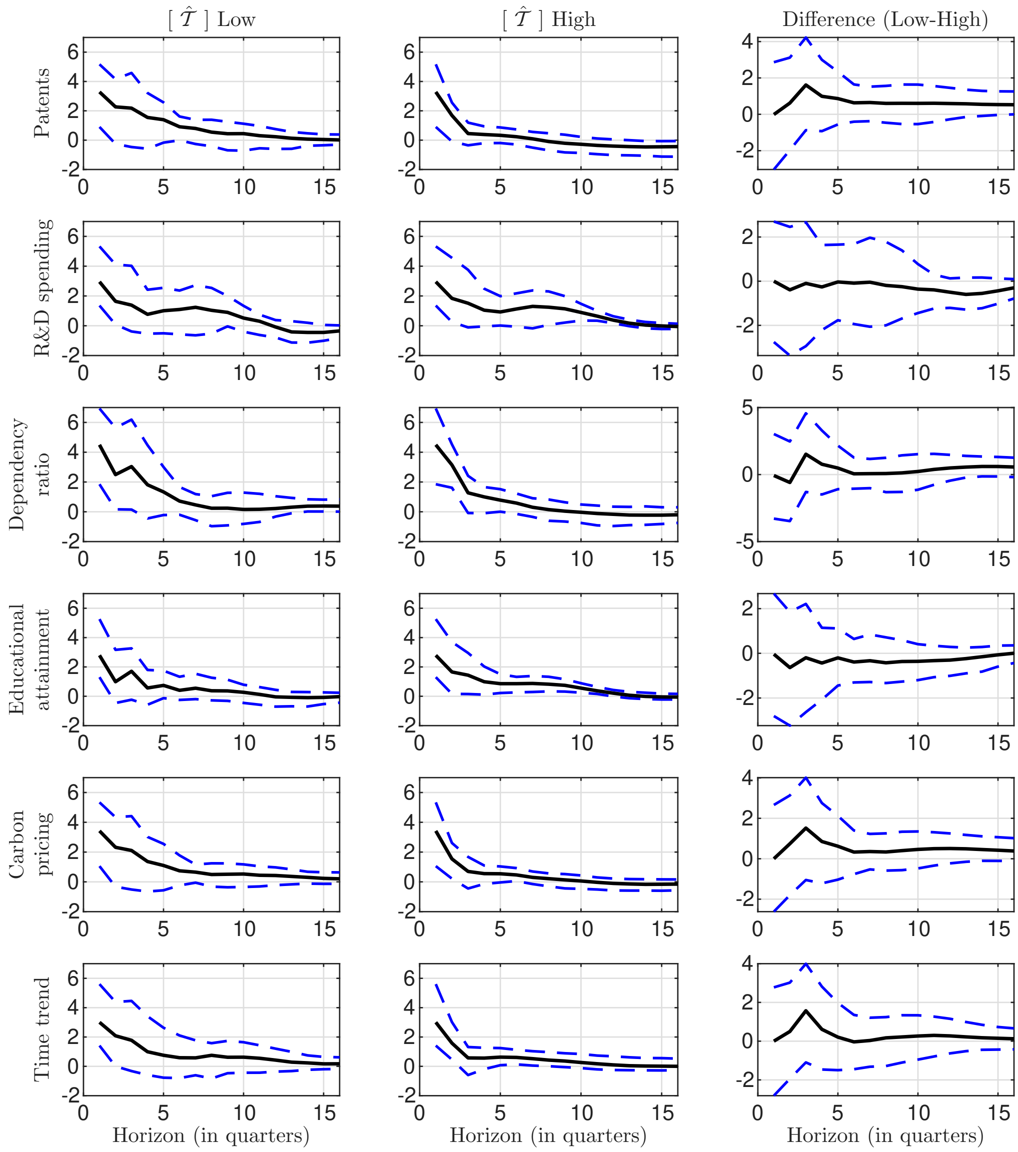

Note: The figure shows the impulse response functions of GDP to a global supply shock. The IRFs are shown for low and high values of the temperature anomaly. Each row shows the IRFs of the extended baseline model, where the extension refers to the variable mentioned on the left of each row being used as additional interaction variable in the IPVAR model.

present IRFs of the GDP deflator and GDP to global demand and supply shocks, varying only the temperature anomaly, to facilitate direct comparison with the baseline IRFs (the nominal interest rate is omitted as it lies outside our primary focus).

To fix ideas, consider the GDP response to the global demand shock in Figure 25. When patents are included as the additional interaction variable, the output IRFs are qualitatively indistinguishable from those in the baseline. This holds both for the IRFs themselves and for the cross-scenario differences across low versus high temperature anomaly states. Hence, adding patents does not alter the baseline conclusions and further supports the robustness of the results with respect to this dimension. An analogous pattern emerges for the other additional interaction variables—R&D spending, the dependency ratio, and educational attainment. In each case, output IRFs remain qualitatively unchanged, and the differences between low and high temperature-anomaly states retain their statistical significance. We therefore conclude that, for GDP responses to demand shocks, the inclusion of these additional interaction variables does not qualitatively affect our baseline findings.

Turning to the GDP deflator under a demand shock, the baseline conclusions are likewise preserved. Across all additional interaction variables, prices increase significantly, and—crucially—the magnitude of the price response rises with higher levels of the temperature anomaly. Moreover, the difference between low and high temperature-anomaly states remains statistically different from zero in every specification considered.

Finally, the supply-shock results confirm our baseline interpretation. While the additional interaction variables leave intact the pattern of an output increase and a price decline in response to the expansionary supply shock, the differences in IRFs across low versus high temperature-anomaly states remain statistically insignificant. Taken together, these findings reinforce the (qualitative) robustness of our baseline results to the inclusion of multiple additional interaction variables.

5.11.2. *Elements shaping inflation.* Bettarelli et al. (2025) investigate the impact of climate change policies on inflation across a large set of countries and find that carbon taxes lead to substantial inflationary pressures. The effect is not negligible: a one standard deviation carbon tax shock—corresponding to a \$5/tCO2 increase in emissions-weighted carbon taxes—results in an increase in the price level of about 0.7 percent one year after the implementation of the policy, and between 1.6 and 4 percent in the medium term. Moreover, they find that emissions trading systems, as well as non-market-based climate change policies (such as R&D subsidies), also have an (though smaller) effect on prices.

Since the introduction of these policies aligns with the steady increase in temperatures, it is possible that our finding—namely, higher inflationary pressure in the wake of a demand shock when temperatures are high—may result from these aforementioned policy measures rather than from climate change itself. To explore this possibility, we extend our empirical model by introducing an additional interaction variable: a measure of the carbon pricing system at the country level, alongside the temperature anomaly variable ($\hat{\mathcal{T}}_t$).

The inclusion of a carbon pricing variable reduces the country coverage of our sample, as not all countries included in our baseline results have implemented a carbon pricing system. Specifically, India and the USA drop out of the sample.[4] For the remaining countries, we introduce a variable capturing carbon pricing. For the European countries, we use the price of emissions allowances (EUA) traded on the European union's Emissions Trading Scheme (ETS), which was introduced in 2005. For Australia, we use the Australian Carbon Credit Unit (ACCU), introduced in 2012. For Japan, we use the Saitama Target Setting Emissions Trading System, introduced in 2011. For New Zealand, we use the New Zealand Emissions Trading Scheme (NZ ETS), which was launched in 2008 and covers roughly half of the country's emissions. For Canada, we use the Canadian National Output-Based Pricing System (OBPS), introduced in 2019.

A challenge arises due to the relatively late start of these time series (2005 for European countries, and even later for others), while our baseline sample begins in 1961. To avoid limiting the sample size for the estimation of the IPVAR model, we extend the carbon price data by filling in the missing entries with zeros. This reflects the fact that, prior to the introduction of carbon pricing systems, no price for carbon emissions was applied, aligning with a zero price for this period.

We re-estimate the IPVAR with the extended carbon pricing series and find that the estimated coefficients on the carbon pricing variable are quantitatively negligible and, crucially, statistically insignificant. To assess the implications

[4] In the USA, a carbon pricing scheme is introduced in various states; however, it is not implemented at the national level. While the USA could be included in the sample with carbon pricing data at the state level (consider the following link in this respect), the fact that on the one hand, only a few states have introduced carbon pricing and on the other hand, it is difficult to aggregate state-level pricing data without quantity information (of carbon emission at the state level), led us to exclude the USA from the sample for this exercise.

of including carbon pricing in the empirical specification more transparently, we also report IRFs of the GDP deflator and GDP to both demand and supply shocks (fifth-row panels in Figures 24–27). The resulting IRFs closely track the baseline responses, exhibiting no qualitative differences and only minor quantitative deviations. These findings indicate that our baseline results are robust to the inclusion of carbon pricing, confirming that this policy variable does not materially alter the main conclusions.

5.11.3. *A linear time trend.* As a final exercise, we include a deterministic time trend as an additional interaction variable. This choice is motivated by the figures in the Introduction, which document a pronounced upward trajectory in temperature anomalies across all countries since 1977. Even with the time trend included, the baseline results are preserved: the coefficient estimates for $\mathcal{Z}^c$ remain quantitatively small and statistically indistinguishable from zero. The implications are corroborated by the IRFs reported in the bottom-row panels of Figures 24–27, which show that the time trend does not materially affect the sensitivity of price and output responses to global demand and supply shocks across different levels of the temperature anomaly.

