

Quantifying climate damages when regions trade: a structural gravity approach

**Jeanne Astier, Geoffrey Barrows, Raphael Calel and
Helene Ollivier**

May 2026

**Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 447**

ISSN 2515-5717 (Online)

The Grantham Research Institute on Climate Change and the Environment was established in 2008 at the London School of Economics and Political Science. The Institute brings together international expertise on economics, as well as finance, geography, the environment, international development and political economy to establish a world-leading centre for policy-relevant research, teaching and training in climate change and the environment. It is hosted by the Global School of Sustainability at LSE and is funded by the Grantham Foundation for the Protection of the Environment.

www.lse.ac.uk/granthaminstitute

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one referee before publication. The views in this paper are those of the authors and do not necessarily represent the position of the Grantham Research Institute's senior management or funders. Any errors and omissions remain those of the authors.

Authors' declaration of AI use: The authors made limited use of AI-based language tools (e.g., ChatGPT, Claude) to refine prose and to assist in checking selected mathematical derivations. All ideas, analyses, and conclusions are the authors' own, and the authors take full responsibility for the content of the manuscript.

This paper was first published in May 2026 by the Grantham Research Institute on Climate Change and the Environment at the London School of Economics and Political Science.

© The authors, 2026

Licensed under [CC BY-NC 4.0](https://creativecommons.org/licenses/by-nc/4.0/)

Suggested citation:

Astier J, Barrows G, Calel R and Ollivier H (2026) *Quantifying climate damages when regions trade: a structural gravity approach*. Grantham Research Institute on Climate Change and the Environment Working Paper X. London: London School of Economics and Political Science

DOI: [10.21953/researchonline.lse.ac.uk.00138243](https://doi.org/10.21953/researchonline.lse.ac.uk.00138243)

Quantifying Climate Damages When Regions Trade: A Structural Gravity Approach*

Jeanne Astier	Geoffrey	Raphael Calel	Helene Ollivier
<i>CREST</i>	Barrows	<i>McCourt School,</i>	<i>PSE-CNRS</i>
	<i>CREST-CNRS</i>	<i>Georgetown</i>	
		<i>University</i>	

This version: April 3, 2026

Abstract

This paper presents a method for estimating treatment effects of local climate shocks when regions trade with each other. Because trade creates spillovers, comparing the change in outcomes of regions with different exposure to shocks leads to biased estimates. We model these between-region spillovers using standard assumptions from international trade theory, and develop a model-consistent strategy for estimating key parameters and deriving counterfactuals. We use our estimation strategy to revisit the literature on the impact of climate change on gross output. We find that accounting for trade spillovers yields substantially larger climate damage projections.

Keywords: climate change , spillovers, trade, gravity

JEL Classification: Q48, L1, L5

*We would like to thank Jean-Marc Bourgeon, Kirill Borusyak, Gregory Corcos, Lucie Gadenne, Matthew Gordon, Christophe Gouel, Florian Grosset-Touba, Olivier Deschênes, Xavier D'Haultfœuille, Andrew Hultgren, Isabelle Méjean, Kyle Meng, Ishan Nath, Mathieu Parenti, Roland Rathelot, Ariell Reshef, Ivan Rudik, and Gonzalo Vázquez-Bare, as well as seminar participants at UCSB, PSE, Basel, EAERE-2025, Milan (ECWAAE), Chicago, PSIA, Venice (WEEW), IFS, Wurzburg, CESifo, for their valuable comments. We thank David A Stainforth (LSE) for providing advice on the use of CMIP data. Ollivier acknowledges support from the EUR grant ANR-17-EURE-0001.

Email (URL): jeanne.astier@ensae.fr (<https://jeanneastier.github.io/>),
geoffrey-masters.barrows@polytechnique.edu (<https://sites.google.com/site/geoffreybarrows>),
raphael.calel@georgetown.edu (<https://sites.google.com/view/calel>),
helene.ollivier@psemail.eu (<https://sites.google.com/site/heleneollivier>).

1 Introduction

The world has warmed considerably in recent decades, and this trend is expected to continue for the foreseeable future (IPCC, 2021). Quantifying the economic effects of this warming is essential for setting optimal policy (Nordhaus, 1992), and for understanding the divergent economic paths of nations (Dell et al., 2012).

A large literature quantifies the economic effects of climate change by estimating historical correlations between local weather shocks and local economic outcomes (Deschênes & Greenstone, 2007; Dell et al., 2012; Deryugina & Hsiang, 2014; Burke et al., 2015; Nath et al., 2024). One appeal of this approach is that variation in weather shocks is arguably orthogonal to variation in unobserved shocks to economic fundamentals, suggesting that the identifying exclusion restriction may hold. However, when regions are economically interconnected through trade, a weather shock in one location can affect outcomes in others via competition, input–output linkages, and general equilibrium demand responses. Such across-unit dependencies undermine the stable unit treatment value assumption (SUTVA)—another identifying condition—thereby complicating the interpretation of the reduced-form correlations documented in this literature (Donaldson, 2015).

In this paper, we propose a novel structural estimator to account for SUTVA violations arising from equilibrium effects due to trade. Like previous work in the quantitative trade literature, we impose parametric restrictions on demand and supply in order to discipline across-region dependencies (Eaton & Kortum, 2002; Caliendo & Parro, 2015). However, in contrast to most structural approaches in this literature, we pin down key model parameters using the observed correlations in our data. The result is an estimator that leverages similar data and sources of identifying variation as reduced-form exercises, but whose outputs are nevertheless consistent with a standard trade framework.

Our paper begins with an analysis of reduced-form estimators in a general framework with flexible and heterogeneous spillovers across units, similar to Hudgens & Halloran (2008) and Adão et al. (2024). We consider the canonical two-way fixed effect (TWFE) estimator, along with the estimator from de Chaisemartin et al. (2024) that allows for heterogeneous effects of weather shocks, the estimator from Das et al. (2022); Feng et al. (2025); Zappalà (2024) that controls for weighted average shocks upstream and downstream, and the global time-series estimator from Bilal & Känzig (2024). In the general framework, there is a fundamental identification problem: there are more parameters to estimate than data points. Using our framework, we show the particular restrictions on spillovers that

are implicit in each of the reduced-form estimators. While these restrictions do solve the identification problem, they are not grounded in theory, unlikely to hold if countries trade, and, ultimately, difficult to test empirically.

Focusing on the canonical TWFE estimator, we provide a novel decomposition that helps to interpret the results from the large literature cited above, and to delineate the appropriate uses of this estimator. Specifically, we show that the slope parameter estimated by the TWFE can be decomposed into two terms. The first term is the slope of the best linear approximation to a function relating the *Individual Overall Causal Effect (IOCE)* to own-region observed shock (e.g., temperature) magnitude.¹ The second term is the covariance between a region’s own observed shock and the equilibrium influence of unobserved shocks in *all* regions on the given region’s outcomes. Under certain conditions (e.g., randomly assigned observed shocks), the later term is zero, in which case the TWFE recovers the slope of the best linear approximation of IOCEs, in expectation. This slope can be used to compute *differences in effects* (i.e. effects measured on an interval scale), knowledge of which may be useful in certain contexts for targeting aid or determining equitable tax burdens. Even then, the TWFE is insufficient to quantify the level of effects, or even the relative effects, since the intercept of the best linear approximation is not identified.

To identify the *level* of the individual overall causal effects, we specify a standard multi-sector trade model with roundabout production, perfect competition, and endogenous labor allocation across sectors, as in [Caliendo & Parro \(2015\)](#). We augment the standard framework with a damage function that links climate variables to average productivity in a given country-sector, and thus to equilibrium economic outcomes, similar to [Cruz & Rossi-Hansberg \(2024\)](#), [Nath \(2025\)](#), and [Rudik et al. \(2022\)](#). The model allows for heterogeneous spillover effects between any two regions, with productivity shocks propagating through the world economy via input-output linkages, labor allocations, output market competition, and income effects. Crucially, spillovers channel through a set of observable sufficient statistics—namely, labor allocation and *multilateral resistance indices* ([Anderson & van Wincoop, 2003](#))—, knowledge of which can be leveraged to solve the identification problem.

We show that counterfactuals can be computed with knowledge of just two structural parameters per sector. The first parameter is the elasticity of productivity to climate

¹The best linear approximation to the IOCE is similar to the average slope of potential outcome functions considered in the heterogeneous-effects literature ([de Chaisemartin et al., 2024](#)), but it allows for spillover effects.

variables—the “climate elasticity,” for short. Generating model-consistent estimates of this parameter is at the heart of our contribution. If productivity were observed—either at the macro level or the micro level—then standard econometric techniques could be used to estimate this key elasticity. In practice, however, reliable measures of productivity are difficult to construct for most industries.² We adapt techniques from the International Trade and Economic Geography fields to estimate the climate elasticity without appealing to auxiliary micro data or instruments.

The second structural parameter is the well-known “trade elasticity”—the elasticity of trade flows to trade costs. This parameter can be identified from variation in tariffs or transportation costs, and has been estimated many times already (Caliendo & Parro, 2015; Shapiro, 2016; Boehm et al., 2023). We merge trade and production data for manufacturing and agriculture from Mayer et al. (2023) and Fontagné et al. (2023) with country-by-sector-specific cost share information from GTAP, bilateral trade costs shifters, and bilateral tariff data, and re-estimate sector-specific trade elasticities for our sample, period, and level of aggregation. We also show robustness to adopting parameter estimates from the literature.³

We assess the performance of reduced-form estimators and our structural procedure through Monte Carlo experiments calibrated to observed features of global trade, production, and environmental conditions. Conditional on assumed parameters, we simulate productivity levels and trade costs, and then solve for equilibrium outcomes every year between 1991 and 2019. We then compare the ability of the different estimators to recover the “true” effects of the observed change in weather realizations from 1991 to 2019 in the simulated data sets.

We find that our structural procedure recovers the full vector of IOCEs from finite samples. Across all metrics and countries, the difference between the estimated IOCE and the true IOCE is near zero, on average, with approximately symmetrical errors. By contrast, the reduced-form estimators return systematically biased results. The TWFE estimator tends to *understate* the IOCEs for most countries and sectors. This is because the TWFE estimator tends to recover the slope of the best linear approximation to the function

²Even in the agricultural sector, wherein researchers often equate output per acre with productivity, producers nevertheless endogenously choose the number of planting seasons, crop rotations, and input levels, implying that output per acre is itself an outcome shaped by market conditions (Fuglie, 2024). As a result, regressing farm-level output per acre on weather shocks likely generates biased estimates due to SUTVA violations, just as when output itself is used as the dependent variable.

³We also present a novel GMM procedure that identifies the trade elasticity from the weather shocks themselves. We document the performance of this estimator in Monte Carlo experiments, but rely on the usual gravity regressions in the empirical exercise to leverage additional statistical power.

relating the IOCE to own-country shock magnitude, but not the intercept. Hence, the damages predicted by the TWFE estimator tend to be shifted towards zero. Extensions to the TWFE estimator perform only marginally better. The heterogeneous-robust estimator and the upstream/downstream estimator both understate country-level damages. The global time-series estimator can either understate or overstate the global damages from climate change, depending on the variance of unobserved shocks.

For the empirical application, we use our structural approach to provide new estimates of the effect of climate change on gross output and welfare for each country in the world. We find that, as a result of the warming that occurred from 1991 to 2019, gross national output in 2019 was on average 4.7% lower than it otherwise would have been, and real income was 5.3% lower. The aggregate effects mask substantial heterogeneity though, with some countries losing as much as 28% of aggregate gross output, and as much as 33% of gross agricultural output. The TWFE, heterogeneous-robust, and upstream/downstream estimators all yield lower damage estimates than what we find with our structural approach. This finding are robust to alternative specifications of climate conditions and alternative estimates of trade elasticities.

Combining our estimates with climate projections from the Coupled Model Intercomparison Project Phase 6 (CMIP6) (Eyring et al., 2016), we project that each extra degree of global mean annual temperature lowers global nominal gross output by 4.6%. This “damage elasticity” is higher than values quoted in the literature, which usually range between 1%-3% (Dell et al., 2012; Nath et al., 2024), but also smaller than the value estimated by Bilal & Känzig (2024) of roughly 20%. This ordering is consistent with our Monte Carlo simulations, in which the TWFE, heterogeneous-robust, and upstream/downstream estimators all tend to understate the effects of climate change, while the global time-series estimator overstates the damages for some parameterizations.

Our estimation procedure builds on a large body of work in quantitative trade (Eaton & Kortum, 2002; Hsieh & Ossa, 2016; Caliendo & Parro, 2015; Redding & Venables, 2004; Donaldson & Hornbeck, 2016; Shapiro & Walker, 2018; Anderson et al., 2020; Adão et al., 2020; Bartelme et al., 2024). Most of this work aims to quantify the effects of changes to trade costs or changes to overall productivity levels. We show how the same tools can be used to quantify the effects of observable determinants of productivity—here, climate realizations—, a pursuit of interest even outside the field of International Trade.

Our paper also builds on previous work analyzing TWFE estimators under general conditions. de Chaisemartin et al. (2024) demonstrate that the TWFE may yield biased esti-

mates of average treatment effects when the assumption of treatment effect homogeneity fails, but they assume SUTVA holds. [Borusyak et al. \(2022\)](#) and [Alves et al. \(2024\)](#) study the TWFE estimator in the context of migration models, and [Barrows et al. \(2025\)](#) study the TWFE in the context of a standard industrial organizational framework, both specific contexts in which SUTVA would be expected to fail. In this paper, our analysis of the TWFE does not depend on any particular model structure, beyond static equilibrium.⁴

Finally, our paper contributes to the literature devoted to quantifying the impact of weather shocks on aggregate economic outcomes. While most existing work relies on TWFE estimators (including the papers cited above), a few papers account for across-unit dependencies arising from trade. [Costinot et al. \(2016\)](#) and [Dingel & Meng \(2025\)](#) study climate effects on agricultural output, taking potential or observed yield per acre as a measure of productivity. [Nath \(2025\)](#) includes multiple sectors, but relies on micro data to estimate the climate elasticity. We show how to estimate the elasticity of productivity to climate in a model-consistent way for any sector, even sectors for which quantities are not observed, and without resorting to micro data. [Cruz & Rossi-Hansberg \(2024\)](#) allow for dynamic channels, but require many pieces of auxiliary information to calibrate parameters of the model, while we estimate all parameters needed to compute causal effects from just the data on weather, economic outcomes, and trade data.⁵ Lastly, [Rudik et al. \(2022\)](#) and [Osberghaus & Schenker \(2022\)](#) also use structural gravity frameworks to estimate the effects of weather shocks, but both papers exploit restrictive assumptions with respect to the joint distribution of input prices and weather shocks, which we relax.