6. Climate change as a business cycle shock

In the following, we present the empirical results for a specification along the lines of equation (2) of the main part. This specification treats climate change, again proxied by the temperature anomaly, as an endogenous variable in the system, and it is subsequently used to introduce an exogenous shock to the temperature anomaly into the system, similar to other "traditional" shocks, such as those resulting from economic policy changes, preference shifts, cost-push shocks, and the like. This specification has been considered in Pretis (2021); Mukherjee and Ouattara (2021); Ciccarelli and Marotta (2024); Huber et al. (2023); Ciccarelli et al. (2024), among others. This analysis serves to examine the implicit assumption used in the idea formalized by the equation (3) of the main part, that is, that climate change is exogenous to the system. Obviously, climate change is endogenous to the economic system if a sufficiently long time horizon is considered. However, for the purpose of our analysis, we focus on business cycles that capture short -(1 to 3 years) to medium-term (3 to 8 years) fluctuations. The plausibility and reliability of our results therefore

FIGURE 28. Shock to global temperature anomaly ($\hat{\mathcal{T}}$)

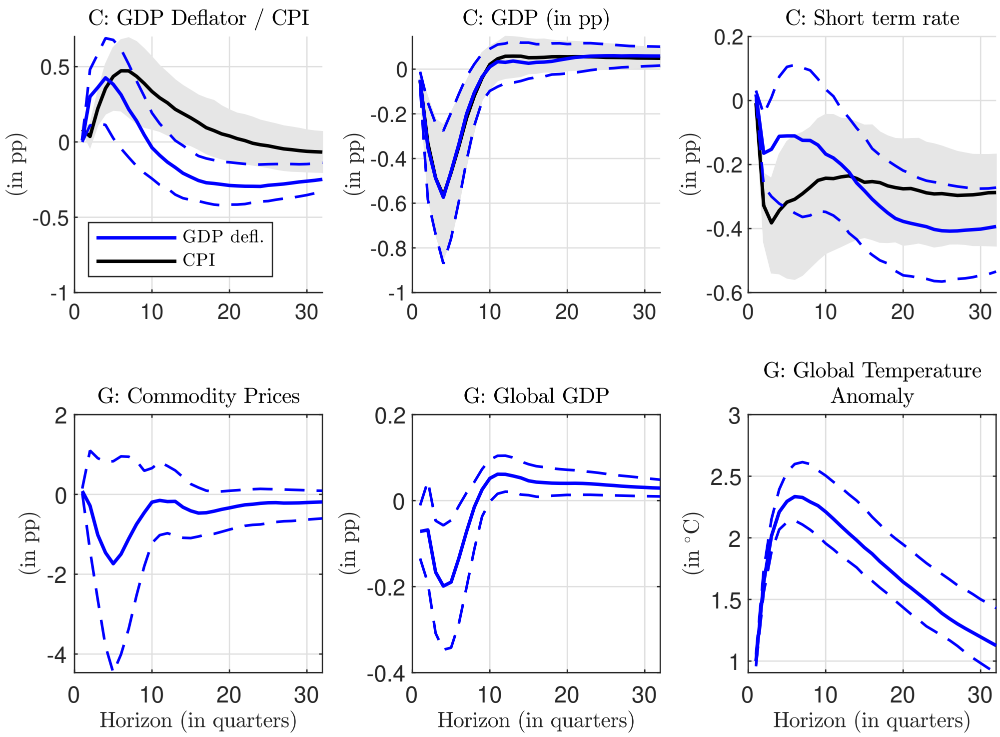

Note: The figure shows the impulse response functions to a shock in the global temperature anomaly. The identification is based on a recursive ordering.

requires that rising temperatures cannot have significant short- and medium-term effects beyond the damage resulting from natural disasters that we have already controlled for in the SVAR. This is discussed in more detail below.

We consider the VAR specification and its block-exogenous structure as presented in the main part as the baseline model and change only one element: the position of the temperature anomaly variable ($\hat{\mathcal{T}}_t$). Whereas in the baseline model it was an interaction variable in the matrix $\tilde{\mathbf{\Psi}}_{c,j}(\hat{\mathcal{T}})$, we now place it in the vector of global variables $\boldsymbol{g}_t$, which is now three-dimensional. Crucially, climate change, as proxied by the global temperature anomaly, is now incorporated into the dynamic system as an endogenous variable, but the feedback from economic activity to climate is limited to the global variables, i.e. we exclude a feedback from home-country variables contained in the vector $\boldsymbol{y}_{c,t}$.

We identify the model by assuming a triangular structure among the contemporaneous variables, which we implement by computing the Cholesky decomposition of the reduced form variance-covariance matrix of the VAR model. In this specification, climate change is ranked first in the vector of endogenous

FIGURE 29. Shock to global temperature anomaly ($\hat{\mathcal{T}}$, with country-temperature anomaly $\hat{\mathcal{T}}_c$)

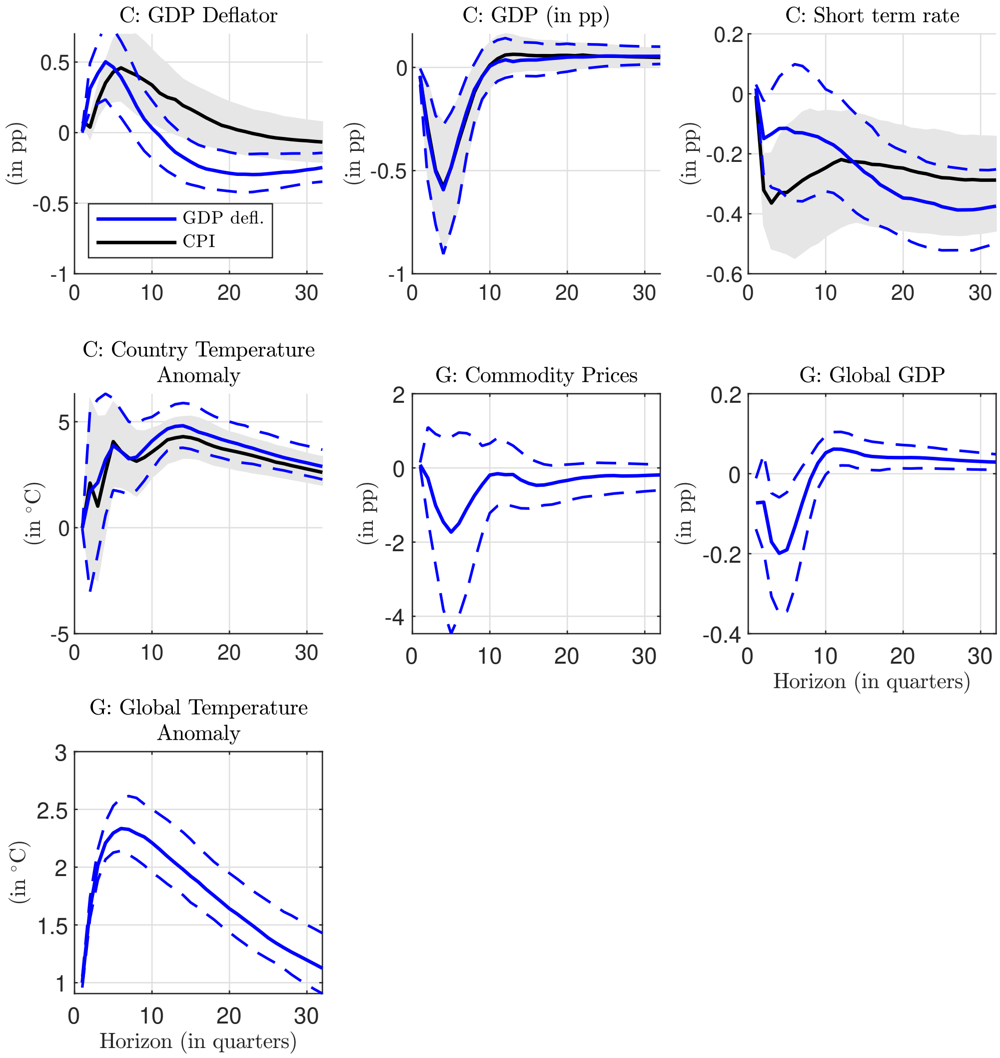

Note: The figure shows the impulse response functions to a shock in the global temperature anomaly. The identification is based on a recursive ordering.

variables $[\boldsymbol{g}_t', \boldsymbol{y}_{c,t}']'$ such that it is the most exogenous variable (given an upper triangular Cholesky decomposition). The objective is to examine whether temperature innovations operate as autonomous business cycle shocks in the short to medium run, and how they propagate to prices, output, and interest rates under a recursive identification scheme.

6.1. **Global temperature anomaly shock.** Figure 28 presents the impulse response results for two model specifications: one using the GDP deflator as the domestic price measure (blue lines) and the other using the CPI (black lines). This distinction only affects the domestic component, as the global price measure is always based on the same commodity index. Importantly, the block-exogenous structure ensures that fluctuations at the country level do not feed back into global dynamics.

The exogenous shock corresponds to an increase in global temperature by one degree Celsius above its average. The temperature anomaly exhibits a hump-shaped adjustment path (median response) before reverting to zero. This trajectory is consistent with the dynamics of a stable higher-order autoregressive process, where the initial rise is followed by a gradual decline, compare Nath et al. (2024).

In terms of global variables, the shock first leads to a decline in both global commodity prices and global GDP, although only the latter displays a response that is statistically significant at the 68 percent confidence level (confidence bands shown by dashed blue lines).

Turning to home-country-level variables, the results are more pronounced. Domestic GDP declines significantly, although the effect is short-lived. For domestic prices, the response varies across specifications. When measured by the CPI, prices increase persistently, and the response is significantly different from zero. In contrast, when measured by the GDP deflator, the initial increase is followed by a decline, resulting in a less clear overall price effect.

6.2. **Global and country temperature anomalies: joint system.** While the results presented in Figure 28 highlight a discernible effect of the global temperature anomaly on the economic system—at least in the short run—an important propagation channel has thus far been ignored in the specification. Specifically, the previous model abstracts from the impact of global temperature fluctuations on the corresponding temperature anomaly at the country level. Yet this propagation mechanism may crucially shape the transmission of climate shocks to domestic economic dynamics, extending the influence of global disturbances beyond their direct effects.

To address this limitation, we enrich the model by incorporating the country-specific temperature anomaly ($\hat{\mathcal{T}}_c$) as an additional element in the vector of country variables, $\boldsymbol{y}_{c,t}$. This extension yields a four-dimensional system of endogenous variables. We then repeat the same experiment as before: introducing a one-degree Celsius increase in the global temperature anomaly. The resulting responses are displayed in Figure 29.

The results reveal that a rise in the global temperature anomaly systematically propagates to the country level, inducing an increase in the domestic anomaly. Notably, the magnitude of the domestic response exceeds that of the global shock and materializes with a delay, indicating that country-specific climatic conditions may amplify and prolong the effects of global disturbances. For both the global and domestic temperature variables, the responses are statistically significantly different from zero at the 68 percent credible interval.

Turning to the macroeconomic variables, the results closely resemble those obtained in the baseline specification (Figure 28). Domestic GDP declines in the aftermath of the temperature shock, while the responses of price measures differ by specification—with the CPI displaying a persistent increase, in contrast to the less conclusive pattern observed for the GDP deflator. Importantly, these qualitative results remain robust across both specifications, underscoring the short-run contractionary effects of climate shocks on real output and pointing to heterogeneous implications for inflation depending on the chosen price index.

Taken together, these findings highlight the relevance of accounting for local climatic dynamics in the transmission process: while global shocks set the direction of adjustment, country-level temperature anomalies can act as amplifiers, shaping both the strength and persistence of the economic response. This feature suggests that climate-induced macroeconomic fluctuations may be more pronounced at the domestic level than what is inferred from global dynamics alone.

6.3. **Country temperature anomaly shock.** As a final exercise, we reconsider the specification underlying Figure 29, but now identify the shock in the country-level temperature anomaly, $\hat{\mathcal{T}}_c$. By construction, the block-exogenous structure of the VAR implies that innovations to $\hat{\mathcal{T}}_c$ affect only the country-level variables, with no feedback to the global block. Consequently, Figure 30 reports impulse responses exclusively for the country variables, even though the full state vector continues to include the global variables.

The shock is calibrated to raise the domestic temperature anomaly by one degree Celsius on impact. As shown in Figure 30, the ensuing dynamics display a rapid reversion of $\hat{\mathcal{T}}_c$ toward zero, indicating relatively limited persistence in country-level temperature anomalies. This contrasts with the more inertial,

FIGURE 30. Shock to country temperature anomaly ($\hat{\mathcal{T}}_c$)

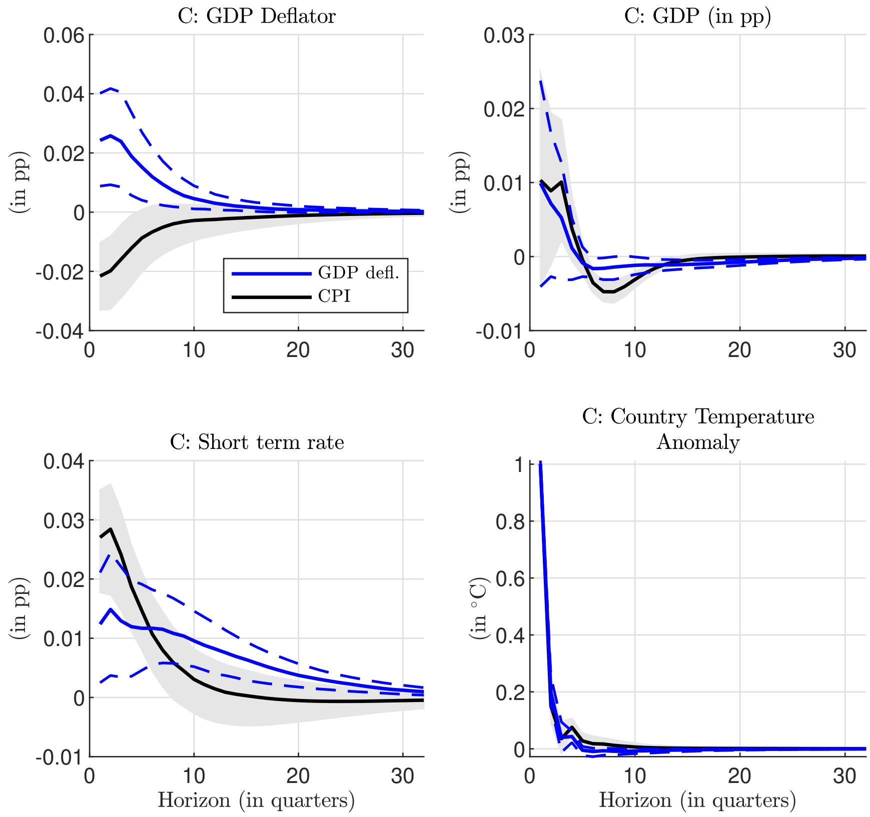

Note: The figure shows the impulse response functions to a shock in the country temperature anomaly. The identification is based on a recursive ordering.

hump-shaped decay observed for global temperature anomalies in Figures 28 and 29, suggesting that local climatic perturbations dissipate more quickly than global ones.

Turning to macroeconomic outcomes, the responses exhibit notable specification sensitivity. We estimate two variants that differ only in the domestic price measure: one using the GDP deflator (blue lines) and the other using the CPI (black lines). Output tends to decline following the temperature shock (i.e., higher temperature), but the decrease attains statistical significance only under the CPI specification, whereas no statistically significant output response emerges when the GDP deflator is used. In both variants, the decline in output is preceded by a short-lived uptick that is not statistically distinguishable from zero.

Price dynamics are even more ambiguous. For the same shock, prices rise when measured by the GDP deflator, but fall when measured by the CPI, with both responses statistically different from zero at the 68 percent level. This divergence underscores that the sign of the inflationary effect depends on the chosen price index, potentially reflecting distinct sensitivities across consumption versus value-added price baskets. Notably, this form of sign reversal does not appear in the interacted VAR specification that conditions responses on temperature via interaction terms, where price effects are comparatively more uniform.

Overall, these results suggest that country-specific temperature shocks are short-lived yet can exert measurable, though fairly small real and nominal effects in the short run, with their perceived inflationary consequences contingent on the price measure employed. This replicates the results in Ciccarelli and Marotta (2024). From a modeling perspective, the contrast between CPI and GDP deflator responses cautions against interpreting "inflation" effects without attention to index construction and sectoral exposure embedded in the respective measures.

7. Estimation of structural parameters of the DSGE model

Our baseline simulation of the DSGE model fundamentally depends on the adverse effect of temperature anomalies on total factor productivity (TFP), captured by the following relationship:

$$\xi_A(\hat{\mathcal{T}}) = e^{-\alpha \hat{\mathcal{T}} + \bar{\xi}_A}. \tag{15}$$

Throughout the preceding analysis, the temperature anomaly $\hat{\mathcal{T}}$ has been exogenously specified and the model has been simulated across a range of values for this parameter. Crucially, variations in $\hat{\mathcal{T}}$ affect TFP through the structural parameter α. In order to enhance the credibility of our simulation exercises and to quantify the impact of temperature anomalies on TFP—a relationship firmly supported by recent empirical studies (Donadelli et al., 2022; Kumar and Maiti, 2024; Drescher and Janzen, 2025)—we now proceed to estimate the structural parameters of the model.

Given the current methodological limitations in panel estimation for DSGE models spanning multiple countries, we conduct the estimation separately for each of the $N = 20$ countries in our sample. Consistent with the baseline configuration, temperature anomalies are treated as strictly exogenous: their observed time series are used to estimate their effect on TFP, but any reverse causality from economic activity (e.g., carbon emissions) to the temperature anomaly is omitted. This approach maintains the short-run perspective pertinent to business cycle analysis, in which the direction of causality runs from temperature anomalies to economic outcomes.

Several key structural parameters governing in-model dynamics are estimated. Except for α—which influences both the steady-state and transitional dynamics—all parameters determine the propagation mechanisms of the model. Specifically, we estimate the following: the Taylor rule coefficients ϕ_π (inflation) and ϕ_y (output); the price stickiness parameter ξ_p; autocorrelation coefficients for exogenous global supply and demand shocks ($\rho_{P_e^*}$ and ρ_{Y^*}, respectively); and the interest rate smoothing parameter ϕ_i associated with monetary policy. Standard prior distributions are adopted: φ_π^{MP} (inflation coefficient) follows a normal distribution (mean 1.5, standard deviation 0.1); φ_y^{MP} (output coefficient) is normal (mean 0.5, standard deviation 0.15); the interest rate smoothing parameter ρ_i is beta (mean 0.5, standard deviation 0.1); AR(1) coefficients $\rho_{P_e^*}$ and ρ_{Y^*} are beta (mean 0.7, standard deviation 0.1); the TFP-temperature sensitivity α is normal (mean 0.0, standard deviation 0.2); and the Calvo price rigidity parameter ξ_p is beta (mean 0.7, standard deviation 0.1).