The rest of the paper is organized as follows. Section 2 provides an analytical examination of standard estimators in a general setting with spillovers. Section 3 presents our quantitative structural approach. Section 4 describes the data used in the Monte-Carlo experiments (section 5), as well as our empirical application (section 6). Section 7 concludes.

2 Existing Estimators in a General Model Setting

In this section we lay out a general model of how observable and unobservable factors influence endogenous outcomes, allowing for flexible spillover effects. Our primary interest

⁴There is also a broader literature in Urban and Network Economics that studies the identification of treatment effects in settings in which SUTVA fails. Canonical works include [Manski \(1993\)](#), [Sobel \(2006\)](#), [Hudgens & Halloran \(2008\)](#). More recent efforts include [Butts \(2021\)](#), [Leung \(2020\)](#), and [Vazquez-Bare \(2023\)](#). In most settings, researchers exploit sparsity of networks to identify spillover effects. In quantitative trade models, all units can potentially affect all other units, so this strategy is not available to us.

⁵[Cruz & Rossi-Hansberg \(2024\)](#) also assume a single final output sector, symmetric trade costs, and no roundabout production, which are all assumptions we relax.

is to understand conditions under which standard estimators recover unbiased estimates of target values. Finally, we derive a novel characterization of the expected output of the TWFE estimator under general conditions with interference.

2.1 General Model

Consider an economy in which a finite number of units $i = 1, \dots, N$ are endowed in each period $t \in \{0, 1, \dots, T\}$ with a realization of an observable factor z_{it} and an unobservable (to the researcher) factor ε_{it} . Units interact in each period, generating endogenous equilibrium outcomes y_{it} . We assume a static equilibrium in each period and write the mapping between inputs and outputs as

$$\ln y_t = f(z_t, \varepsilon_t), \quad (1)$$

where $\ln y_t \equiv \{\ln y_{1t}, \ln y_{2t}, \dots, \ln y_{Nt}\}$, $z_t \equiv \{z_{1t}, z_{2t}, \dots, z_{Nt}\}$, and $\varepsilon_t \equiv \{\varepsilon_{1t}, \varepsilon_{2t}, \dots, \varepsilon_{Nt}\}$, denote the vectors of the endogenous outcome, the observed factor, and the unobserved factor, respectively.⁶ The supports of z_t and ε_t are bounded subsets of \mathbb{R}^N . The mapping $f(\cdot)$ nests many economic environments and data generating processes (DGP), all of which admit heterogeneous effects of z_{it} and ε_{it} on $\ln y_{it}$, as well as spillover effects of z_{jt} and ε_{jt} on $\ln y_{it}$, for all $i \neq j$.⁷

In our empirical application, i indexes countries, z is a climate variable, ε is unobserved productivity, and y is gross output. More generally, depending on the application, i could index other kinds of geographical units, or even firms or workers. The observable factor z could be policies or operating conditions such as local taxes, infrastructure, or natural disasters, the effects of which we would like to quantify. The unobservable factor ε might include local productivity levels, amenities, or worker skills, while the equilibrium outcome could represent welfare, revenues, or wages. For simplicity, our exposition here considers a single observable factor and a single unobservable factor, though in our empirical application we allow for multiple observables and unobservables.

Our primary object of interest is the causal effect of the change in observables from z_0 to z_T . This effect can be expressed as

$$\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T) \equiv f_i(z_T, \varepsilon_T) - f_i(z_0, \varepsilon_T), \quad (2)$$

where $\Delta \ln y_i^\dagger$ indicates the i th element of the vector of all causal effects $\Delta \ln y^\dagger(\cdot)$, and $f_i(\cdot)$

⁶We take the natural log of the outcome variable because we study growth rates in our application, but the model could just as easily be expressed in levels.

⁷Adão et al. (2024) adopt the same general structure to study equilibrium effects of observable factors, as do many papers in the network econometric literature (Hudgens & Halloran, 2008; Leung, 2020).

denotes the i th element of the vector valued function $f(\cdot)$.⁸ In the language of [Hudgens & Halloran \(2008\)](#), equation (2) defines the *Individual Overall Causal Effect* (IOCE)—the effect on unit i resulting from the change in observables from z_0 to z_T , evaluated at ε_T .⁹

While researchers would ideally like to recover the entire vector of treatment effects, an identification problem arises from the large number of between-regions dependencies. To illustrate the identification problem, we take a Taylor series expansion of (1) around (z_0, ε_0) , assuming $f(\cdot)$ is real-analytic in (z, ε) at (z_0, ε_0) . For each i , we have

$$\begin{aligned} \Delta \ln y_i &= \sum_{a=1}^N \frac{\partial f_i}{\partial z_a} \Big|_{(z_0, \varepsilon_0)} \Delta z_a + \sum_{b=1}^N \frac{\partial f_i}{\partial \varepsilon_b} \Big|_{(z_0, \varepsilon_0)} \Delta \varepsilon_b \\ &+ \sum_{k=2}^{\infty} \sum_{r=0}^k \frac{1}{r!(k-r)!} \sum_{a_1, \dots, a_r=1}^N \sum_{b_1, \dots, b_{k-r}=1}^N \frac{\partial^k f_i}{\partial z_{a_1} \cdots \partial z_{a_r} \partial \varepsilon_{b_1} \cdots \partial \varepsilon_{b_{k-r}}} \Big|_{(z_0, \varepsilon_0)} \prod_{u=1}^r \Delta z_{a_u} \prod_{v=1}^{k-r} \Delta \varepsilon_{b_v}, \end{aligned} \quad (3)$$

The first line on the right hand side of (3) collects the first-order terms while the second line collects higher-order terms. It is convenient to write this expansion in matrix form, denoting the higher-order terms with the vector c .

$$\begin{pmatrix} \Delta \ln y_1 \\ \Delta \ln y_2 \\ \vdots \\ \Delta \ln y_N \end{pmatrix} = \underbrace{\begin{pmatrix} b_{11}^z & b_{12}^z & \cdots & b_{1N}^z \\ b_{21}^z & b_{22}^z & \cdots & b_{2N}^z \\ \vdots & \vdots & \ddots & \vdots \\ b_{N1}^z & b_{N2}^z & \cdots & b_{NN}^z \end{pmatrix}}_{\equiv B^z} \begin{pmatrix} \Delta z_1 \\ \Delta z_2 \\ \vdots \\ \Delta z_N \end{pmatrix} + \underbrace{\begin{pmatrix} b_{11}^\varepsilon & b_{12}^\varepsilon & \cdots & b_{1N}^\varepsilon \\ b_{21}^\varepsilon & b_{22}^\varepsilon & \cdots & b_{2N}^\varepsilon \\ \vdots & \vdots & \ddots & \vdots \\ b_{N1}^\varepsilon & b_{N2}^\varepsilon & \cdots & b_{NN}^\varepsilon \end{pmatrix}}_{\equiv B^\varepsilon} \begin{pmatrix} \Delta \varepsilon_1 \\ \Delta \varepsilon_2 \\ \vdots \\ \Delta \varepsilon_N \end{pmatrix} + \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_N \end{pmatrix} \quad (4)$$

with the elements of the B^z matrix defined as $b_{ij}^z \equiv \frac{\partial f_i}{\partial z_j} \Big|_{(z_0, \varepsilon_0)}$ and the elements of the B^ε matrix defined as $b_{ij}^\varepsilon \equiv \frac{\partial f_i}{\partial \varepsilon_j} \Big|_{(z_0, \varepsilon_0)}$ for all i, j .

This matrix expression shows clearly that the dimensionality of the problem makes it very difficult to recover the causal effect of the observable factors on the endogenous outcomes through purely reduced-form estimation. Even ignoring higher order terms, there are N^2 elements in the B^z matrix, whereas researchers only observe data for N units. It is

⁸Many other counterfactuals could be of interest, such as the one focusing on the change in z between two intermediate periods $t \in \{0, T-1\}$ and $t+1$ or the one substituting the vector of observed shocks at t with an alternative vector z'_t . Here, we restrict attention to the IOCE defined in (2) to match our counterfactual exercise in the application, where we focus on the effect of moving from a baseline climate realization in year $t=0$, to a different, possibly temporally distant climate, as realized in year $t=T$.

⁹We evaluate these effects at end-period vector of unobserved shocks, ε_T , because of how we build our counterfactuals, but one could also evaluate them at ε_0 .

therefore infeasible to estimate the b_{ij}^z coefficients without imposing further structure on the problem.¹⁰

Instead of trying to identify the entire vector of treatment effects, many estimators target the average treatment effect. As a benchmark for these estimators, we define the economy-wide *Best Linear Approximation* (BLA) of IOCEs as

$$(\alpha^{BLA}, \beta^{BLA}) \equiv \arg \min_{a,b} \sum_{i=1}^N \left(\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T) - a - b \Delta z_i \right)^2, \quad (5)$$

Solving (5) yields

$$\beta^{BLA}(z_T, z_0; \varepsilon_T) = \sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T)}{\Delta z_i} \right) \quad (6)$$

$$\alpha^{BLA}(z_T, z_0; \varepsilon_T) = \overline{\Delta \ln y^\dagger} - \overline{\Delta z} \times \beta^{BLA}(z_T, z_0; \varepsilon_T) \quad (7)$$

with $\zeta_i \equiv \frac{\Delta z_i (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})}$ and $\sum_i \zeta_i = 1$, and where $\overline{\Delta \ln y^\dagger}$ and $\overline{\Delta z}$ indicate the average values of $\Delta \ln y_i^\dagger$ and Δz_i across all units.

The BLA estimand extends the notion of an average treatment effect to settings with continuous treatment and spillovers. Specifically, the BLA is the best linear fit between each unit's own observed shock and its IOCE (see Figure 1 for an illustration). The parameter β^{BLA} measures the average difference in the IOCE between units exposed to larger and smaller shocks. It is defined as a weighted average slope of IOCEs. The parameter α^{BLA} measures the average spillover effect, which is conceptually equivalent to the causal effect of the entire vector of observed shocks on units whose own Δz_i is zero.

Under what conditions do existing estimation strategies recover information about the BLA estimand? In other words, what untestable restrictions on the general model must hold so that the estimators tell us about the intercept and slope of the best linear approximation of IOCEs? That is the question we turn to now.

2.2 Some Existing Estimators

Two-way Fixed Effect Estimator One of the most commonly-used estimators in empirical economic research is the TWFE estimator. Given our interest in the long-difference

¹⁰Exploiting variation in the intermediate periods $t \in \{1, T-1\}$ does not necessarily help unless one assumes the b_{ij}^z coefficients are constant over time. Since these coefficients represent partial derivatives evaluated at a specific equilibrium, they likely change over time as the equilibrium changes. Furthermore, one would need very long time series to be able to estimate the b_{ij}^z coefficients non-parametrically, a point raised by Adão et al. (2024).

estimand (2) for our empirical application, the corresponding TWFE econometric model can be written as

$$\Delta \ln y_i = \beta^{FE} \Delta z_i + \alpha^{FE} + \Delta \xi_i^{FE}, \quad (8)$$

where α^{FE} denotes a constant that absorbs common trends and $\Delta \xi_i^{FE}$ denotes the error term.¹¹

The estimator is computed in the usual way

$$\check{\beta}^{FE} = \frac{\sum_i (\Delta \ln y_i - \overline{\Delta \ln y}) (\Delta z_i - \overline{\Delta z})}{\sum_i (\Delta z_i - \overline{\Delta z}) (\Delta z_i - \overline{\Delta z})} = \frac{\sum_i \Delta \ln y_i (\Delta z_i - \overline{\Delta z})}{\sum_i \Delta z_i (\Delta z_i - \overline{\Delta z})} \quad (10)$$

$$\check{\alpha}^{FE} = \overline{\Delta \ln y} - \check{\beta}^{FE} \overline{\Delta z} \quad (11)$$

where $\overline{\Delta \ln y}$ indicates the average value of $\Delta \ln y_i$ in the economy.

The general model represented in equation (4) simplifies to equation (8) if we impose the following restrictions:

1. *No spillover effects*: $b_{ij}^z = 0$ for all $j \neq i$,
2. *Treatment effect homogeneity*: $b_{ii}^z = \beta^{FE}$ for all i , and
3. *No higher-order terms*: $c_i = 0$ for all i .

If these restrictions hold, and we further assume $E[\Delta \varepsilon_i | \Delta z_i] = \kappa$, then we have $E[\check{\beta}^{FE} \Delta z_i] = \beta^{FE} \Delta z_i = \Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T)$. However, it is hard to justify these restrictions (especially assumption 1) if regions trade with each other. Other estimation strategies therefore endeavor to relax some of these identifying assumptions.

Heterogeneous-Robust Estimators de Chaisemartin et al. (2024) develop an estimator for the case that treatment effects are heterogeneous along unobserved dimensions, and the shock is continuously distributed in every period with no “stayer” observations—in our empirical application. de Chaisemartin et al. (2024) identify a weighted average of the slopes of units’ potential-outcome function by comparing the observed changes in

¹¹To increase statistical power, researchers sometimes exploit variation from the intermediate periods $t \in \{1, T-1\}$ and use the following model in first differences (i.e., $\Delta x_{it} \equiv x_{it} - x_{it-1}$ for any variable x):

$$\Delta \ln y_{it} = \beta^{FE,p} \Delta z_{it} + \alpha_t^{FE,p} + \Delta \xi_{it}^{FE,p}, \quad (9)$$

where $\alpha_t^{FE,p}$ denotes time fixed effects and $\Delta \xi_{it}^{FE,p}$ denotes the error term. Even though this specification is, of course, not literally a TWFE model, but a pooled first-difference regression, we will continue to use the term “TWFE” to refer to both the long-difference estimator (8) and the pooled first-difference estimator (9).

outcomes to estimated changes in outcomes under the counterfactual of no treatment.