Addressing potential issues of stochastic singularity (Ruge-Murcia, 2007) and to ensure a non-singular covariance structure in the likelihood function, the vector of structural shocks ($\boldsymbol{u}_t$) is augmented. In addition to global energy price and output shocks, we include: a domestic demand shock (affecting demand for domestically produced consumption goods), a domestic supply shock (embedded in the New-Keynesian Phillips Curve), a monetary policy shock (affecting the Taylor rule), and a temperature anomaly shock (modeled as $\hat{\mathcal{T}}_t \to \hat{\mathcal{T}}_t + \nu_t^{\hat{\mathcal{T}}}$). We expand the shock vector to explicitly include these aforementioned domestic components. The standard deviations of these structural shocks are also estimated, each with an inverse gamma prior: foreign energy price shock ($\nu_{P_e^*}$; scale 0.1, shape 2), foreign output shock (ν_{Y^*}; scale 0.85, shape 2), monetary policy shock (ν_i; scale 0.1, shape 2), domestic supply shock (ν_{DS}; scale 0.1, shape 2), domestic demand shock (ν_{DD}; scale 0.1, shape 2), and temperature anomaly shock ($\nu^{\mathcal{T}}$; scale 0.1, shape 2).

Table 2. Estimates for α across countries.

Country		CI_L (10%)	Median	CI_U (90%)
Australia	AU	0.13	0.15	0.15
Austria	AT	0.39	0.46	0.51
Belgium	BE	0.18	0.21	0.24
Canada	CA	-0.02	0.05	0.11
Finland	FI	-0.05	0.15	0.27
France	FR	0.16	0.20	0.24
Greece	GR	0.01	0.03	0.06
India	IN	0.02	0.05	0.08
Italy	IT	0.05	0.07	0.08
Japan	JP	0.55	0.66	0.78
Korea	KR	-0.05	-0.02	0.02
Mexico	MX	-0.01	0.08	0.13
Netherlands	NL	0.13	0.16	0.20
Norway	NO	0.16	0.20	0.26
New Zealand	NZ	0.41	0.44	0.47
Portugal	PT	0.04	0.11	0.19
Spain	ES	0.12	0.14	0.16
Sweden	SE	-0.02	0.02	0.05
Utd. Kingdom	UK	-0.06	-0.01	0.03
USA	US	0.03	0.04	0.08

These specifications for the priors are selected to reflect standard empirical practice and to ensure structural parameters remain within empirically plausible ranges (Herbst and Schorfheide, 2016; Filippeli et al., 2020). Meanwhile, country-specific parameters describing economic structure—such as the share of domestically produced goods, energy consumption and production shares, the share of renewables, and the consumption-output ratio—are calibrated directly from available country data and held fixed in the estimation.

All empirical and model-based variables used within the estimation exercise are transformed into year-over-year growth rates, except for the nominal interest rate and temperature anomaly, which are included in levels. The set of observables aligns with those employed in estimation of the IPVAR model: global and domestic output and price indices, the domestic nominal interest rate, and the temperature anomaly. Table 2 presents the estimated posterior distributions for α, the parameter governing the impact of temperature anomalies on TFP. For each country, the table reports the lower bound (10%), median, and upper bound (90%) of the posterior credible interval.

The estimated values of α across countries are summarised in Table 2. Several observations are worth highlighting.

Range and cross-country consistency. The median posterior estimates range from -0.02 (Korea) to 0.66 (Japan), with the majority of countries concentrated between 0.03 and 0.46. For 18 out of 20 countries the median is positive; the two exceptions—Korea (-0.02) and the United Kingdom (-0.01)—have credible intervals that include zero and are therefore statistically indistinguishable from zero. Negative point estimates thus do not constitute evidence against the core assumption; they merely reflect estimation uncertainty in countries for which the temperature–TFP link appears weak. The sign pattern is broadly consistent across countries and corroborates the directional prediction of equation (11): higher temperature anomalies depress TFP.

Economic plausibility. To assess plausibility, it is useful to translate the estimates into implied TFP losses. Given the functional form $\xi_A(\hat{\mathcal{T}}) = e^{-\alpha\hat{\mathcal{T}}+\bar{\xi}_A}$, a median estimate of $\alpha = 0.03$—close to the baseline calibration—implies approximately a 3% reduction in TFP per 1°C rise in the temperature anomaly. This is well aligned with the cross-country estimate of Kumar and Maiti (2024), who document a 3.2% TFP reduction per degree, and with the evidence in Donadelli et al. (2022) and Drescher and Janzen (2025). The cross-country median of the individual median estimates is approximately 0.13, implying a TFP loss of around 13% per degree for the average country, which is larger than the aggregate panel estimate but consistent with the finding in Donadelli et al. (2022) that advanced economies tend to exhibit stronger productivity responses to temperature. Countries with large estimates—such as Japan (0.66) and Austria (0.46)—are predominantly temperate, capital-intensive economies where even modest temperature deviations may impair productivity noticeably. While the upper end of the distribution is economically large, the estimates are not implausible given the documented non-linearities in temperature–productivity relationships and the long sample period used for estimation.

Implications for the validity of the core assumption. Taken together, these results support the key structural assumption embedded in equation (11): that TFP responds negatively to temperature anomalies. The baseline calibration of $\alpha = 0.03$ is conservative—it lies near the lower bound of the cross-country distribution—and can therefore be regarded as a lower-bound estimate of the mechanism's magnitude. The consistency of the sign across 18 of 20 countries, combined with statistical significance in most cases, provides empirical backing for this assumption and reinforces the conclusions drawn from the DSGE simulations in the main part. These results reinforce the conclusions of previous empirical studies (Donadelli et al., 2022; Kumar and Maiti, 2024; Drescher and Janzen, 2025) and provide robust evidence for the key transmission channel from temperature anomalies to economic activity via TFP.

Finally, we note that in the interest of clarity and brevity, we confine the presentation of our estimation results to the parameter α, which is the pivotal structural coefficient shaping the influence of temperature anomalies on economic activity through TFP. This parameter constitutes the central transmission channel emphasized in both our modeling framework and the extant empirical literature. While our estimation procedure yields posterior distributions for all model parameters described above for each of the $N = 20$ countries, reporting the full set of results would necessitate the inclusion of twenty additional tables (one for each country), thereby detracting from our primary focus and obscuring the key mechanism of interest. Detailed estimates for the remaining structural parameters—which mainly govern the propagation dynamics and monetary policy transmission, rather than the direct climate-economy nexus—are therefore omitted. Complete estimation results are available from the authors upon request.

8. Empirical assessment of the TFP channel

The theoretical model of Section 3 of the main part rests on the assumption that the temperature anomaly enters the economy through total factor productivity (TFP), captured by equation (11) of the main part, which posits

$$\xi_A(\hat{\mathcal{T}}) = e^{-\alpha \hat{\mathcal{T}} + \bar{\xi}_A}$$

with $\alpha > 0$. This channel is closely related to the mechanism recently proposed by Bilal and Känzig (2026), who argue that global mean temperature—and rising ocean surface temperature in particular—depresses aggregate economic activity largely because it raises the frequency and severity of extreme climatic events (extreme heat, drought, extreme wind speed, and extreme precipitation). In their accounting exercise, these extreme events account for roughly half of the estimated cumulative impact of a global temperature shock on GDP, a magnitude they attribute to physical destruction of productive capacity and to disruptions of production that map naturally into a decline in measured productivity.

Our empirical framework already speaks directly to this mechanism, since the baseline interacted VAR separately controls for realized natural disasters using the EM-DAT-based damage measure discussed in Section 4.5, over and above the temperature anomaly itself. This allows us to assess whether the disaster channel emphasized by Bilal and Känzig (2026) shows up as a negative effect on output or productivity in our sample. As reported in Section 4.5, we do not find the negative output response to disasters that Bilal and Känzig (2026) document at the global level; if anything, for some countries (e.g., France) the estimated short-run response of output to disaster realizations is *positive* and statistically significant, a pattern we interpret as reflecting reconstruction and repair activity—investment aimed at restoring destroyed capital stock—that temporarily stimulates measured output in the immediate aftermath of a disaster, rather than evidence against a longer-run damaging effect of climate change on productive capacity. This is not necessarily inconsistent with Bilal and Känzig (2026): our interacted panel VAR is estimated at quarterly frequency over business-cycle-relevant horizons (roughly one to eight years), whereas the negative TFP consequences of accumulating climate damage that Bilal and Känzig (2026) identify from global temperature variation may operate at lower frequencies and over longer horizons than our system is designed to capture. We therefore do not treat the short-run reconstruction effect documented for the disaster variable as informative about the sign or magnitude of the temperature–TFP link that is the object of interest in the remainder of this section; instead, we hold the disaster control fixed throughout (as in Section 4.5) precisely so as to isolate the direct effect of the temperature anomaly on productivity, uncontaminated by the separate, and empirically distinct, effect of realized disaster damage.

Against this background, the present section provides direct, reduced-form empirical evidence on the temperature–TFP link itself, complementing the indirect, model-based evidence obtained in Section 7, where the Bayesian estimation recovers α as a structural parameter and reports its posterior distribution country by country. Two complementary tests are reported: a panel regression of country labor productivity on the temperature anomaly, with the country rainfall anomaly entering as a sensitivity check (Section 8.1); and an interacted panel BVAR that adds labor productivity as an additional endogenous variable, used to assess whether the temperature dependence of the transmission of global shocks documented in the main part shows up directly in the productivity equation (Section 8.2). Together, the two exercises pin down the TFP channel from outside the structural model, in which we proxy TFP by labor productivity in both cases.

8.1. **Panel regression: labor productivity and the temperature anomaly.** The simplest direct test of the TFP channel is to ask whether a measure of productivity at the country–quarter level is reduced by the temperature anomaly. Since TFP is not directly observable at quarterly frequency for a broad cross-section of countries, we use labor productivity—the ratio of real GDP to employment—as the natural empirical proxy. The contemporaneous use of the rainfall anomaly as a sensitivity/specificity check serves a complementary purpose: as Section 5.2 shows, the IPVAR results documented in the main part lose their state dependence once the temperature anomaly is replaced by the rainfall anomaly. If the underlying channel is the TFP effect of temperature, the benchmark regression should produce a coefficient on rainfall that is small relative to the coefficient on temperature, consequently it will not survive the inclusion of the temperature anomaly as a competing regressor.

For each country $c = 1, \ldots, N$ and quarter t we construct labor productivity as

$$LP_{c,t} = \frac{GDP_{c,t}}{EMP_{c,t}}, \tag{16}$$

where $GDP_{c,t}$ is the real GDP series used throughout the empirical analysis and $EMP_{c,t}$ is the employment level. Because employment is only partially observed at quarterly frequency for several countries in the sample—in particular for the earlier part of the sample period—we back-fill the missing observations using the country-by-country regression

$$\log EMP_{c,t} = a_c + b_1^c \log IP_{c,t} + b_2^c \log UNEMP_{c,t} + g^c\, t + \varepsilon_{c,t}, \tag{17}$$

where IP is the industrial production index, and this equation is estimated on the periods where all three series are jointly observed. For the five countries in which the unemployment level is itself unavailable for the relevant period (Spain, Portugal, Greece, Mexico, and South Korea), we drop $\log UNEMP_{c,t}$ from the right-hand side and fit the corresponding IP-only specification. The in-sample R^2 ranges from 0.71 to 0.91 across countries, indicating that the back-filled series tracks the observed employment level closely.

We estimate the country-fixed-effects panel regressions

$$\log LP_{c,t} = \alpha_c + b\,\hat{\mathcal{T}}_{c,t} + g\,t + e_{c,t}, \tag{18}$$

$$\Delta_4 \log LP_{c,t} = \alpha_c + b\,\hat{\mathcal{T}}_{c,t} + e_{c,t}, \tag{19}$$

where $\hat{\mathcal{T}}_{c,t}$ is the country temperature anomaly and $\Delta_4 \log LP_{c,t} = \log LP_{c,t} - \log LP_{c,t-4}$ is the year-on-year log change in labor productivity. We then re-estimate the same two equations after replacing the temperature anomaly $\hat{\mathcal{T}}_{c,t}$ with the country rainfall anomaly $\hat{\mathcal{R}}_{c,t}$ (z-scored within country, as in Section 4.4), and finally as a horse race with both anomalies entered jointly as regressors. Standard errors are one-way cluster-robust with clusters defined by the country, following the Cameron–Gelbach–Miller correction (Cameron et al., 2011). The sample comprises $N = 20$ countries and up to $T = 265$ quarters per country (between 5,020 and 5,122 observations depending on the specification).

Table 3 reports the point estimates of the coefficient b, together with cluster-robust standard errors and the implied t-statistic, for each of the two specifications and for both the univariate (one anomaly at a time) and the joint (both anomalies entered together) version of the regression. Specification (19) uses the year-on-year change in labor productivity and therefore eliminates the deterministic trend between rising temperatures and slowly growing productivity that is present in specification (18); it is consequently the cleanest identification of the temperature–productivity link.

Three findings stand out. *First*, the coefficient on the temperature anomaly is negative and statistically significant across all four specifications. The

TABLE 3. Panel regression of labor productivity on the temperature and rainfall anomalies

	Univariate		Joint	
	$\hat{\mathcal{T}}_{c,t}$	$\hat{\mathcal{R}}_{c,t}$	$\hat{\mathcal{T}}_{c,t}$	$\hat{\mathcal{R}}_{c,t}$
Spec. (18): $\log LP_{c,t}$	-0.0226	0.0192	-0.0227	0.0179
	(-1.90)	$(\ \ 1.21)$	(-1.98)	$(\ \ 1.04)$
Spec. (19): $\Delta_4 \log LP_{c,t}$	-0.0048	-0.0012	-0.0047	-0.0010
	(-6.50)	(-1.16)	(-6.26)	(-1.09)

Notes: The table reports the coefficient on the country temperature anomaly $\hat{\mathcal{T}}_{c,t}$ and on the country rainfall anomaly $\hat{\mathcal{R}}_{c,t}$ from country-fixed-effects panel regressions of (18) the log of labor productivity and (19) its year-on-year log change. "Univariate" columns report estimates when each anomaly is entered separately as the only regressor; "Joint" columns report estimates when the two anomalies are entered together as regressors. Specification (18) additionally includes a linear time trend; specification (19) does not. Standard errors are one-way cluster-robust with clusters equal to countries (Cameron–Gelbach–Miller correction); t-statistics in parentheses. The sample comprises $N = 20$ countries over up to $T = 265$ quarters; total joint observations 5,020 to 5,122. The rainfall anomaly enters in country-specific z-scores constructed as described in Section 4.4.

cleanest estimate, obtained from the year-on-year-growth specification (19), is $b = -0.0048$ with a t-statistic of -6.50. The point estimate implies that a 1 °C increase in the country temperature anomaly reduces year-on-year labor productivity growth by approximately 0.48 percentage points. This is consistent with the cross-country evidence on the temperature–productivity nexus reported in Donadelli et al. (2022), Kumar and Maiti (2024) and Drescher and Janzen (2025): Kumar and Maiti (2024), in particular, document a 3.2% reduction in TFP per degree Celsius, of which our annual-growth estimate captures a substantial share of the implied trajectory. The result therefore corroborates, in a fully reduced-form setting, the negative temperature–TFP relationship that is the central building block of equation (11) in the theoretical model of the main part. It underscores that the modelling assumption $\alpha > 0$ embedded in $\xi_A(\hat{\mathcal{T}}) = e^{-\alpha\hat{\mathcal{T}}+\bar{\xi}_A}$ is empirically well-grounded.

Second, the rainfall anomaly is a substantially weaker predictor of labor productivity. In the year-on-year-growth specification, the coefficient is -0.0012 with a t-statistic of -1.16, that is, less than half the absolute size of the temperature coefficient and far less statistically robust. In the level specification, the rainfall coefficient even flips sign, suggesting that the apparent positive association in specification (18) is contaminated by the deterministic trend in the data rather than reflecting a stable empirical relationship. The contrast with the temperature coefficient, which is essentially unchanged across the four specifications, is informative: it indicates that the productivity-relevant climatic variation is temperature-driven, not rainfall-driven (compare Bilal and Känzig, 2026). This finding provides an independent, empirical rationalisation of the central result of Section 5.2: the IPVAR state-dependence vanishes once the temperature anomaly is replaced by the rainfall anomaly precisely because the rainfall anomaly is not a sufficient indicator for the underlying TFP channel.

Third, in the horse-race regressions (right-hand panel of Table 3), the temperature coefficient is essentially unchanged when the rainfall anomaly is included as an additional regressor, while the rainfall coefficient is further attenuated. Temperature absorbs the productivity-relevant climatic variation, but the converse is not true—confirming that the two anomalies are not interchangeable proxies for the same underlying mechanism.

The panel-regression evidence is consistent with the country-level posterior distributions of α reported in Table 2. A median estimate of $\alpha = 0.03$, which is close to the conservative baseline calibration in Table 1 of the main part, implies a steady-state TFP loss of $1 - e^{-0.03} \approx 3\%$ per 1 °C anomaly, while the cross-country median of the country medians is approximately 0.13, implying a TFP loss of around 13% per 1 °C—comparable in magnitude and direction to the labor-productivity effect identified here. These magnitudes also align well with the independent, macro time-series evidence in Bilal and Känzig (2026), who structurally estimate a damage function for TFP from the response of world output to global temperature shocks and obtain a peak short-run productivity loss of approximately 4% following a transitory 1 °C global temperature shock. Despite relying on an entirely different identification strategy—aggregate, global time-series variation in ocean and land temperature rather than country-level panel variation in the temperature anomaly—their estimate falls squarely within the single-digit range implied by our conservative baseline calibration of α, while our cross-country median estimate suggests an even more pronounced productivity response for the average country in our sample. More important than the precise magnitudes is the directional agreement: the panel regression finds a robust negative reduced-form effect of temperature on productivity, the country-level Bayesian estimation finds a positive α (and therefore a negative TFP effect) for 18 of 20 countries, the structural damage function of Bilal and Känzig (2026) implies the same negative and economically sizable productivity response to global temperature, and the cross-country literature converges on a negative temperature–TFP relationship. The model-based and the reduced-form (panel-regressions) evidence converge on the same conclusion: the temperature anomaly is a quantitative drag on productivity, the rainfall anomaly is not, and equation (11) of the theoretical model is the appropriate channel through which the temperature state of the economy shapes the propagation of macroeconomic shocks documented in the main part.

8.2. **Interacted panel BVAR including labor productivity.** The panel regression in Section 8.1 establishes the temperature–productivity link in a fully reduced-form setting. A natural follow-up question is whether the same link materialises *within* the structural VAR framework used in the main part—that is, whether the temperature-dependent transmission of global demand and supply shocks documented in Sections 4.3 of the main part also shows up directly in a productivity equation. If the TFP channel of equation (11) of the main part is the operative mechanism, the same temperature state that amplifies the inflation response and dampens the output response to a global demand shock should also dampen the labor-productivity response to that shock. This subsection runs that test by augmenting the country block of the baseline IPVAR with labor productivity as a fourth endogenous variable.