In our notation, [de Chaisemartin et al. \(2024\)](#) first estimate $E[f(z_0, \varepsilon_T) - f(z_0, \varepsilon_0)|z_0]$, and then compute

$$\check{\beta}^{HR} = \frac{\frac{1}{N} \sum_i \text{sign}(\Delta z_i) \times (\Delta \ln y_i - \check{E}[f(z_0, \varepsilon_T) - f(z_0, \varepsilon_0)|z_0])}{\frac{1}{N} \sum_i |\Delta z_i|}. \quad (12)$$

This estimator requires that the first and third assumptions listed above for the TWFE hold, but relaxes the requirement of treatment effect homogeneity (assumption 2). It is notable that this approach still requires that there are no spillover effects, which is implausible in applications where units trade. See [Appendix B](#) for details of a parametric implementation.

Estimator with Upstream and Downstream Spillovers One approach that allows for spillover effects is to include “upstream” and “downstream” exposure measures to shocks in other units ([Das et al., 2022](#); [Feng et al., 2025](#); [Zappalà, 2024](#)). A typical estimation equation takes the form

$$\Delta \ln y_i = \beta^{Own} \Delta z_i + \beta^{Upstream} \sum_{j \neq i} h_{ij,0} \Delta z_j + \beta^{Downstream} \sum_{k \neq i} h_{ki,0} \Delta z_k + \alpha^{UD} + \xi_i, \quad (13)$$

where $h_{ij,0}$ indicates the upstream exposure of unit i to unit j (j sells to i) measured at time $t = 0$, and $h_{ki,0}$ indicates the downstream exposure of unit i to unit k (k buys from i) measured at time $t = 0$.

Viewed purely as an econometric model, this approach seems quite flexible: researchers can estimate spillovers without imposing structural assumptions on supply and demand. But an implicit assumption behind (13) is that exposure measures are fixed.¹² This assumption is hard to rationalize in a context in which regions trade with each other. Even within a highly stylized quantitative trade model—i.e., the model by [Armington \(1969\)](#) without any input-output linkages—production and consumption shares adjust endogenously with productivity shocks. Hence, we would not expect this parsimonious representation of spillovers to fully account for the types of across-unit dependencies embedded

¹²Alternatively, (13) can be interpreted as a first-order approximation to a structural model, in which case the exposure coefficients can be thought of as *endogenous* sufficient statistics, i.e., not fixed. While such a linear representation always exists, these sufficient statistics may not be expressible as functions of observable metrics, like trade shares. In [Appendix C](#) we discuss a special case based on an amended version of [Acemoglu et al. \(2016\)](#), in which exposure *can* be expressed as functions of observable metrics. But this special case relies on Cobb-Douglas utility and production function in country-sector nodes. For more general production structures, there is no closed-form expression for the exposure measures.

in a data-generating process in which units trade with each other. Whether the regression nonetheless provides a useful approximation when the true data-generating process features endogenous reallocation is an open empirical question.

Without a structural model, there is no theoretical guidance on how to construct the exposure measures in equation (13). Absent theoretical constraints, we follow what appears to be the more common approach of proxying the exposure terms $h_{ij,0}$ and $h_{ki,0}$ with direct trade shares (Das et al., 2022; Feng et al., 2025). This econometric model maps into the B^z matrix with: (i) $b_{ii}^z = \beta^{Own}$ for all i , (ii) $b_{ij}^z = \beta^{Upstream}h_{ij,0} + \beta^{Downstream}h_{ji,0}$ for all $j \neq i$, and (iii) $c_i = 0$ for all i .

Global Spillovers Yet another approach to incorporating spillover effects is to aggregate to a broader geographic level, as in Bilal & Känzig (2024), who analyze globally aggregated data. In our notation, their econometric model is:

$$\tilde{y}_{t+h} - \tilde{y}_t = \alpha_h + \beta_h \tilde{z}_t^{shock} + \sum_{l=1}^L \gamma_{h,l} X_{t-l} + \xi_{t+h} \quad (14)$$

where $\tilde{y}_t \equiv \sum_{i=1}^N y_{it}$ denotes aggregate gross output, $\tilde{y}_{t+h} - \tilde{y}_t$ corresponds to world-wide growth of \tilde{y} between period t and horizon h , \tilde{z}_t^{shock} is a one-year-ahead forecast error of world-average \tilde{z} , β_h is the causal effect of interest at horizon h , and X_{t-l} is a vector of controls up to l lags. Equation (14) is estimated by OLS using global time series data.

While global aggregation means the estimate β_h includes spillover effects to some extent, it does not fully resolve the identification problem. Expressing (1) in levels and computing world-average \tilde{z}_t^{shock} as a weighted sum of country-specific z_{it} (using pre-period population weights $\omega_{i,0}$), we derive (14) from aggregation if and only if each i th column of B^z sum to $\beta_h/\omega_{i,0}$ and $c_i = 0$ for all i . This would mean that two shocks in different parts of the world that have the same effect on the world-average \tilde{z}_t^{shock} would also have the same effect on world-level outcome \tilde{y} , regardless of the location of the shocks. In quantitative trade models, the global effect of a shock depends on a region's position in the network (i.e., its size, centrality, and input-output structure). Equal contributions to global averages therefore do not imply equal equilibrium effects. Hence, if the underlying DGP features discrete regional units that trade with each other, it is unlikely that the global estimator recovers the IOCEs.

One can of course specify yet other econometric models that implicitly populate the B^z matrix with different off-diagonal elements than the ones presented above. However, in lieu

of some theoretical grounding, all of these approaches ultimately rely on *ad hoc* restrictions on the off-diagonal elements that are difficult to test. Quantifying general equilibrium effects of shocks requires *some* structure. Our approach, below, is to derive an estimator that imposes restrictions consistent with a standard trade model.

2.3 Two-Way Fixed Effect Estimator in the General Model

Let us return momentarily to the TWFE estimator. Although the identifying assumptions for equation (8) are unlikely to hold in applications with spillovers, we will demonstrate that the TWFE estimator may still be informative under weaker conditions, when the restrictions presented above fail.

In full generality, we can decompose the observed change in outcomes between period $t = 0$ and T into the individual overall causal effect (IOCE) $\Delta \ln y_i^\dagger$, and residual terms that depend on changes in the unobserved factors:

$$\Delta \ln y_i = \Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T) + f_i(z_0, \varepsilon_T) - f_i(z_0, \varepsilon_0). \quad (15)$$

The residual terms together identify the *individual overall effect of unobserved shocks*, denoted $\Delta \vartheta_i \equiv f_i(z_0, \varepsilon_T) - f_i(z_0, \varepsilon_0)$. Note that, even if the unobserved shocks ε_i were i.i.d., $\Delta \vartheta_i$ would not necessarily be i.i.d., since the latter also depends on equilibrium interactions. If $\Delta \vartheta_i$ were i.i.d., however, we would have a version of the standard unbiasedness result.

Proposition 1. *The slope and constant coefficients from the TWFE model (8) can be written as:*

$$\begin{aligned} \check{\beta}^{FE} &= \beta^{BLA}(z_T, z_0; \varepsilon_T) + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta})(\Delta z_i - \overline{\Delta z})}{\sum_i (\Delta z_i - \overline{\Delta z})^2} \\ \check{\alpha}^{FE} &= \alpha^{BLA}(z_T, z_0; \varepsilon_T) + \overline{\Delta \vartheta} - \overline{\Delta z} \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta})(\Delta z_i - \overline{\Delta z})}{\sum_i (\Delta z_i - \overline{\Delta z})^2}, \end{aligned}$$

so $\mathbb{E}[\check{\beta}^{FE}] = \mathbb{E}[\beta^{BLA}]$ and $\mathbb{E}[\check{\alpha}^{FE}] = \mathbb{E}[\alpha^{BLA}] + \mathbb{E}[\overline{\Delta \vartheta}]$ if and only if $\text{Cov}(\Delta \vartheta, \Delta z) = 0$.

Proof: See Appendix A.1.

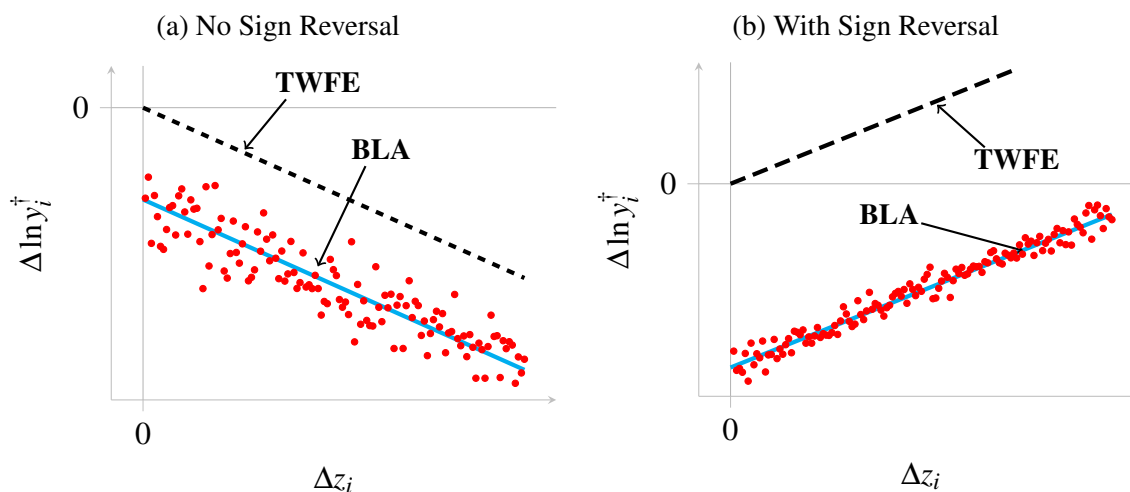
Proposition 1 at once shows the TWFE estimator's surprising power and its clear limitations.¹³ On the one hand, the proposition shows that TWFE estimator can recover meaningful comparisons of treatment effects across units—even in the presence of arbitrary

¹³An analogous proposition for TWFE models estimated on multi-period panel data is included in Appendix A.2.

spillovers—provided that the covariance restriction holds. This is a striking result given the potentially complex ways in which spillovers may propagate across units.

On the other hand, Proposition 1 reveals that the TWFE estimator fails to identify the intercept of the best linear approximation to the IOCEs (see Figure 1). This “missing intercept problem”, as it is sometimes referred to in the literature, has important consequences for quantifying the levels of treatment effects, or indeed relative effects. Without knowing α^{BLA} , it is not possible to pin down the magnitudes, or even the sign of treatment effects. For instance, suppose all IOCEs are negative, with a negative intercept and a positive slope (indicating stronger treatment effects for units receiving small Δz_i). Under the covariance restriction, the TWFE estimator would then recover the positive slope, but without knowing the intercept is negative, one could wrongly conclude that the treatment effects are all positive (as in Figure 1b). This possibility is discussed in [de Chaisemartin et al. \(2024\)](#) as a “sign reversal”, and is generally considered a major drawback of the TWFE estimator.

Figure 1: Illustration of the IOCEs, the BLA, and the TWFE estimates under Proposition 1



Notes: Figure provides an illustration of IOCEs (red dot), the BLA estimand (blue line), and the TWFE estimated treatment effects in expectation (black line), obtained if Proposition 1 holds, for the case of $\beta^{BLA} < 0$ (panel a), and $\beta^{BLA} > 0$ (panel b).

Even if the TWFE estimator fails to identify the *level* of treatment effects, it may still provide useful information about the *difference* in effects between units (i.e. by how many percentage points they differ, in expectation). The critical question is whether or not the covariance restriction holds.¹⁴ Under what conditions is this restriction likely to hold, and

¹⁴The usefulness of Proposition 1 also depends on how well the linear approximation represents the true relationship between IOCEs and own-unit shock magnitude. If the relationship between IOCE and own-unit

when might it fail?

Trivially, if unobserved shocks correlate with observed shocks, then the covariance restriction fails, and the TWFE yields biased estimates of the slope of the BLA. But in such cases, identification is compromised from the outset, since any reduced-form estimator relying on variation in Δz_i would be contaminated by correlated unobservables.

More interestingly, even if $E[\Delta \varepsilon_i | \Delta z_i] = \kappa$, the covariance restriction might still fail. Taking expectations over $\Delta \varepsilon_i$, we can express the restriction in terms of *sensitivity* to the influence of unobserved shocks, rather than the realized influence of unobserved shocks (see Appendix A.1). In this case, we have a simpler orthogonality restriction: units experiencing stronger observed shocks are neither more nor less sensitive to all units' unobserved shocks. This covariance is guaranteed to be zero if the observed shocks Δz_i are i.i.d. When Δz_i are not i.i.d., the correlations between Δz_i and unit characteristics could lead to non-zero covariance.

In the climate change example, we expect that country-level temperature growth correlates with geographical characteristics (e.g., latitude, longitude, altitude, etc.) that might also influence the sensitivity to countries' exogenous productivity shocks. For example, if Δz_i is higher for countries that are far from the equator and geographically isolated, then countries with higher values of Δz_i would also be less sensitive to other countries' exogenous productivity shocks. The reverse could also be true: geographically isolated countries might be *more* sensitive to the influence of all unobserved productivity shocks, since they are more dependent on the supply of close trading partners.

It is difficult to have clear intuitions about the nature of the correlation between Δz_i and countries' sensitivity to unobserved shocks, but these examples give reason to doubt that the condition from Proposition 1 holds when regions trade with each other. In Monte Carlo experiments below, we further discuss the conditions of Proposition 1 and the magnitude of the “missing intercept” problem in a standard trade framework.

In summary, all of the reduced-form estimation strategies we have reviewed impose restrictions on the statistical relationships between inputs and outputs to circumvent an identification problem. We have shown that, under weaker conditions, the TWFE estimator can recover *differences in treatment effects*, at most. To make progress on estimating the *level* of treatment effects, we propose a theory-based solution to the identification problem.

treatment Δz_i is highly nonlinear, the TWFE estimator may deviate significantly from the truth even if it recovers the slope of the best linear approximation.