The baseline IPVAR of Section 4.1 of the main part uses a three-variable country block $\boldsymbol{y}_{c,t} = (\pi_{c,t}, \hat{y}_{c,t}, i_{c,t})'$ together with a two-variable global block $\boldsymbol{y}_{g,t} = (P^*_{e,t}, Y^*_t)'$, where c indexes country-specific and g global variables. We extend the country block by adding labor productivity ($\widehat{lprod}_{c,t}$), defined as the year-on-year log change of $LP_{c,t} = GDP_{c,t}/EMP_{c,t}$, with $EMP_{c,t}$ the back-filled employment series introduced in Section 8.1. The augmented country block is therefore

$$\boldsymbol{y}_{c,t} = \big(\pi_{c,t},\ \hat{y}_{c,t},\ \widehat{lprod}_{c,t},\ i_{c,t}\big)'. \tag{20}$$

As in the baseline specification, the model retains the natural disaster variable described in Section 4.5 as a strictly exogenous regressor, so that the productivity equation, like the output, inflation, and interest rate equations, is estimated conditional on realized disaster damage rather than confounding the two channels. The lag structure, the COVID-19 dummy, the block-exogeneity restriction with respect to the global variables, and the IAV-interaction structure are exactly as in the baseline IPVAR. The interaction term takes the same form as in equation (30) of the main part,

$$\boldsymbol{B}_{11}^{c}(\mathcal{T}_t)\,\boldsymbol{y}_{c,t-1} = \left(\boldsymbol{\Gamma}^{c}(\bar{\mathcal{T}}) + \boldsymbol{\mathcal{F}}^{c}\,\hat{\mathcal{T}}_{c,t}\right)\boldsymbol{y}_{c,t-1}, \tag{21}$$

but $\boldsymbol{\Gamma}^c$ and $\boldsymbol{\mathcal{F}}^c$ are now 4×4 rather than 3×3, and the interaction term enters as provided in equation (21), with the same lag length $\mathcal{P} = 3$ used as in the baseline.

Because output and labor productivity are, almost by construction, positively correlated at business-cycle frequencies, adding $\widehat{lprod}_{c,t}$ alongside $\hat{y}_{c,t}$ in the same country block introduces a degree of collinearity that is absent from the baseline three-variable system. In practice, this shows up as somewhat wider posterior credible bands around the impulse responses of both output and labor productivity relative to their baseline counterparts, reflecting the additional estimation uncertainty associated with disentangling two closely co-moving series from a limited number of observations per country. Reassuringly, however, this loss of precision is modest: the point estimates remain well identified and the qualitative pattern of results discussed below is unaffected, indicating that collinearity between output and productivity attenuates the sharpness, but not the substance, of the augmented IPVAR's conclusions.

Shocks are identified by sign restrictions on the contemporaneous impact matrix, extended to the four-variable country block in the natural way. The sign convention is the one used in the main part: an expansionary global demand shock raises inflation and output at home and abroad, raises the domestic policy rate, and we leave labor productivity unrestricted; a contractionary global supply shock raises home and global commodity-price inflation, lowers home and global output, and again we leave labor productivity unrestricted. The supply shock is plotted in expansionary normalisation, as in Figure 6 of the main part.

The baseline IPVAR aggregates 20 country-level posteriors via the mean-group estimator of Pesaran and Smith (1995). Adding labor productivity forces us to drop five countries from the sample—Spain, Portugal, Greece, Mexico, and South Korea—for which the quarterly employment series in Section 8.1

is unfortunately too short (more than half of the observations would have to be back-filled). The remaining sample comprises $N = 15$ countries. With a smaller cross-section, the mean-group estimator averages noisier country-level posteriors. We therefore depart from the baseline convention here and estimate the IPVAR by *pooling* the country panels into a single Bayesian VAR with country-common slope coefficients, as in the alternative panel methodology already documented for the baseline specification (see Section 5, "Alternative panel methodology"). The pooled estimator stacks the $N \times T$ observations and produces tighter posterior bands at the cost of imposing slope homogeneity across countries—exactly the pooled-vs-mean-group trade-off discussed for the baseline IPVAR. Two further considerations support pooling for this exercise: (i) the TFP channel under test is a *common* structural mechanism whose strength should not, in principle, vary across countries in a manner correlated with sample size; and (ii) the joint estimation of the slope coefficients on the four-variable country block, the two-variable global block, and the interaction terms with the temperature anomaly costs $4 + 2 + (4 \cdot \mathcal{P}) = 18$ parameters per equation, which is comfortable with a 5,000-observation pooled panel but stretches the country-level samples on which the mean-group estimator relies. The pooled BVAR is implemented using the same priors and the same Gibbs-sampling scheme as in Section 5; the only modification is the inclusion of labor productivity as a fourth country variable in the stacked endogenous vector.

Figures 31 and 32 report the impulse responses to the global demand shock and to the (expansionary) global supply shock, respectively. Each figure has the same layout used for Figures 4 and 6 of the main part, evaluated at the 5th and 95th percentiles of the pooled temperature anomaly distribution ($\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$).

Three findings stand out. First, the augmented IPVAR reproduces the central qualitative results of the baseline IPVAR. Following an expansionary global demand shock (Figure 31), domestic inflation rises and the inflation response is markedly larger when the temperature anomaly is high than when it is low; the corresponding "Difference (Low–High)" panel for inflation lies systematically below zero. Following an expansionary global supply shock (Figure 32), domestic prices fall, output and labor productivity rise, the short-term interest rate declines, and the Low–High difference is essentially indistinguishable

FIGURE 31. Demand shock: pooled IPVAR with labor productivity

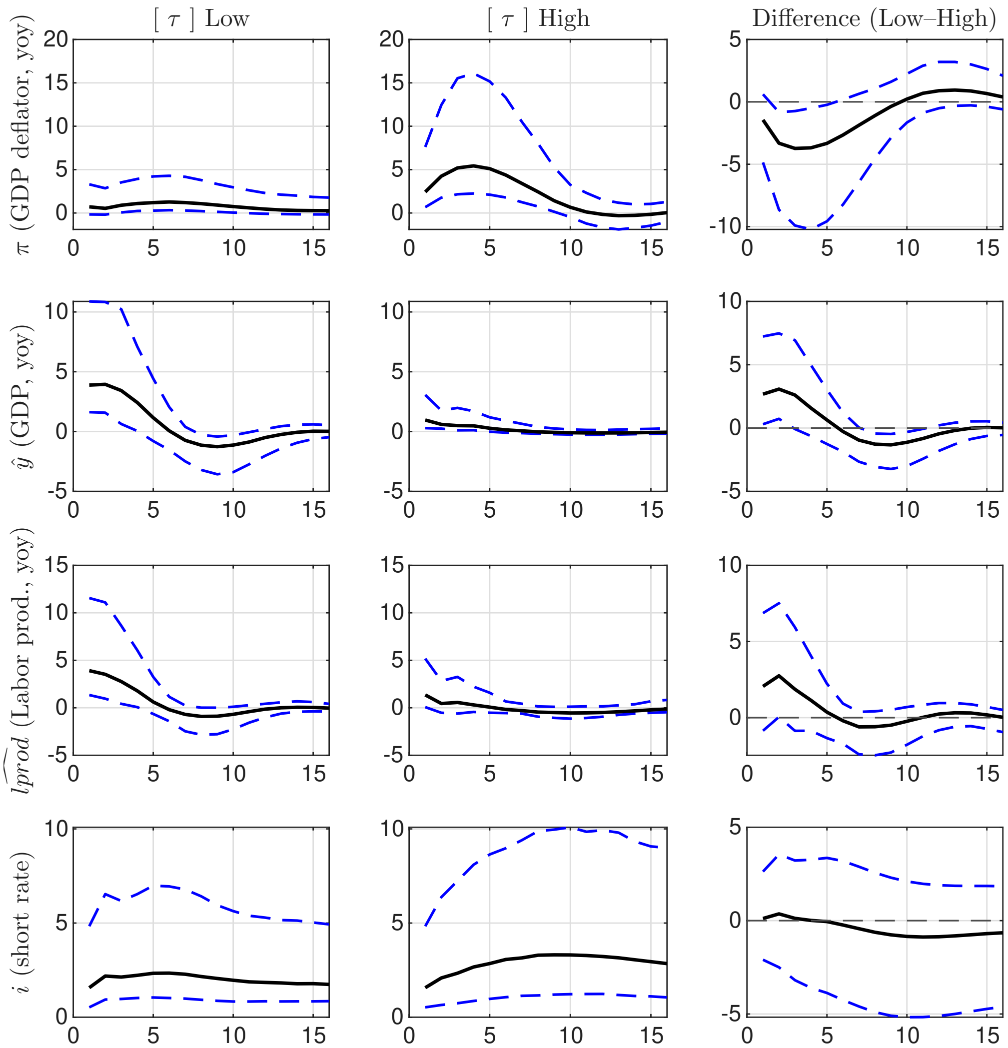

Note: Impulse responses to a global demand shock from the pooled IPVAR extended with labor productivity as a fourth country-block endogenous variable, identified by sign restrictions following the convention of Section 4.1.3 of the main part (with $\widehat{lprod}+$ on impact in the demand-shock column). Columns show low and high values of the country temperature anomaly ($\hat{\mathcal{T}}_{\text{Low}} = -2$ and $\hat{\mathcal{T}}_{\text{High}} = 3.8$) and the difference between the two. Solid lines: posterior medians; dashed lines: 68 percent credible bands. The sample comprises $N = 15$ countries (Spain, Portugal, Greece, Mexico, and South Korea are excluded because of a lack of data).

from zero across all variables. The baseline message—temperature anomalies amplify the inflationary impact of demand shocks while leaving supply-shock propagation largely unchanged—is therefore preserved when labor productivity is added to the country block.