3 A Structural Approach Based on Quantitative Trade

In this section we use a standard quantitative trade model to evaluate the effect of a change in the vector of observable factors on the outcomes of all units. The advantage of this approach over reduced-form approaches is that the economic assumptions are internally consistent and explicit. Most of the elements of the model are standard, so we leave the majority of the derivations for the Appendix. We explain in detail how to estimate parameters in a model-consistent way with minimal data requirements, which is the core of our contribution.

3.1 Model

The world is divided into N distinct regions, which we refer to as “countries”, and S economic sectors. In each period t , each country n is endowed with a fixed number of worker-consumers L_{nt} that endogenously sort into one of the S sectors. All worker-consumers receive the country-time-specific marginal product of their labor, w_{nt} , for inelastically supplying their one unit of labor. Worker-consumers are assumed to be immobile between countries, but perfectly mobile between sectors, such that the wage varies across countries but not within countries.¹⁵

Worker-consumers from each country n have Cobb-Douglas preferences over the different sectors $s \in \{1, \dots, S\}$, spending a share α_{ns}^C of their income on goods from sector s , with $\sum_s \alpha_{ns}^C = 1$. Within each sector s , worker-consumers in all countries have identical preferences over a continuum of varieties $j \in (0, 1)$, and a constant elasticity of substitution between varieties $\sigma_s > 0$. The resulting Cobb-Douglas price index of country n at time t is: $p_{nt} = \prod_{s=1}^S (p_{nst} / \alpha_{ns}^C)^{\alpha_{ns}^C}$, where p_{nst} is the CES price index of sector s in country n at date t , written as $p_{nst} = \left(\int p_{nst}(j)^{1-\sigma_s} dj \right)^{\frac{1}{1-\sigma_s}}$.

At time t , producers in country i and sector s produce output Q_{ist} using Cobb-Douglas technology with constant returns to scale, requiring labor and intermediate inputs in proportions $\eta_{is} \in [0, 1]$ and $1 - \eta_{is}$, respectively. A proportion α_{ish}^M of sector h 's output is used as an intermediate input of sector s , with $\sum_h \alpha_{ish}^M = 1$. Sector s combines these intermediate inputs using Cobb-Douglas technology.

Consumers (both final good consumers and firms) in each country i in year t source variety j from the lowest cost supplier. Firms engage in perfect competition. With constant

¹⁵As in most trade models, goods are internationally mobile while labor is not. Allowing for international migration is a potentially interesting extension, but we abstract from it because bilateral migration flows are small and migration data are noisy (Parsons et al., 2007).

returns to scale, the cost of producing a unit of good j in sector s of country i at date t is $c_{ist}/v_{ist}(j)$, where $v_{ist}(j)$ denotes country i 's efficiency in producing good j of sector s at date t , and c_{ist} is given by the Cobb-Douglas combination $w_{it}^{\eta_{is}} \prod_{h=1}^S (p_{iht})^{\alpha_{ish}^M (1-\eta_{is})}$.

Firms from sector s in country i selling their goods in country n at date t face iceberg trade costs $\tau_{nist} = \kappa_{nist} (1 + \text{tariff}_{nist})$, where κ_{nist} are non-pecuniary frictions related to distance or lack of common language, for instance, and tariff_{nist} is the *ad valorem* tariff charged to ship goods in sector s from i to n . Hence, delivering a unit of good j from sector s produced in country i to country n at date t costs

$$p_{nist}(j) = \frac{w_{it}^{\eta_{is}} \left(\prod_{h=1}^S p_{iht}^{\alpha_{ish}^M} \right)^{1-\eta_{is}} \tau_{nist}}{v_{ist}(j)}. \quad (16)$$

We assume $v_{ist}(j)$ is a realization of a random variable drawn independently for each country-sector from a specific Fréchet distribution, with CDF $F_{ist}(v) = \exp(-A_{ist} v^{-\theta_s})$ for $v > 0$. The parameter $A_{ist} > 0$ governs the location of the distribution, and $\theta_s > 1$ governs the dispersion.

The country-sector-time-specific productivity parameter A_{ist} is not observed by the researcher, but is modeled as a function of observable factors z_{it}^v , $v \in \{1, \dots, \mathcal{V}\}$, as well as unobserved factors of productivity:

$$A_{ist} = \exp \left(\sum_{v=1}^{\mathcal{V}} \mu_s^v z_{it}^v \right) \exp \left(\psi_{is} + \iota_{st} + \omega_{ist} \right). \quad (17)$$

The observed productivity shifters z_{it}^v include variables like temperature, while the unobserved shifters are grouped into those that are country-sector-specific but time-invariant, ψ_{is} (e.g., latitude, longitude, and institutions), those that are time-varying but common across all countries for a particular sector, ι_{st} (such as technological progress), and those that vary across countries, sectors, and time, ω_{ist} . The set of parameters μ_s^v govern how observed factors affect local productivity for a given sector.¹⁶

Buyers in country n would pay $p_{nist}(j)$ if they chose to buy good j from sector s of country i at date t ; but, under perfect competition, the price buyers in country n actually

¹⁶We abstract from climate effects on population or labor supply and assume that climate influences production solely through productivity. This is consistent with evidence that climate-related mortality primarily affects non-working-age groups (Carleton et al., 2022), implying limited direct consequences for effective labor inputs. Any indirect repercussions on productivity are absorbed in the reduced-form terms in (17).

pay for a good j is the lowest across all sources i . Given the Fréchet distribution, the probability that country i supplies a good of sector s at the lowest price in country n can be written as:

$$\pi_{nist} = \frac{A_{ist} w_{it}^{-\theta_s \eta_{is}} \left(\prod_{h=1}^S p_{iht}^{\alpha_{ish}^M} \right)^{-\theta_s (1-\eta_{is})} \tau_{nist}^{-\theta_s}}{\Phi_{nst}}, \quad (18)$$

with

$$\Phi_{nst} = \sum_{k=1}^N A_{kst} \left(w_{kt}^{\eta_{ks}} \left(\prod_{h=1}^S p_{iht}^{\alpha_{ksh}^M} \right)^{(1-\eta_{ks})} \tau_{nkst} \right)^{-\theta_s}. \quad (19)$$

In the literature, Φ_{nst} is referred to as the “inward multilateral resistance”—a measure of accessibility to suppliers for consumers in n . The probability π_{nist} is also the fraction of sector s ' goods that country n buys from country i at date t , i.e. $\pi_{nist} = X_{nist}/X_{nst}$, where X_{nist} is the value of goods purchased by buyers in n from sellers in i in sector s and time t , and X_{nst} represents total expenditures in n on goods from sector s at time t , which includes final good purchased by household, intermediate inputs purchased by firms, and recycled tariff revenues. Finally, real income per worker is equal to $(w_{nt} + \frac{1}{L_{it}} \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{nkht} X_{nkht}) / p_{nt}$.

Given the distribution of labor across countries, the country-sector-specific productivity draws, the vector of observable factors, and the bilateral sectoral trade costs, imposing market clearing $Y_{ist} = \sum_k \pi_{kist} X_{kst}$ yields a static equilibrium for each period t that satisfies the following equations (see Appendix D.1 for details):

$$Y_{ist} = \exp \left(\sum_{v=1}^{\mathcal{V}} \frac{\mu_s^v}{1 + \theta_s \eta_{is}} z_{it}^v \right) L_{ist}^{\frac{\theta_s \eta_{is}}{1 + \theta_s \eta_{is}}} \Omega_{ist}^{\frac{1}{1 + \theta_s \eta_{is}}} \left(\prod_{h=1}^S \Phi_{iht}^{\alpha_{ish}^M \frac{\theta_s (1-\eta_{is})}{1 + \theta_s \eta_{is}}} \right) \exp(\varepsilon_{is} + \varepsilon_{st} + \varepsilon_{ist}) \quad (20)$$

$$X_{ist} = \sum_{h=1}^S \left(\eta_{ih} \alpha_{is}^C + (1 - \eta_{ih}) \alpha_{ihs}^M \right) Y_{iht} + \alpha_{is}^C \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} X_{ikh} \quad (21)$$

$$\frac{L_{ist}}{L_{it}} = \frac{\eta_{is} Y_{ist}}{\sum_h \eta_{ih} Y_{iht}} \quad (22)$$

$$X_{nist} = \frac{Y_{ist}}{\Omega_{ist}} \frac{X_{nst}}{\Phi_{nst}} \phi_{nist} \quad (23)$$

with

$$\Phi_{nst} = \sum_k \frac{Y_{kst}}{\Omega_{kst}} \phi_{nkst}, \quad \Omega_{nst} = \sum_k \frac{X_{kst}}{\Phi_{kst}} \phi_{knst}. \quad (24)$$

In these equations, Y_{ist} denotes country i 's income from selling goods of sector s at date t , L_{ist} denotes the labor used in country i in sector s at date t , Ω_{ist} denotes the “outward

multilateral resistance” (a measure of market access for producers in country i), and $\phi_{nist} \equiv \tau_{nist}^{-\theta_s}$ represents bilateral accessibility (the inverse of trade frictions). The reduced-form terms ε_{is} , ε_{st} , and ε_{ist} are country-sector, sector-time, and country-time-specific unobserved factors, respectively, each of which is a different combination of ι_{st} , ψ_{is} , ω_{ist} , and structural parameters.¹⁷

From equation (20), we can see that a country-sector’s gross output depends directly on that country’s observed and unobserved determinants of productivity, and through labor allocations and multilateral resistance terms, depends indirectly on the observable and unobservable determinants of productivity in all country-sectors. These across-unit dependencies invalidate the SUTVA, as in the general framework from Section 2.

3.2 Estimation

We are now ready to outline our five-step procedure for estimating all required model parameters and computing counterfactual outcomes. First, we flexibly estimate trade frictions, ϕ_{nist} , and multilateral resistance terms, Φ_{nst} and Ω_{ist} . Second, we project these trade frictions onto tariffs and shifters of bilateral trade cost to estimate θ_s .¹⁸ Third, we plug all of these estimates into (20) and solve for an observable and sufficient proxy for productivity, A_{ist} . Fourth, we project the estimated values of productivity on all observable determinants z^v in a panel structure to estimate μ_s^v —the elasticities of productivity to observable shocks. Finally, we compute counterfactual outcomes for alternative realizations of the observable shocks via exact hat algebra (Dekle et al., 2007).

Step 1: Estimate Trade Frictions and Multilateral Resistance Terms. Following Egger & Nigai (2015), equations (23) and (24) form a system that can be solved up to a normalization without imposing a functional form on the inverse of trade frictions, ϕ_{nist} . We take a slightly different approach to solving this system from Egger & Nigai (2015), which we find to be more computationally stable. Taking logs of (23), we have:

$$\ln X_{nist} = \delta_{ist} + \delta_{nst} + \delta_{nist} \quad (25)$$

¹⁷Specifically, we have $\varepsilon_{is} \equiv \frac{1}{1+\theta_s\eta_{is}} (-\theta_s(1-\eta_{is})\sum_{h=1}^S \alpha_{ish}^M \ln \gamma_h - \theta_s\eta_{is} \ln \eta_{is} + \psi_{is})$, $\varepsilon_{st} \equiv \frac{1}{1+\theta_s\eta_{is}} \iota_{st}$, and $\varepsilon_{ist} \equiv \frac{1}{1+\theta_s\eta_{is}} \omega_{ist}$, with $\gamma_s = [\Gamma(\frac{\theta_s+1-\sigma_s}{\theta_s})]^{-\frac{1}{1-\sigma_s}}$, and $\Gamma(\cdot)$ represents the gamma function.

¹⁸Appendix E outlines a novel GMM procedure that accomplishes the same task without using tariff data. The cost of using this approach is that it requires greater variation in observed productivity shifters in order to identify all of the parameters from less data. We do not have sufficient variation in recent climate records to utilize this GMM procedure in our application, but we choose to report it in the appendix for researchers to use for other suitable applications.

where trade frictions are in the residual, $\ln \delta_{nist} = \phi_{nist}$, and where country-sector-time fixed effects absorb equilibrium outcomes at the origin and destination, $\delta_{ist} = \ln \left(\frac{Y_{ist}}{\Omega_{ist}} \right)$ and $\delta_{nst} = \ln \left(\frac{X_{nst}}{\Phi_{nst}} \right)$. Estimating (25) by OLS and extracting the residual, we can build an estimate of normalized trade frictions, $\check{\phi}_{nist}$:

$$\check{\phi}_{nist} \equiv \frac{\exp(\check{\delta}_{nist})}{\exp(\check{\delta}_{nnst})} = \frac{\phi_{nist}}{\phi_{nnst}} \times \frac{\phi_{nst}^e}{\phi_{ist}^e} \quad (26)$$

where $\check{\delta}_{nist}$ is the estimated residual for trade flows from i to n , normalized by self-trade-flow estimated residual. This implies a normalization using average trade frictions: $\phi_{nst}^e \equiv \exp \left(\frac{1}{N} \sum_k \ln \phi_{knst} \right)$ and $\phi_{ist}^e \equiv \exp \left(\frac{1}{N} \sum_k \ln \phi_{kist} \right)$. For any trade flow with $X_{nist} = 0$, we set $\check{\phi}_{nist} = 0$.¹⁹

Next, we solve the system defined in (24), using $\check{\phi}_{nist}$ in place of ϕ_{nist} :

$$m_{nst} = \sum_k \frac{Y_{kst}}{e_{kst}} \check{\phi}_{nkst} \quad , \quad e_{ist} = \sum_k \frac{X_{kst}}{m_{kst}} \check{\phi}_{kist} \quad (27)$$

and normalize by m_{1st} . We can solve for m_{nst} and e_{ist} for all countries by fixed point iteration. Given the normalizations, estimated multilateral resistance terms relate to true multilateral resistance terms in the following way:

$$\check{\Phi}_{nst} \equiv m_{nst} = \frac{\Phi_{nst} \phi_{nst}^e \phi_{11st}}{\Phi_{1st} \phi_{1st}^e \phi_{nnst}} \quad , \quad \check{\Omega}_{ist} \equiv e_{ist} = \frac{\Omega_{ist}}{\phi_{ist}^e \phi_{11st}} \Phi_{1st} \phi_{1st}^e \quad (28)$$

The estimated values $\check{\phi}_{nist}$, $\check{\Phi}_{nst}$, and $\check{\Omega}_{ist}$ thus solve the system (23) and (24).²⁰ As we show below, the normalization terms do not influence the computation of counterfactuals.

Step 2: Estimating θ from Gravity. We can now use the estimated trade frictions, $\check{\phi}_{nist}$, to recover the country-sector specific shape parameter of the Fréchet distributions, θ_s . Following the gravity literature, we model the trade frictions as a product of tariffs and non-pecuniary trade costs, κ_{nist} , with the latter parameterized as an exponential function of log distance, crossing a border, being contiguous, sharing a common language, and sharing a

¹⁹This assumption rationalizes zero trade flows with near-infinite trade costs.

²⁰First, to show that (23) holds, compute for any sector-period:

$$\hat{X}_{ni} = \frac{Y_i}{\hat{\Omega}_i} \frac{X_n}{\hat{\Phi}_n} \hat{\phi}_{ni} = \left(\frac{Y_i}{\Omega_i} \frac{\phi_{11}}{\Phi_1} \frac{\phi_i^e}{\phi_i^e} \right) \left(\frac{X_n}{\Phi_n} \Phi_1 \frac{\phi_1^e}{\phi_n^e} \frac{\phi_{nn}}{\phi_{11}} \right) \left(\frac{\phi_{ni}}{\phi_{nn}} \frac{\phi_n^e}{\phi_i^e} \right) = \frac{Y_i}{\Omega_i} \frac{X_n}{\Phi_n} \phi_{ni} = X_{ni}.$$

Second, since \hat{X}_{ni} coincides cell-by-cell with X_{ni} , the adding-up constraints follow immediately.

colonial link. Imposing this functional form on equation (28) and taking logs, we have

$$\begin{aligned} \ln \check{\phi}_{nist} &= \tilde{\delta}_{ist} + \tilde{\delta}_{nst} - \theta_s v_s^{dist} \ln dist_{ni} - \theta_s v_s^{Contig} Contig_{ni} - \theta_s v_s^{Border} Border_{ni} \\ &- \theta_s v_s^{ComLang} ComLang_{ni} - \theta_s v_s^{Colonial} Colonial_{ni} - \theta_s \ln(1 + tariff_{nist}) + \rho_{nist} \end{aligned} \quad (29)$$

where country-sector-time fixed effects $\tilde{\delta}_{ist}$ and $\tilde{\delta}_{nst}$ absorb the unobserved terms $\phi_{nist}^e / \phi_{nnst}$ and ϕ_{ist}^e , respectively, and ρ_{nist} represents unobserved components of trade costs. In this context, the trade elasticity θ_s is equivalent to the shape parameter of the country-sector-specific distribution of production efficiencies, $F_{ist}(v)$, encountered in equation (16).

Equation (29) can be estimated by OLS to recover the parameters associated to each bilateral cost shifter. The trade elasticity, θ_s , cannot be recovered from the coefficients on the non-tariff gravity variables, since these coefficients combine θ_s with the elasticities of trade costs to each of those gravity variables, the v_s 's. However, since *ad valorem* tariffs increase costs by a multiplicative factor, the elasticity of log trade frictions to $\ln(1 + tariff_{nist})$ directly identifies θ_s . In practice, we use lagged tariffs as an instrument for current levels to address measurement error.

Two points about this procedure are worth mentioning. First, the estimated coefficients from equation (29) are identical to those from a corresponding regression that puts trade flows directly on the left hand side. So why not simply estimate equation (29) with trade flows on the left hand side and construct the trade frictions, ϕ_{nist} , from the gravity variables and the estimated parameters? Many trade applications follow exactly this recipe. The benefit of this alternative would be that we could skip Step 1 of our estimation procedure. The downside is that this alternative approach can only recover estimates of trade frictions for bilateral pairs with non-missing gravity variables; and in practice, bilateral tariff data are often missing.²¹ By contrast, we recover a complete matrix of trade frictions, as well as estimates of multilateral resistance terms that are model consistent, by estimating trade frictions separately in Step 1.²²

Second, [Silva & Tenreyro \(2006\)](#) argue that log-linear gravity regressions can generate inconsistent estimates when heteroskedasticity induces dependence between the multiplicative error term and the regressors. This critique does not compromise identification

²¹The International Trade Center in Geneva has made extensive effort to address incomplete coverage and misreporting of tariff rates with the MAcMap dataset, but MAcMap still only covers a few years: 2001, 2004, 2007, 2010, 2016, and 2019 ([Fontagné et al., 2022](#)). Outside of these years, bilateral tariff rates are missing for a high share of country-good cells, even aggregated to the sector level.

²²Alternatively, we could skip estimating (29) entirely, if we simply adopt estimates of θ_s from the literature, of which there are many ([Shapiro, 2016](#); [Fontagné et al., 2022](#); [Boehm et al., 2023](#)).

of θ_s when estimating equation (29) by OLS because the error term ρ_{nist} has a structural interpretation as an unobserved multiplicative trade cost. Under the assumption that ρ_{nist} is independent of the observable gravity variables, the regression error satisfies the separability condition required for consistency of the log-linear estimator.²³ Moreover, Egger & Nigai (2015) show that the Pseudo-Poisson Maximum Likelihood (PPML) estimator recommended by Silva & Tenreyro (2006) is inconsistent in the presence of unobserved trade costs. Our log-linear specification therefore provides a consistent alternative when there are unobserved components of trade cost.

Step 3: Solving for a productivity proxy. Having estimated Φ_{ist} , Ω_{ist} , and θ_s , we combine these estimates with country-sector labor production shares, η_{is} , labor allocations, L_{ist} , expenditure shares, α_{ish}^M , and consumption shares, α_{is}^C , which we assume are available from input-output tables. Taking logs and moving these observed components to the left hand side of equation (20), we get:

$$\begin{aligned}
(1 + \theta_s \eta_{i,s}) \ln Y_{ist} & - \theta_s \eta_{i,s} \ln L_{ist} - \theta_s (1 - \eta_{i,s}) \sum_{h=1}^S \frac{\alpha_{ish}^M}{\theta_h} \ln \check{\Phi}_{iht} - \ln \check{\Omega}_{ist} \\
& = \sum_{v=1}^{\mathcal{V}} \mu_s^v z_{it}^v + \psi_{is} + \iota_{st} + \omega_{ist} + \left(\ln (\Phi_{1st} \phi_{1st}^e / \phi_{11st}) - \ln \phi_{ist}^e \right) \\
& + \theta_s (1 - \eta_{i,s}) \sum_h \frac{\alpha_{ish}^M}{\theta_h} \left(\ln \frac{\phi_{iht}^e}{\phi_{iist}} - \ln \left(\frac{\Phi_{1ht} \phi_{1ht}^e}{\phi_{11st}} \right) \right) \tag{30}
\end{aligned}$$

The left hand side of this equation is an adjusted output measure, leveraging the model structure to subtract off all the general equilibrium objects from gross output. We will refer to this object later simply as ξ_{ist} . Meanwhile, notice that the right hand side of this equation collects the log of unobserved productivity A_{ist} (as defined in equation 17), as well as components related to unobserved normalization terms from Step 1. Equation (30) thus provides an equivalence between an observed quantity, ξ_{ist} , and another object that, for the purpose of estimating μ_s , turns out to be an adequate proxy of unobserved productivity.

²³ Silva & Tenreyro (2006) consider the model $X_{ni} = \exp(G_{ni}\beta) + \varepsilon_{ni}$ where G_{ni} is a vector of gravity variables and ε_{ni} is an error term that rationalizes imperfect fit between X_{ni} and $\exp(G_{ni}\beta)$. In this case, the log-linear representation is valid only if $\varepsilon_{ni} = \exp(G_{ni}\beta)\zeta_{ni}$ with ζ_{ni} representing a variable independent of G_{ni} . Hence, assuming that $\zeta_{nist} = \exp(\rho_{nist}) - 1$ satisfy the condition required for consistency of the log-linear estimator in Silva & Tenreyro (2006) because ρ_{nist} , and thus ζ_{nist} , are independent of G_{ni} .

Step 4: Estimating μ_s . Taking the first difference of equation (30) yields our estimation equation:

$$\Delta \ln \xi_{ist} = \delta_{is} + \delta_{st} + \sum_{v=1}^{\mathcal{V}} \delta_s^v \Delta z_{it}^v + \Delta \ln \tilde{\xi}_{ist} \quad (31)$$

where $\Delta \ln \xi_{ist}$ is the first difference of ξ_{ist} . This term is essentially an adjusted measure of output growth, subtracting off all the general equilibrium effects. The terms that remain on the right hand side of equation (31) are all exogenous—observed and unobserved components of productivity, as well as components related to unobserved normalization terms from Step 1 for the residual: $\Delta \ln \tilde{\xi}_{ist} \equiv \Delta \omega_{ist} - \Delta \ln \phi_{ist}^e + \sum_h \theta_s (1 - \eta_{i,s}) \alpha_{ish}^M \left(\Delta \ln \left(\frac{\phi_{iht}^e}{\phi_{iist}^e} \right) \right) - \sum_h \theta_s (1 - \eta_{i,s}) \alpha_{ish}^M \left(\Delta \ln \left(\frac{\Phi_{1ht} \phi_{1ht}^e}{\Phi_{11ht}} \right) \right)$.

If we could assume that Δz_{it}^v were i.i.d. random variables, then $E \left[\Delta \ln \tilde{\xi}_{ist} \mid \Delta z_{it}^v \right] = 0$, and hence $E \left[\check{\delta}_s^v \right] = \mu_s^v$, where $\check{\delta}_s^v$ indicates parameters estimated by OLS. However, when Δz_{it}^v represents climate variables, two complications arise.

First, Δz_{it}^v may correlate with unobserved measures of export frictions $\Delta \ln \phi_{ist}^e$. For example, common trends in transportation costs may differentially lower bilateral transportation costs for places that are more centrally located (Borusyak & Hull, 2023). Centrally located countries tend to be located away from the equator, where climate may evolve differently from countries closer to the equator. This type of correlation would lead to $E \left[\Delta \ln \tilde{\xi}_{ist} \mid \Delta z_{it}^v \right] \neq 0$, and hence $E \left[\check{\delta}_s^v \right] \neq \mu_s^v$.

Second, some climate variables, such as the growth in the number of extreme heat days, likely correlate with baseline climate, which may correlate with unobserved trends in productivity $\Delta \omega_{ist}$, as argued by Jones et al. (2025). For example, countries that are hotter on average *in levels* may have slower productivity growth because heat lowers R&D efficiency, and may also experience a larger increase in extreme heat days due to climate change. As above, this leads to bias in OLS estimation of (31).

Both sources of bias can be addressed by controlling for *expected* values of Δz_{it}^v . We follow the procedure from Jones et al. (2025) for computing precisely these expectations based on historical values of z_{it}^v and control linearly for these measures in (31). We also include these controls for all alternative estimators (see Appendix H). In this case, identification relies on the difference between the realized and the expected Δz_{it}^v being orthogonal to unobserved productivity shocks and growth in average export frictions. We also build controls of $(1 - \eta_{i,s}) \alpha_{ish}^M$ interacted with sector-year time fixed effects to control for the last

term in the definition of $\Delta \ln \tilde{\xi}_{ist}$.

Note that, had we simply regressed shocks on output growth, SUTVA would be violated and the results would be biased. We can estimate the $\check{\delta}_s^v$'s with OLS only because we have already stripped out all the general equilibrium effects from the dependent variable. As long as weather shocks in one region do not affect *productivity* elsewhere, we have $E \left[\check{\delta}_s^v \right] = \mu_s^v$.

Step 5: Computing Counterfactuals The last step is to compute outcomes under some counterfactual vector of weather shocks. For a generic outcome g_t , we can solve for the counterfactual g'_t for any z'_t . For the ratio of counterfactual and observed outcomes, we use the notation $\hat{g}_t \equiv \frac{g'_t}{g_t}$.

Armed with estimates of multilateral resistance terms, trade frictions, and structural parameters, we solve the system below for $\hat{\Phi}_{nst}$, $\hat{\Omega}_{nst}$, \hat{Y}_{nst} , \hat{L}_{nst} , and \hat{X}_{nst} , by fixed point iteration:

$$\hat{\Phi}_{nst} = \frac{\sum_k \frac{\hat{Y}_{kst} Y_{kst}}{\hat{\Omega}_{kst} \check{\Omega}_{kst}} \check{\Phi}_{nkst}}{\sum_k \frac{Y_{kst}}{\check{\Omega}_{kst}} \check{\Phi}_{nkst}}, \quad \hat{\Omega}_{nst} = \frac{\sum_k \frac{\hat{X}_{kst} X_{kst}}{\hat{\Phi}_{kst} \check{\Phi}_{kst}} \check{\Phi}_{knst}}{\sum_k \frac{X_{kst}}{\check{\Phi}_{kst}} \check{\Phi}_{knst}} \quad (32)$$

$$\hat{Y}_{ist} = \exp \left(\sum_{v=1}^{\mathcal{V}} \frac{\mu_s^v}{1 + \theta_s \eta_{is}} (z_{it}^{v'} - z_{it}^v) \right) \hat{L}_{ist}^{\frac{\theta_s \eta_{is}}{1 + \theta_s \eta_{is}}} \hat{\Omega}_{ist}^{\frac{1}{1 + \theta_s \eta_{is}}} \left(\prod_{h=1}^S \hat{\Phi}_{iht}^{\alpha_{ish}^M \frac{\theta_s (1 - \eta_{is})}{1 + \theta_s \eta_{is}}} \right) \quad (33)$$

$$\hat{X}_{ist} X_{ist} = \sum_h (\alpha_{is}^C \eta_{ih} + (1 - \eta_{ih}) \alpha_{ihs}^M) \hat{Y}_{iht} Y_{iht} + \alpha_{is}^C \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} \hat{X}_{ikh} X_{ikh} \quad (34)$$

$$\hat{L}_{ist} L_{ist} = L_{it} \frac{\eta_{is} \hat{Y}_{ist} Y_{ist}}{\sum_h \eta_{ih} \hat{Y}_{iht} Y_{iht}} \quad (35)$$

As mentioned above, the unobserved normalization terms cancel out in this system, so at the true parameter values θ_s and μ_s^v , this system recovers the true vector of counterfactual outcomes (see appendix D.2).