Second, the new productivity equation provides the direct, model-internal corroboration of the TFP channel that motivated the exercise. The impact response of labor productivity to the demand shock is positive at both low

FIGURE 32. Supply shock (expansionary): pooled IPVAR with labor productivity

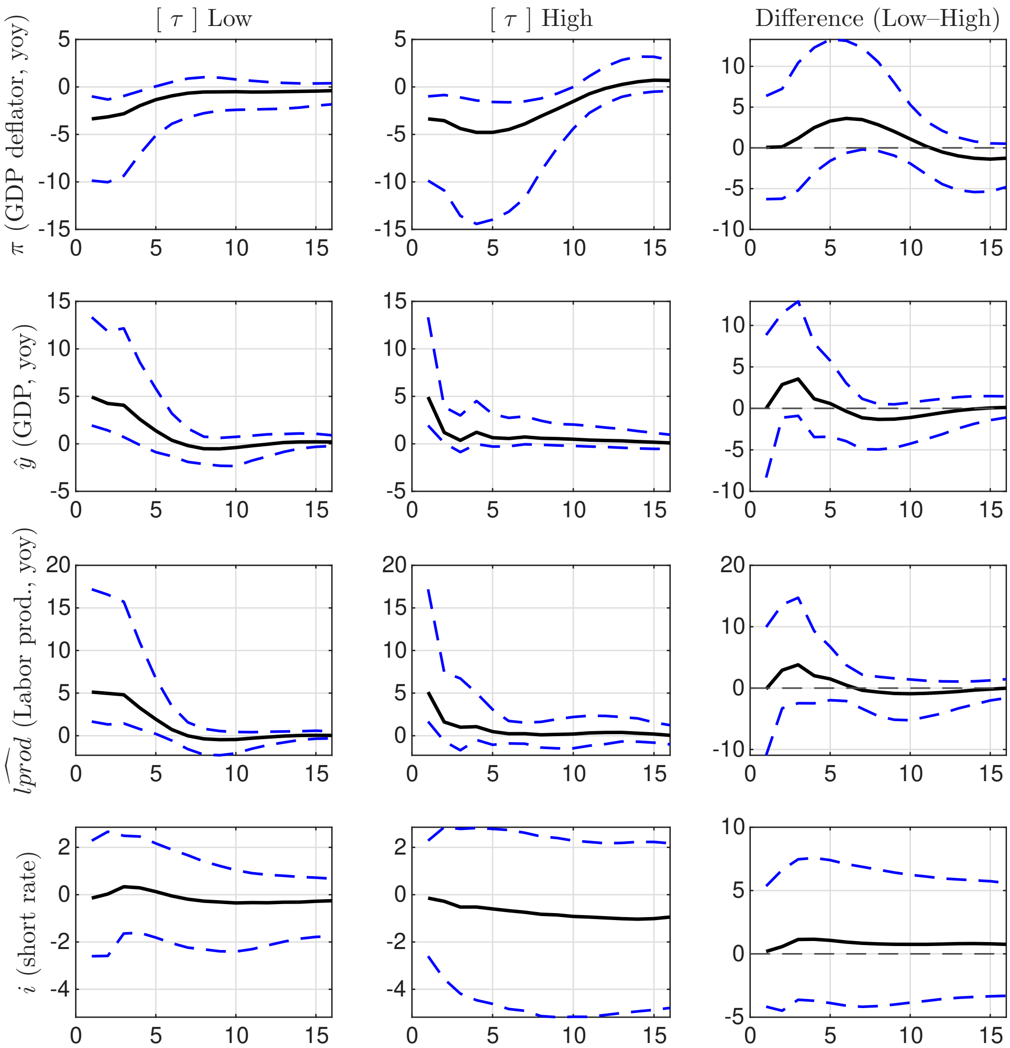

Note: Impulse responses to a global supply shock (plotted in expansionary normalisation, mirroring Figure 6 of the main part) from the pooled IPVAR extended with labor productivity as a fourth country-block endogenous variable. Layout and bands as in Figure 31.

and high temperature anomalies, reflecting the standard cyclical comovement of productivity with output. The temperature dependence emerges in the propagation dynamics: at high temperature the productivity response is more contained than at low temperature over the first few quarters, which is exactly what one would predict from a model in which higher temperature anomalies depress TFP and therefore prevent productivity from rising as much for a given expansion of aggregate demand. For the supply shock, by contrast, the productivity response is qualitatively the same across temperature states—consistent with the theoretical prediction that the TFP channel reshapes the slope of the Phillips curve and hence shapes the transmission of *demand* shocks, but does not materially affect the transmission of supply shocks.

Third, the magnitudes of the temperature-dependent effects are in line with the country-level Bayesian estimates of α in Table 2. At the cross-country median estimate of α, the implied TFP penalty for a 1 °C anomaly is of the order of at least 3 percent, depending on which moment of the cross-country posterior is used. The pooled IPVAR translates this structural penalty into a noticeable, but bounded, attenuation of the productivity response to demand shocks at high temperature, which is the empirical counterpart of the steepened Phillips curve documented in the simulations of Section 3 of the main part.

The three pieces of evidence assembled in this Appendix—the panel regression of Section 8.1, the Bayesian estimation of α in Section 7, and the IPVAR with labor productivity reported here—are deliberately complementary in their assumptions and their identification. The panel regression is a reduced-form correlation conditional on country fixed effects and a linear trend; the Bayesian estimation is a structural model-based recovery of α under the maintained DSGE specification; the IPVAR extension is a structural exercise that disciplines the productivity response via sign restrictions derived from the theory but lets the temperature-state dependence be governed by the data. The three exercises converge on the same conclusion: temperature anomalies depress labor productivity / TFP, the magnitude is in the single-digit percent range per degree Celsius, and the consequence for macroeconomic dynamics is exactly the asymmetry between demand and supply shocks that the baseline IPVAR documents. This pattern is also consistent with the broader cross-country evidence on the temperature–productivity nexus reported in Donadelli et al. (2022), Kumar and Maiti (2024), Drescher and Janzen (2025), and—in a different identification scheme—in the temperature-shock literature surveyed by Bilal and Stock (2025). Relative to that body of work, the contribution of the present paper is to show that the temperature–TFP link not only reduces the *level* of productivity but also reshapes the *propagation* of conventional business-cycle shocks: it makes prices more sensitive and output less sensitive to demand disturbances, with labor productivity tracing out the intermediate step in this transmission.

8.3. **The Temperature Anomaly and the Global Productivity Slowdown.** A growing body of evidence suggests that the post-2005 deceleration in labor productivity and total factor productivity (TFP) is not attributable to any single cause but rather reflects the confluence of several mutually reinforcing structural forces. Goldin et al. (2024) provide a comprehensive, quantitative synthesis of these forces for five advanced economies—France, Germany, Japan, the United Kingdom, and the United States—and conclude that the observed slowdown in TFP growth of 0.8 to 1.8 percentage points per year is most parsimoniously explained by a decline in capital deepening, lower spillovers from intangible capital, a slowdown in global trade, and reduced allocative efficiency. Crucially, their decomposition identifies the TFP residual as the single largest contributor to the aggregate productivity slowdown across all countries considered.

Within their framework, Goldin et al. (2024) further distinguish between cyclical and secular components of the capital deepening slowdown—the former driven primarily by the aftermath of the global financial crisis, the latter by deeper structural forces including the compositional shift of capital toward intangible assets, declining competitive intensity, and changes in corporate governance and globalization (see their Tables 11 and 12). Importantly, climate change and its consequences for productive capacity are not explicitly accounted for in their analysis. Our previous findings suggest that rising temperature anomalies potentially constitute such a secular force: one that is persistent by nature, geographically pervasive, and—as our theoretical and empirical analysis demonstrates—operates directly through a suppression of TFP. This climate-induced TFP drag therefore fits naturally within what Goldin et al. (2024) term "Capital: secular trends," a category that, by their own account, encompasses structural determinants of productivity.

The quantitative implications of this conjecture are non-trivial. Our panel regression results in Section 8.1 indicate that a one-degree-Celsius increase in the temperature anomaly is associated with a statistically significant reduction in labor productivity growth of approximately 0.48 percentage points per year. Given that the temperature anomaly across the countries in our sample has risen by roughly one degree Celsius over the post-1977 period—the structural break identified in our empirical specification—this estimate implies a climate-induced drag on productivity of a magnitude that is comparable to the "missing" productivity growth documented by Goldin et al. (2024). Moreover, the temperature anomaly satisfies each of the three criteria that Goldin et al. (2024) themselves propose for a credible explanation of the productivity slowdown: it is quantitatively significant at the scale required, it exhibits a clear trend that predates the post-2005 break (thereby satisfying the sequencing criterion), and its effects are geographically broad in scope, spanning advanced and emerging economies alike.

Taken together, our findings suggest that the role of climate change as a driver of secular productivity decline deserves recognition in the broader literature on the productivity slowdown. While Goldin et al. (2024) conclude that no single explanation suffices, the temperature anomaly—operating through the TFP channel documented in this paper—constitutes a quantitatively relevant and theoretically well-grounded addition to the set of factors that collectively account for the post-2005 deceleration in labor productivity and TFP. Integrating climate change into the growth-accounting framework advanced by Goldin et al. (2024) thus represents an interesting direction for future research.