With the solution to this system in hand, the proportional change in the real income is computed, without loss of generality, as²⁴

$$\hat{R}I_{it} = \frac{\hat{w}_{it} + \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} \hat{X}_{ikh} X_{ikh}}{\hat{p}_{it}} = \frac{\hat{Y}_{i1t} / \hat{L}_{i1t} + \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} \hat{X}_{ikh} X_{ikh}}{\prod_{h=1}^S \hat{\Phi}_{iht}^{-\alpha_{ih}^C / \theta_h}}. \quad (36)$$

²⁴Since the wage is equalized across sectors, we have $\hat{w}_{it} = \hat{Y}_{ist} / \hat{L}_{ist}$ for any s .

3.3 Comparisons to the Literature

Since there are many alternative procedures for estimating μ_s^v , it is worth comparing our proposed multi-step procedure to related efforts.

First, [Rudik et al. \(2022\)](#) estimate μ_s^v directly from a gravity regression like equation (25), where bilateral trade flows are divided by self-trade flows. The exclusion restrictions are thus that the pairwise differences between countries’ weather shocks, wages, observable trade barriers, are each orthogonal to the residuals, which absorb the pairwise differences in intermediate input prices, unobservable trade barriers, and unobserved productivity shocks.²⁵ If wages and intermediate input prices are jointly determined based on unobserved productivity shocks and weather shocks, as they are in our model, there is necessarily correlation between the residual and the explanatory variables. In this case, identifying μ_s^v would require additional means of purging endogeneity, such as instrumental variables. This is why we develop a strategy based on estimating equation (31) instead.

Second, researchers in International Trade typically estimate the effects of “monadic” cost shifters—i.e., factors that shift the cost of serving all locations—via a two-step procedure ([Eaton & Kortum, 2002](#); [Head & Ries, 2008](#)). With this approach, researchers first estimate gravity regressions like equation (25) by parameterizing δ_{nist} , then extract the fixed effects and project them onto explanatory variables. This procedure can also identify μ_s^v in principle, but requires data on all components of costs, such as wages and input prices, as well as instruments for these measures, as in the approach from [Rudik et al. \(2022\)](#). The procedure we propose circumvents these data requirements.

Finally, [Nath \(2025\)](#) estimates μ_s^v from firm-level regressions, exploiting variation in temperature across regions within a given country. While micro data may lead to higher precision in estimated coefficients because of lower measurement error, micro data does not solve the identification problem posed by violations of SUTVA. Firms compete within a country as well, so climate-induced productivity shocks in one region likely affect the

²⁵To illustrate, [Rudik et al. \(2022\)](#) divide bilateral trade flows by self-trade flows to derive (using our notation)

$$\ln \left(\frac{X_{nist}}{X_{nnt}} \right) = \sum_{v=1}^{\mathcal{V}} \mu_s^v (z_{it}^v - z_{nt}^v) + (\psi_{is} - \psi_{ns}) + (\omega_{ist} - \omega_{nst}) - \theta_s \ln \left(\frac{w_{it}^{\eta_{is}}}{w_{nt}^{\eta_{ns}}} \right) - \theta_s \ln \left(\frac{\sum_{h=1}^S P_{iht}^{\alpha_{ish}^M (1-\eta_{is})}}{\sum_{h=1}^S P_{nht}^{\alpha_{nsh}^M (1-\eta_{ns})}} \right) + \phi_{nist}.$$

[Rudik et al. \(2022\)](#) parameterize ϕ_{nist} in the usual way and then estimate μ_s^v s directly from the equation above via PPML, controlling for relative wages and observable trade barriers. Lacking data on intermediate input prices, unobserved productivity shocks, and unobservable trade barriers, these variables are effectively absorbed by the error term.

outcomes of firms in other regions within the country. As a result, TWFE estimates based on micro data likely suffer from the same bias as TWFE at the macro level (Barrows et al., 2025). The bias in micro data could conceivably be even more severe, as trade within countries is usually more extensive than trade between countries.

4 Data

This section describes the data that will be used in the empirical application and to guide our Monte Carlo experiments. Further details are relegated to Appendix F and descriptive statistics are presented in Table F.5.

Trade and Production Data. We use data from Mayer et al. (2023) and Fontagné et al. (2023) to measure gross output and bilateral trade flows in the agriculture and manufacturing sectors at the country-level over the period 1991-2019. This dataset combines bilateral trade data from the UN Commodity Trade Statistics Database (COMTRADE) and the Food and Agriculture Organization of the United Nations Statistics Division (FAOSTAT), with production data from the UNIDO Industrial Statistics database INDSTAT and FAOSTAT, thereby recording self-trade. We aggregate the underlying manufacturing data, available at the 6-digit level, into a single manufacturing sector, and likewise aggregate the 4-digit level agriculture data into one agricultural sector.

Weather Data. To measure the observable determinants z_{it}^y , we use temperature data from the global reanalysis ERA-5 dataset compiled by the European Centre for Medium-Range Weather Forecasts (Hersbach et al., 2023). We start with daily data from 1991 to 2019, available for a $0.25^\circ \times 0.25^\circ$ resolution grid (corresponding to cells of 30km x 30km at the equator). For each grid cell, we compute the number of days in the year where maximum daily temperature falls within a given temperature bin. This specification imposes minimal restrictions on the response functional form (Schlenker & Roberts, 2006; Deschênes & Greenstone, 2011), allowing us to capture any nonlinear effects of weather on productivity. In order not to over-tax the variation in the output data, we compute just three temperature bins: $(-\infty; 0^\circ\text{C}]$, $(0^\circ\text{C}; 30^\circ\text{C}]$, $(30^\circ\text{C}; +\infty)$.²⁶ Having aggregated from daily to annual data for each grid cell, we then aggregate from cells to country level, as recommended by Hsiang (2016), using population weights of year 2000.²⁷

²⁶In robustness checks, we also consider the annual average of the daily mean temperature, which can be interacted with the initial-period average temperature to allow for heterogeneous effects.

²⁷We use population weights of year 2000 from the Socioeconomic Data and Application Center's UN WPP-Adjusted Gridded Population of the World dataset (CIESIN, 2018).

Gravity Data. The gravity variables required to estimate equation (29) come from the Gravity database developed by CEPII (Conte et al., 2022). Each observation corresponds to an importer-exporter pair in a particular year. Data span from 1948 to 2019, and include 252 countries (with a history of past territorial configurations of countries). We consider the usual characteristics, such as geographical distance between a pair of countries, whether they are contiguous, whether they share a common language, and whether they were in a colonial relationship or shared a common colonizer.

Tariff Data. We use bilateral tariff information from the MAcMap-HS6 dataset (Fontagné et al., 2022), which provides applied *ad valorem* equivalent tariffs (including non-tariff barriers) at the HS6 level for importer-exporter pairs in benchmark years 2001, 2004, 2007, 2010, 2013, 2016 and 2019. We aggregate the HS6-level tariffs to one manufacturing sector and one agricultural sector, using trade expenditures as weights.²⁸

GTAP Data. We compute production and consumption shares from the Global Trade Analysis Project (GTAP) Data Base version 9 (Aguilar et al., 2016). We consider domestic and import expenditures at purchaser’s price for the reference year 2011. The database contains information on 57 commodities for 116 countries and 24 aggregate regions, which we aggregate into three broad sectors: agriculture, manufacturing, and other, while attributing values of the aggregate regions to their individual constitutive countries. We compute the consumption shares α_{is}^C by dividing country i ’s household consumption in each sector s by total consumption. The expenditure shares α_{ish}^M are computed by dividing input expenditures in each sector h coming from sector s in country i by total expenditures. The labor production shares η_{is} are computed as the ratio of labor expenses by firms in country i and sector s over the sum of their total variable input expenditures.

Labor Data. Total employment is calculated as the sum of International Labor Organization modeled estimates (ILOEST) of agriculture and manufacturing employment for each country and year. These data cover 274 countries from 1991 to 2023.

Final Dataset. The merged dataset contains annual data on 132 countries from 1991 to 2019, covering trade, production, and consumption for both the agricultural and manufacturing sectors, as well as total employed labor and weather variables. Table F.5 provides descriptive statistics on the main variables in our dataset.

²⁸To compute counterfactual outcomes for a given year (in our case, 2019), we need to know the tariff rates for all bilateral pairs and all goods in that year. We therefore complement MAcMap with a procedure described in Appendix F for that year. Since imputed tariff rates are only used to compute counterfactuals, we run a robustness check where we assume that tariff revenues are not rebated to households and are thus a pure deadweight loss, and this leaves our results virtually unchanged.

The average country experienced a 9% increase in population-weighted average temperature from 1991 to 2019. This shift in the temperature distribution is also captured by the increase in the number of days where the maximum temperature is above 30°C (+58% for the average country) and the reduction in days where the maximum temperature is below 0°C (-86%). There is a large amount of variation across countries, though, which is what we will rely on to estimate key parameters. See appendix F for further discussion.

5 Monte Carlo Experiments

In this section we perform Monte Carlo experiments to assess the finite-sample performance of the TWFE estimator, extensions thereof, and of our structural procedure. The model from section 3 serves as our data generating process, and the data from section 4 guides our calibration so that our experiments resemble the application in section 6.

5.1 Set Up

Our simulations take as given the set of 132 countries in our merged data, and their endowments of geography, population, production and consumption parameters, as described in section 4. We compute bilateral trade costs using values of $\theta_{ag} = 1.5$ and $\theta_{manuf} = 4.6$, estimated using our gravity equation (29).²⁹ We set $\sigma_{ag} = 2$ and $\sigma_{manuf} = 3$, so that $\theta_s > \sigma_s - 1$.³⁰

For simplicity, we consider a single continuous treatment variable—the number of days in the year with maximum temperature above 30°C.³¹ We set ψ_{is} proportional to baseline temperature and construct ω_{ist} to follow a random walk process, with $\Delta\omega_{ist} \sim N(0, 0.3)$. We then construct country-sector-year productivity A_{ist} using equation (17), assuming values for the climate elasticities of $\mu_{ag}^{30^\circ\text{C}} = -0.002$ and $\mu_{manuf}^{30^\circ\text{C}} = -0.0005$.

We draw unobserved trade costs $\rho_{nist} \sim N(0, 0.02)$ and build $\tau_{nist} = \kappa_{nist} (1 + \text{tariff}_{nist}) \rho_{nist}$, with $\kappa_{nist} = e^{\left(v_s^{dist} \ln dist_{ni} + v_s^{Contig} Contig_{ni} + v_s^{Border} Border_{ni} + v_s^{ComLang} ComLang_{ni} + v_s^{Colonial} Colonial_{ni}\right)}$ with

²⁹See Appendix G for details on the estimation. Our estimates are higher than estimates from Boehm et al. (2023) ($\theta_{ag} = 0.8$, $\theta_{manuf} = 1.1$), but lower than estimates from Shapiro (2016) ($\theta_{ag} = 3.3$, $\theta_{manuf} = 8.5$) and Fontagné et al. (2022) ($\theta_{ag} = 6.9$, $\theta_{manuf} = 8.3$). In the empirical application, we test the sensitivity of our results to these alternative values from the literature, but treat the estimates from Table G.6 as our preferred specification.

³⁰These values are not far from the average estimates of σ across 3-digit industries computed by Broda & Weinstein (2006), which are 3.59 in manufacturing (SITC Rev3 3-digit codes 231 - 971) and 3.62 in the agricultural sector (SITC Rev3 3-digit codes 001 - 223), excluding one outlier industry (SITC 017, MEAT AND EDIBLE MEAT OFFAL, PREPARED OR PRESERVED N.E.S.).

³¹This is properly understood as a low-dimensional summary statistic representing whatever aspects of climate are important for economic activity. It commits us only to the view that number of hot days is correlated with the those features. Later on, we consider alternative specifications.

parameters v_s^{dist} , v_s^{Contig} , v_s^{Border} , $v_s^{ComLang}$, $v_s^{Colonial}$ constructed by dividing parameter estimates in Table G.6 by the assumed θ_s .³²

Given the constructed values of A_{ist} and τ_{nist} , along with the assumed values for σ_s , θ_s , α^M , α^C , we solve the system of equations (20) - (24) in levels, period by period, for $\{X_{st}\}, \{Y_{st}\}, \{\Phi_{st}\}, \{\Omega_{st}\}, \{L_{st}\}$, by fixed point iteration. This process yields many simulated datasets, or “replications”, each approximating the true data. We simulate equilibrium levels each year from 1991 until 2019 using observed temperature realizations, and then simulate the counterfactual equilibrium in the last year of the sample imposing the treatment vector from the first year of the sample, $z'_{i,2019} = z_{i,1991}$. Taking the difference between these two equilibria yields the IOCEs for each replication.

We implement the TWFE estimator on each of our simulated data sets, as well as the heterogeneous-robust estimator, the upstream/downstream estimator, the global time-series estimator, and our structural estimator. Each time we compare the estimated effects to the “true” IOCEs to assess finite-sample performances of these estimators.

5.2 Simulation Results

Figure 2 summarizes the results for total gross output from 200 simulated replications. The black dots, which are repeated in each subplot, show the median IOCE for each country across replications.³³ It shows that changing the climate treatment vector from 1991 to 2019 values induces a reduction in total gross output in all countries, ranging from -2% to -20%. Breaking this down by sector, the output losses are proportionately greater in agriculture than in manufacturing, though the manufacturing losses have a greater effect on total gross output due to the sectors’ relative size (see Appendix Figure I.1).