References

Baumeister, Christiane, and James D. Hamilton (2015) 'Sign restrictions, structural vector autoregressions, and useful prior information.' *Econometrica* 83(5), 1963–1999

Bettarelli, Luca, Davide Furceri, Loredana Pisano, and Pietro Pizzuto (2025) 'Greenflation: Empirical evidence using macro, regional and sectoral data.' *European Economic Review* p. 104983

Bilal, Adrien, and Diego R Känzig (2026) 'The macroeconomic impact of climate change: Global versus local temperature.' *The Quarterly Journal of Economics* 141(2), 889–944

Bilal, Adrien, and James H. Stock (2025) 'A guide to macroeconomics and climate change.' NBER Working Papers 33567, National Bureau of Economic Research, Inc, Mar

Blanchard, Olivier J., and David R. Johnson (2017) *Macroeconomics* (Pearson)

Botzen, WJ Wouter, Olivier Deschenes, and Mark Sanders (2019) 'The economic impacts of natural disasters: A review of models and empirical studies.' *Review of Environmental Economics and Policy*

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (2015) 'Global non-linear effect of temperature on economic production.' *Nature* 527(7577), 235–239

Byrne, Joseph P., and Prince Asare Vitenu-Sackey (2024) 'The macroeconomic impact of global and country-specific climate risk.' *Environmental and Resource Economics* 87(3), 655–682

Caggiano, Giovanni, Efrem Castelnuovo, and Giovanni Pellegrino (2017) 'Estimating the real effects of uncertainty shocks at the zero lower bound.' *European Economic Review* 100, 257–272

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller (2011) 'Robust inference with multiway clustering.' *Journal of Business & Economic Statistics* 29(2), 238–249

Canova, Fabio, and Matteo Ciccarelli (2013) 'Panel vector autoregressive models: A survey.' CEPR Discussion Papers 9380, C.E.P.R. Discussion Papers, Mar

Casey, Gregory, Stephie Fried, and Ethan Goode (2023) 'Projecting the impact of rising temperatures: The role of macroeconomic dynamics.' *IMF Economic Review* 71(3), 688–718

Cevik, Serhan, and Joao Jalles (2024) 'Eye of the storm: the impact of climate shocks on inflation and growth.' *Review of Economics* 75(2), 109–138

Chavas, Jean-Paul, Corbett Grainger, and Nicholas Hudson (2016) 'How should economists model climate? Tipping points and nonlinear dynamics of carbon dioxide concentrations.' *Journal of Economic Behavior & Organization* 132(PB), 56–65

Ciccarelli, Matteo, and Fulvia Marotta (2024) 'Demand or Supply? An empirical exploration of the effects of climate change on the macroeconomy.' *Energy Economics* 129(C), S0140988323006618

Ciccarelli, Matteo, Friderike Kuik, and Catalina Martínez Hernández (2024) 'The asymmetric effects of temperature shocks on inflation in the largest euro area countries.' *European Economic Review* p. 104805

Dell, Melissa, Benjamin Jones, and Benjamin Olken (2014) 'What do we learn from the weather? The new climate-economy literature.' *Journal of Economic Literature* 52(3), 740–98

Deschênes, Olivier, and Michael Greenstone (2011) 'Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US.' *American Economic Journal: Applied Economics* 3(4), 152–185

Donadelli, Michael, Marcus Jüppner, and Sergio Vergalli (2022) 'Temperature variability and the macroeconomy: A world tour.' *Environmental and resource economics* 83(1), 221–259

Drescher, Katharina, and Benedikt Janzen (2025) 'When weather wounds workers: The impact of temperature on workplace accidents.' *Journal of Public Economics* 241, 105258

Emediegwu, Lotanna E., and Chisom L. Ubabukoh (2023) 'Re-examining the impact of annual weather fluctuations on global livestock production.' *Ecological Economics* 204, 107662

Filippeli, Thomai, Richard Harrison, and Konstantinos Theodoridis (2020) 'Dsge-based priors for bvars and quasi-bayesian dsge estimation.' *Econometrics and Statistics* 16(C), 1–27

Food, and Agriculture Organization of the United Nations (2023) 'Temperature change statistics 1961–2022: Global, regional and country trends.' Technical Report, FAOSTAT Analytical Brief 63, Rome. ISSN 2709-006X (Print), ISSN 2709-0078 (Online)

Food and Agriculture Organization of the United Nations (2025) 'Temperature change statistics 1961–2024: Global, regional and country trends.' Technical Report, FAOSTAT Analytical Brief No. 101, Rome. ISSN 2709-006X (Print), ISSN 2709-0078 (Online)

Fry, Renée, and Adrian Pagan (2011) 'Sign Restrictions in Structural Vector Autoregressions: A Critical Review.' *Journal of Economic Literature* 49(4), 938–960

Giorgi, Filippo, Francesca Raffaele, and Erika Coppola (2019) 'The response of precipitation characteristics to global warming from climate projections.' *Earth System Dynamics* 10(1), 73–89

Glocker, Christian, and Pascal Towbin (2015) 'Reserve requirements as a macroprudential instrument - Empirical evidence from Brazil.' *Journal of Macroeconomics* 44(C), 158–176

Glocker, Christian, and Philipp Piribauer (2025) 'Consumer preferences and inflation diffusion.' *Macroeconomic Dynamics* 29, e80

Goldin, Ian, Pantelis Koutroumpis, François Lafond, and Julian Winkler (2024) 'Why is productivity slowing down?' *Journal of Economic Literature* 62(1), 196–268

Granger, Clive WJ (1998) 'Overview of nonlinear time series specification in economics.' In 'Berkeley NSF-Symposia'

Hamilton, James D. (1994) *Time Series Analysis* (Princeton University Press)

Hassler, John, and Per Krusell (2012) 'Economics and climate change: Integrated assessment in a multi-region world.' *Journal of the European Economic Association* 10(5), 974–1000

Hassler, John, Per Krusell, and Jonas Nycander (2016) 'Climate policy.' *Economic Policy* 31(87), 503–558

Held, Isaac M, and Brian J Soden (2006) 'Robust responses of the hydrological cycle to global warming.' *Journal of climate* 19(21), 5686–5699

Herbst, Edward, and Frank Schorfheide (2016) *Bayesian Estimation of DSGE Models,* 1 ed. (Princeton University Press)

Huber, Florian, Tamás Krisztin, and Michael Pfarrhofer (2023) 'A Bayesian panel vector autoregression to analyze the impact of climate shocks on high-income economies.' *The Annals of Applied Statistics* 17(2), 1543–1573

Jordà, Òscar (2005) 'Estimation and Inference of Impulse Responses by Local Projections.' *American Economic Review* 95(1), 161–182

Kumar, Naveen, and Dibyendu Maiti (2024) 'Long-run macroeconomic impact of climate change on total factor productivity — evidence from emerging economies.' *Structural Change and Economic Dynamics* 68(C), 204–223

Manne, Alan, Robert Mendelsohn, and Richard Richels (1995) 'MERGE: A model for evaluating regional and global effects of GHG reduction policies.' *Energy Policy* 23(1), 17–34

Mittnik, Stefan (1990) 'Modeling nonlinear processes with generalized autoregressions.' *Applied Mathematics Letters* 3(4), 71–74

Mukherjee, K., and Bazoumana Ouattara (2021) 'Climate and monetary policy: Do temperature shocks lead to inflationary pressures?' *Climatic Change* 167(3), 1–21

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (2024) 'How much will global warming cool global growth?' NBER Working Papers 32761, National Bureau of Economic Research, Inc, July

Noy, Ilan (2009) 'The macroeconomic consequences of disasters.' *Journal of Development economics* 88(2), 221–231

Olea, José Luis Montiel, Mikkel Plagborg-Moller, Eric Qian, and Christian Wolf (2024) 'Double robustness of local projections and some unpleasant VARithmetic.' NBER Working Papers 32495, National Bureau of Economic Research, Inc

Palagi, Elisa, Matteo Coronese, Francesco Lamperti, and Andrea Roventini (2022) 'Climate change and the nonlinear impact of precipitation anomalies on income inequality.' *Proceedings of the national academy of sciences* 119(43), e2203595119

Parker, Miles (2018) 'The impact of disasters on inflation.' *Economics of Disasters and Climate Change* 2(1), 21–48

Pesaran, M. Hashem, and Ron Smith (1995) 'Estimating long-run relationships from dynamic heterogeneous panels.' *Journal of Econometrics* 68(1), 79–113

Plagborg-Møller, Mikkel, and Christian Wolf (2021) 'Local projections and VARs estimate the same impulse responses.' *Econometrica* 89(2), 955–980

Pretis, Felix (2021) 'Exogeneity in climate econometrics.' *Energy Economics* 96, 105122

Rubio-Ramírez, Juan F, Daniel F Waggoner, and Tao Zha (2010) 'Structural vector autoregressions: Theory of identification and algorithms for inference.' *The Review of Economic Studies* 77(2), 665–696

Ruge-Murcia, Francisco (2007) 'Methods to estimate dynamic stochastic general equilibrium models.' *Journal of Economic Dynamics and Control* 31(8), 2599–2636

Sachs, Jeffrey D., and Felipe Larraín B. (1993) *Macroeconomics in the Global Economy* (Prentice Hall)

Sapir, Debarati G., and Claudine Misson (1992) 'The development of a database on disasters.' *Disasters* 16(1), 74–80