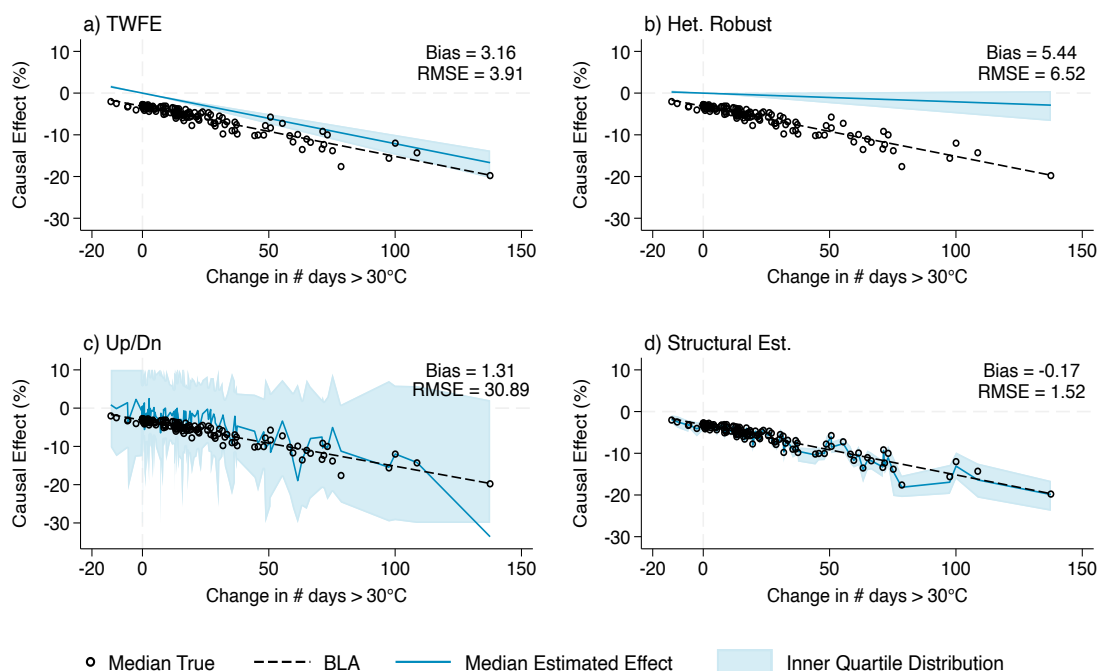
The fact that output changes even in countries with zero treatment is indicative of spillover effects. Moreover, the fact that output falls across the board, even in countries with *negative* treatment (i.e., reduction in the number of extreme heat days) suggests that spatial linkages tend to amplify treatment effects in this setting, as in Adão et al. (2020). Hence, when higher temperatures harm productivity in any one country, they also tend to reduce output in other countries.³⁴

³²After deflating the parameter estimates in Appendix Table G.6 by the assumed θ_s , some of the constructed κ_{nist} are very large, which impedes convergence during the simulation of the equilibrium level of outcomes. To aid convergence, we deflate v_s^{dist} by 40% for these simulations.

³³Although we hold the temperature profiles, total labor, production structure, and geography, constant across the replications, there is nevertheless some variation in the true IOCEs across replications (see Appendix Figure I.1). Hence, we plot the median IOCE across replications to visually represent the true IOCE.

³⁴This pattern does not hold for all countries and all outcomes. Agricultural output, for example, can increase in some countries because of temperature-induced productivity drops in other countries. But for most

Figure 2: Monte Carlo Results on Total Gross Output, Four Estimators



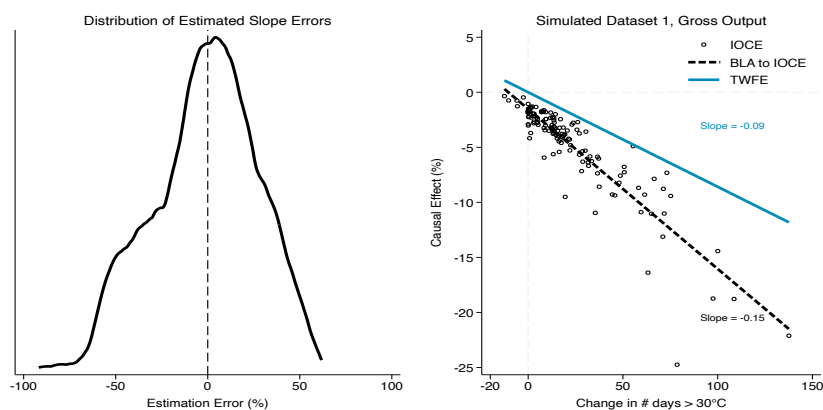
Notes: Figure plots the true median IOCE on total gross output in % on the y-axis against the change in the number (#) of days with maximum temperature above 30°C between 1991 and 2019 (repeated for each subplot) along with the median estimated IOCE for the TWFE (panel a), the heterogeneous robust estimator (panel b), the upstream/downstream estimator (panel c), and the structural estimator (panel d). The black dashed line plots best linear approximation (BLA) to the median IOCEs. The blue solid line presents the median estimated IOCEs, and the blue shaded region depicts the interquartile range of the distribution of these estimates. Simulations include 132 countries and 29 time periods and two sectors. Interquartile range for the upstream/downstream has been censored at 10% and -30% for ease of visual comparison.

Figure 2a overlays the median and interquartile range of estimates from the TWFE model exploiting year-by-year panel variation (equation 9), across the 200 replications. We see that the TWFE closely tracks the slope of best linear approximation to the function relating IOCE to own-unit shock magnitude, but misses the negative intercept.³⁵ The TWFE thus understates the harmful effect of extreme heat days by 3.16 percentage points for the average country in our simulations, and as much as 10 percentage points for some countries (see Figures I.4 and I.5). This appears to be in line with the result in Proposition 1: the TWFE estimator recovers difference in treatment effects, but misses the level of these

countries and outcomes, spatial linkages tend to amplify treatment effects in these simulations, indicating that reduced demand and higher input costs from negative productivity shocks dominate the substitution effects.

³⁵Note that Figure 2a compares the slope of the BLA to the IOCE to the slope coefficient from the panel TWFE (9). The mismatch in time periods adds extra terms to the decomposition from Proposition 1. See Appendix A.2 for the extension.

Figure 3: TWFE vs Best Linear Approximation (BLA)



Notes: Left panel plots the distribution of the difference between the TWFE slope and the slope of the BLA expressed in percent of the absolute value of the slope of the BLA across 200 replications. The right panel plots IOCEs by country (hollow dots), best linear approximation (BLA) (dashed line), and the TWFE (solid line) against the long difference change in # days above 30°C for a single simulated data set.

effects, on average. This result is not inconsistent with the TWFE often being very far from the slope of the IOCE in any given replication of the data (see Figure 3).³⁶


Extensions to the TWFE tend to exhibit qualitatively similar biases in our simulations. The heterogeneous-robust estimator systematically understates the harm by 5.44 percentage points for the average country (see Figure 2b), with errors of nearly 20 percentage points for some places (see Figure I.5b). The upstream/downstream estimator has the same direction of bias (see Figure 2c). Though it produces smaller errors on average (see Figure I.5c), the variation across replications is substantially larger than either the TWFE or the heterogeneous-robust estimator. These results indicate that reduced-form estimators fail to capture the level of the treatment effects from extreme heat days, in the presence of spillovers.

Figure 4 plots the aggregate treatment effects from the global time-series estimator of

³⁶Figure 3 plots the deviations between the TWFE and the BLA slopes across our 200 replications. The TWFE estimate may overstate or understate the slope of the BLA in any individual replication, by as much as 80%. The variance of these errors could be higher or lower in any given empirical setting, depending on the unobservables. In the left subplot, the mean of the distribution is -2.66%, which is indicative of a small downward bias. This reveals that the covariance condition does not hold, as the average covariance term is positive, but small, across replications. The right subplot illustrates the deviation for one replication, using the same display as Figure 2. Unfortunately, researchers cannot rule out larger covariance terms, hence larger biases in the slope of the TWFE, from looking at the data. The covariance between observed shocks and country's sensitivity to all unobserved shocks is not an empirical object and is hard to assess intuitively as it results from a general equilibrium.

Bilal & Känzig (2024).³⁷ The global time-series estimator does not produce reliable results in our simulations—the typical estimate has the wrong sign, and the range of errors extends over 100 percentage points in both directions. When we reduce the variance of unobserved productivity shocks, the mean estimate is instead large and negative (see Figure I.6b).³⁸ In our simulations, at least, the global-time series estimator behaves erratically in the presence of spillovers, and it is certainly capable of greatly exaggerating the treatment effect.

Figure 4: Estimated vs True Treatment Effects, Global Times Series Estimator



Notes: Figure plots the distribution of treatment effects on global aggregate gross output across 200 replications, expressed in % of “observed” final-year global aggregate gross output. The black dashed line plots the true causal effect. The solid blue plots the estimated effects via the global time-series approach.

Finally, Figure 2d plots results for our structural estimator. Across all metrics, the estimated IOCEs closely track the true IOCEs. The estimated bias across our 200 replications is close to zero, and the root mean squared error is the smallest by a substantial margin. The estimated errors within each replication are also small (see Figure I.3).

Since our structural estimator is designed for this data generating process, it is perhaps not surprising that the estimator performs well. Nevertheless, it is not obvious *ex ante* that the estimator would be successful in recovering unbiased estimates in finite samples, or that the errors would be consistently small under plausible parameterizations. These Monte Carlo experiments indicate that the structural estimator can succeed in samples like the one we encounter in our empirical application.³⁹

³⁷To implement their estimator, we aggregate gross output and number of extreme heat days globally for each replication before estimating equation (14). Since our model is static, we abstract from the dynamic considerations in Bilal & Känzig (2024).

³⁸These results do not appear to be driven by small sample bias, as we find approximately the same results when we double the length of the panel (see Figure I.6, panels c and d).

³⁹Figures I.2b and I.3b show that the alternative GMM estimator from Appendix E also performs quite

Additionally, our structural approach recovers the effects on real income (see Figure I.2). Real income is typically not observed, so we cannot use reduced-form estimators for this purpose. By virtue of equation (36), our structural estimator recovers these effects and thereby enables direct inference about the welfare consequences of climate change.⁴⁰

6 Empirical Results

In this section we estimate the effect of the change in global climate between 1991 and 2019 on agricultural gross output, manufacturing gross output, total gross output, and real income. We also compute potential future outcomes using climate projections from CMIP6 (Eyring et al., 2016), the outputs of which are used to trace out a global damage function.

6.1 Point Estimates and Marginal Effects

We implement the TWFE, heterogeneous-robust, and upstream/downstream estimators by regressing year-on-year growth in agricultural and manufacturing gross outputs, and their sum, on the change in the annual number of days with a maximum temperature above 30°C. In all regressions, we control for flexible continent-by-year and country fixed effects, as well as country-level variables $(1 - \eta_{is}) \alpha_{ish}^M$ interacted with year fixed effects, for consistency with our structural approach (31). We also control for the expected year-on-year change in the annual number of days with a maximum temperature above 30°C, computed via the procedure from Jones et al. (2025) (see Appendix H). Finally, we weight all regressions by baseline population. Standard errors are clustered by country. For the heterogeneous-robust estimator, standard errors are bootstrapped by sampling from the data with replacement (see Appendix B).

The first row of Table 1 reports point estimates and standard errors for the TWFE estimator. Column 1 shows that the number of extreme heat days is negatively correlated with agricultural gross output, as in many previous studies (Burke et al., 2015; Deryugina & Hsiang, 2014). An additional day with maximum temperature above 30°C is associated with 0.198% lower annual agricultural output, which is quite close to the values estimated in Deryugina & Hsiang (2014) using US county-level data.⁴¹ These point estimates are statistically significant at the 1% level.

well in our Monte Carlo experiments.

⁴⁰In principle, one could construct measures of real output by deflating country-sector nominal gross output using appropriate price indices. In practice, price indices are difficult to construct and are subject to a well-known set of measurement issues (Broda & Weinstein, 2006; Barrows et al., 2023). Given these limitations, we refrain from treating deflated gross output as an empirical object, and only compute real income via the structural assumptions of the model.

⁴¹The dependent variable is in natural log scale, so the exact calculation is $\% \Delta Y_{ist} = 100 \times (\exp(-0.00198) - 1) \approx -0.1978\%$

In column 2, we find that an additional day with maximum temperature above 30°C is associated with only 0.033% lower manufacturing output, which is statistically indistinguishable from zero. With manufacturing representing a far greater share of nominal output than agriculture in most cases, the coefficient on total gross output is closer to the latter than the former (-0.052%), and again, statistically indistinguishable from zero (column 3). Further analyses reveal that these TWFE point estimates mask significant heterogeneity (see Figure J.1). There is as much as a ten-fold larger coefficient on additional extreme heat days for initially-cooler places (10th percentile of the temperature distribution) versus initially-hotter places (90th percentile). This pattern is familiar from the literature (Hultgren et al., 2025), and consistent with a model in which warmer places have a lower response to additional extreme heat days.

In rows 2 and 3, we find that the heterogeneous-robust and upstream/downstream estimators yield similar patterns of correlation as the TWFE estimator—large negative effects in agriculture, and smaller and imprecise effects in manufacturing. Table 1 reports only the direct effect as point estimates for the upstream/downstream estimator, but the full set of point estimates are reported in Table J.1.

The bottom panel of Table 1 reports point estimates for structural parameters $\mu_{ag}^{30^\circ\text{C}}$ and $\mu_{manuf}^{30^\circ\text{C}}$, given different values of θ_{ag} and θ_{manuf} . In the fourth row of Table 1, the dependent variables are computed as described in Section 3.2, using our estimated values of $\theta_{ag} = 1.5$ and $\theta_{manuf} = 4.6$ (see Appendix G). As with the TWFE estimator, we find large and statistically significant effects in agriculture, with much smaller and statistically insignificant effects in manufacturing. This pattern is robust to using different values for θ_{ag} and θ_{manuf} drawn from the literature (subsequent rows), though unsurprisingly, the magnitudes of the estimated μ_s do depend on the value of θ_{ag} and θ_{manuf} .

We find little heterogeneity in the marginal effects of extreme heat days on productivity based on initial average temperature using our structural approach (see Figure J.1b). Comparing these results with the corresponding results for the TWFE (see Figure J.1a) suggests that the heterogeneity in marginal effects found with the TWFE estimator likely originates from a source other than direct mitigation of one’s own sensitivity. For example, the pattern of heterogeneity in the TWFE results could instead be due to differential exposure to trading partner shocks—i.e., network position.

Figure J.1 panels (a) and (b) also present correlations between cold days (maximum temperature less than 0°C) and nominal output and productivity. Cold days correlate with both nominal output and productivity, but in different ways. Using the TWFE estimator,

Table 1: Regression Point Estimates, Number of Extreme Heat Days

	Agriculture (1)	Manufacturing (2)	Gross Output (3)
<i>Reduced-Form Models</i>			
TWFE	-0.00198 (0.00045)	-0.00033 (0.00041)	-0.00052 (0.00047)
Heterogeneous-Robust	-0.00099 (0.00048)	-0.00042 (0.00049)	-0.00012 (0.00024)
Upstream/Downstream	-0.00633 (0.00226)	0.00133 (0.00077)	
<i>Structural Model</i>			
Preferred Specification ($\theta_{ag} = 1.5, \theta_{manuf} = 4.6$)	-0.00273 (0.00052)	-0.00044 (0.00065)	
Boehm et al (2023) ($\theta_{ag} = 0.8, \theta_{manuf} = 1.1$)	-0.00306 (0.00059)	-0.00065 (0.00057)	
Shapiro (2016) ($\theta_{ag} = 3.3, \theta_{manuf} = 8.5$)	-0.00335 (0.00070)	-0.00095 (0.00075)	
Fontagné (2022) ($\theta_{ag} = 6.9, \theta_{manuf} = 8.3$)	-0.00485 (0.00117)	-0.00129 (0.00076)	

Note: Dependent variable is year-on-year growth in agricultural and manufacturing gross outputs (columns 1 and 2), and the sum of both (column 3). Independent variable is the change in the annual number of days with a maximum temperature above 30°C. In all regressions, we control for flexible continent-by-year and country fixed effects, as well as country-level variables $(1 - \eta_{is}) \alpha_{ish}^M$ interacted with year fixed effects. We also control for the expected year-on-year change in the annual number of extreme heat days, computed via the procedure from Jones et al. (2025). All regressions are weighted by baseline population. Standard errors are clustered by country. For the heterogeneous-robust estimator, standard errors are bootstrapped by sampling from the data with replacement. Sample includes 132 countries over the period 1991-2019. In row 3, the point estimates for the Upstream/Downstream estimator reflect just the direct effects, whereas all other point estimates are reported in Table J.1.

we find that each additional cold day lowers nominal output for agriculture and manufacturing for countries with low initial temperatures, but increases both measures at higher initial average temperatures. For the structural estimator, we find the same pattern for manufacturing, but the sign of the correlation is reversed for agriculture—cold days improve agricultural productivity at low initial average temperatures, but harm agricultural productivity at high initial average temperatures.

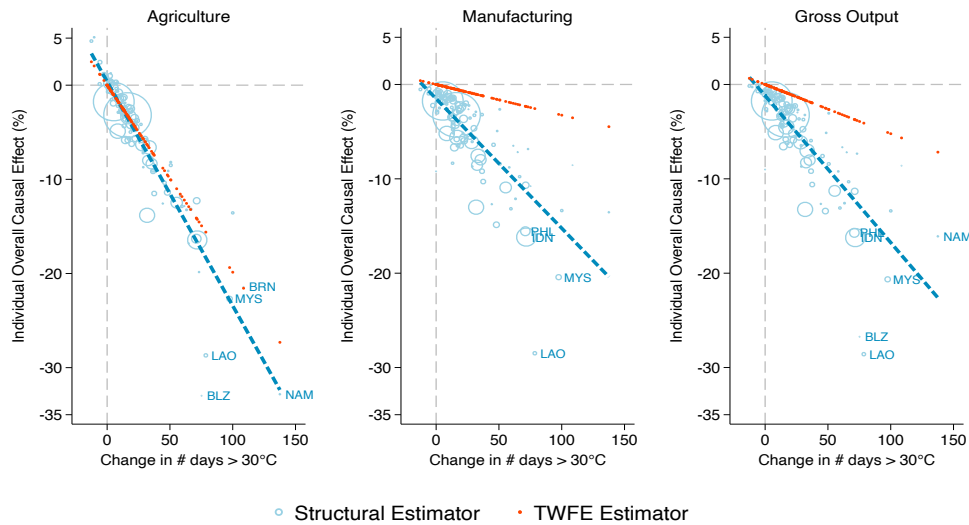
Lastly, for comparison, we also estimate specifications where the climate shock is parameterized as the annual average temperature, instead of hot and cold days. The annual figures are interacted with baseline average temperature, as in Nath et al. (2024). Using the TWFE, we find that an increase in annual average temperature correlates positively with

nominal output growth in agriculture for initially cold countries, but negatively for initially hot countries (see Figure J.2a). Using the structural approach, the effect of an increase in annual average temperature is also decreasing in initial average temperature for agricultural productivity, but the marginal effects are never positive (see Figure J.2b).

6.2 The Effects of Warming between 1991 and 2019

For the structural estimator, we plug in our estimates of μ_s and solve the system (32)-(35) for counterfactual change in outcomes between 1991 and 2019.⁴² In our baseline specification, we consider only the effect of extreme heat days and we impose $\theta_{ag} = 1.5$ and $\theta_{manuf} = 4.6$, based on our estimates in Appendix G. For comparison, we also present estimates based on the TWFE, heterogeneous-robust, and upstream/downstream estimators, as well as additional structural estimates that use alternative specifications for z and alternative measures of θ_{ag} and θ_{manuf} .

Figure 5: Structural Estimator vs TWFE Estimator



Notes: Figure plots the country-specific percentage estimated effects on agriculture gross output (left panel), manufacturing gross output (middle panel), and aggregate gross output (right panel), against observed change in the annual number of extreme heat days. Structural estimates are depicted with blue circles, which are scaled by the size of the workforce in 1991. The blue dashed line depicts the BLA to the structurally-estimated IOCEs, and the TWFE estimates are in red.

Figure 5 summarizes our structural IOCE estimates, with the TWFE estimates plotted

⁴²To account for trade imbalances, we proceed in two steps, as recommended by Ossa (2016). First, we solve a counterfactual, setting trade balances to zero. Second, we solve another counterfactual, assuming we had seen a recurrence of 1991 weather in 2019. The difference between the two counterfactual equilibria yields estimates of the IOCEs by country and sector.

for comparison. The left panel shows that the warming between 1991 and 2019 generated economically significant losses in agriculture for most countries. The mean loss was 5.0% (median loss of 3.2%), with the largest losses in Belize (32.9%), Namibia (32.8%), and Laos (28.6%). The losses were larger in the Tropics and Southern hemisphere, which saw the biggest increases in hot days (see Figure J.4), while regions that experienced smaller increases (such as Canada and much of Northern Europe) benefited in terms of agricultural output. The effects on manufacturing output (Figure 5, middle panel) and total gross output (right panel) were smaller, but still substantial—a mean loss of 4.6% of manufacturing output (median 2.9%), and 4.7% of total national gross output (median 3.0%).

We see the best linear approximation to the structurally-estimated IOCEs (the blue dashed line in Figure 5) passes near the origin, especially for agriculture. It would be a mistake, however, to infer that spillovers are negligible. Even though a non-zero intercept is sufficient to show the presence of spillovers, a zero intercept does not prove their absence. To assess the role of spillovers more rigorously, we decompose the overall IOCEs into a component related to own-country warming and a component related to other-countries warming:

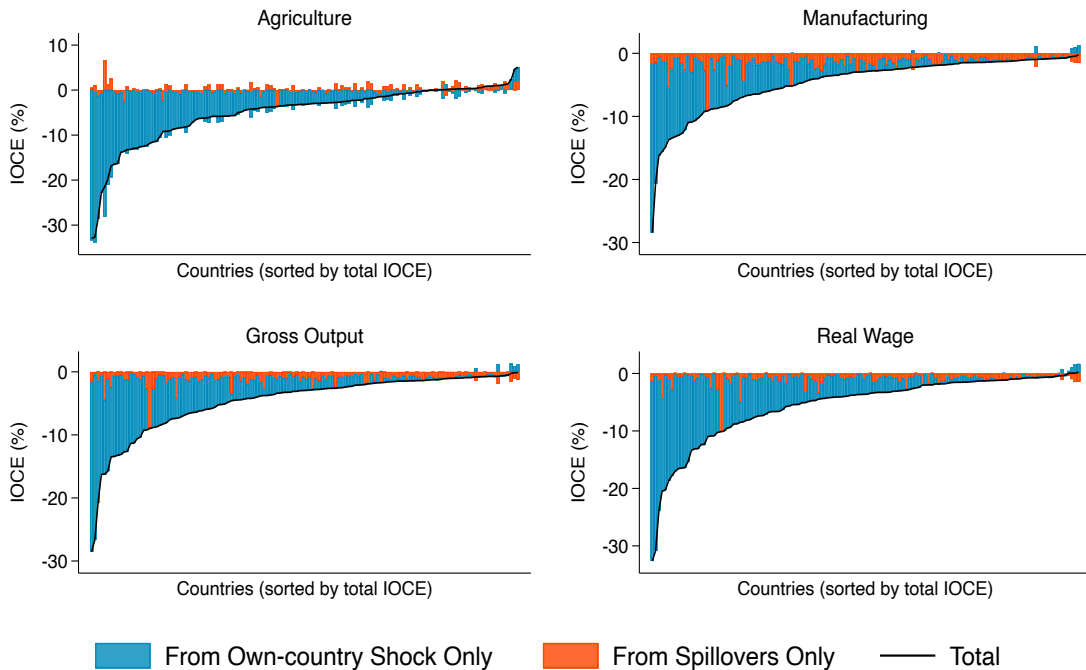
$$\Delta \ln y_i^\dagger = \underbrace{\ln y_{i,2019} - f_i(z'', \epsilon_{2019})}_{\text{influence from own-country warming}} + \underbrace{f_i(z'', \epsilon_{2019}) - f_i(z_{i,1991}, \epsilon_{2019})}_{\text{influence from other-country warming}} \quad (37)$$

with z'' indicating the vector z_{2019} with the i th element replaced with $z_{i,1991}$. The results of this exercise, summarized in Figure 6, reveal that the largest impacts are typically due to own-country shocks, but spillovers account for the lion's share of the effect in less exposed countries. For the median country, approximately 20% of the IOCE for the agricultural sector is due to spillover effects. In this sector, there is a mix of positive and negative spillovers, which happen to average out to a near-zero intercept in the aggregate. For manufacturing and total gross output, the spillovers are consistently negative, and represent a greater share of the overall effect—for the median country, 44% of the manufacturing IOCE and 31% of the total output IOCE. As a result, the estimated BLA intercept is larger in magnitude for manufacturing (-1.4%) and total gross output (-1.0%) than for agriculture.⁴³

⁴³Using our model, we can further decompose the overall IOCEs into the different mechanisms of action. The results of this decomposition are summarized in Figures J.7 and J.8. There are two main findings. First, cross-country and cross-sector spillovers account for a significant part of the IOCEs. While, for agricultural output, most of the IOCEs come through a direct productivity effect, changes in market access and agricultural prices are important mechanisms for both agricultural and manufacturing output. Second, labor reallocation plays very little role in our results. This is in contrast to Nath (2025)'s prediction that climate change may

Given these spillover patterns, we would expect the TWFE to perform better for agriculture than for manufacturing and total gross output. This is, indeed, what we see. The BLA to the structurally estimated IOCE for agriculture has a near-zero intercept, and the TWFE matches the slope of this line quite well. By contrast, with larger and more negative spillovers for manufacturing and total gross output, both the slope and intercept of the TWFE deviate further from the structural estimates. The TWFE appears to substantially understate the IOCEs, with a mean loss of only 1.2% of total gross output (median 0.7%). This is in line with published estimates that range from 1-3%, but is only 26% of the structurally-estimated IOCE (24% for the median).

Figure 6: Contribution to IOCE from Own-country Shock vs Spillovers



Notes: Figure plots the % change in 2019 agriculture gross output, manufacturing gross output, total gross output, and real wage, resulting from the change in # days with maximum temperature above 30°C between 1991 and 2019 estimated using our structural procedure (black line). We set $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$. In blue, we plot the contribution to IOCE from own-country temperature changes ($z_{i,2019}^{>30} - z_{i,1991}^{>30}$). To compute the contribution from own-country temperature changes, we solve for each country the counterfactual in which only the country's own temperature changes. The difference between this counterfactual and the total IOCE yields the contribution to the IOCE from other-country temperature changes, plotted in orange. Countries are sorted on the x-axis by outcome-specific IOCE.

induce large labor swings across sectors, unless preferences are non-homothetic. We do not observe such swings, at least for this counterfactual.

The other reduced-form estimators do not perform better. The heterogeneous-robust estimator delivers roughly the same results as the TWFE (see Fig. J.3a), which suggests that it may not be sufficient to allow for heterogeneous effects of own-country warming in the presence of spillovers. Estimates based on the upstream/downstream estimator display a lot of variation, conditional on own-country warming, but the results are, on average, closer to the TWFE and the heterogeneous-robust than to our structural estimator (see Fig. J.3b).

With our structural model, we can also assess the impacts on real income. The results, summarized in Figure J.4, indicate that real income losses were similar in magnitude and spatial distribution to gross output effects. Aggregating across the entire world, we find that real income was 5.5% lower in 2019, as a result of the warming from 1991.

Finally, we check robustness of these results by re-estimating under different climate shock specification and using alternative sets of trade elasticities (see Fig. J.5). We allow for productivity effects from both hot and cold days, or specify productivity damages as a function of annual average temperature. For each specification, we allow for heterogeneous effects depending on baseline average temperature, and consider imposing zero productivity effects in manufacturing. Across all 12 resulting specifications, the TWFE and heterogeneous-robust estimators yield consistently smaller estimates of climate damages compared with our preferred structural specification (means of -0.98%, -0.29%, and -3.3%, respectively). Allowing for effects of cold days attenuates the damages slightly, while an interaction with baseline temperature does little. Structural estimates based on trade elasticities from Boehm et al. (2023) are somewhat larger in magnitude (mean of -6.1%), while using trade elasticities from Shapiro (2016) and Fontagné et al. (2023) yields estimates that are slightly smaller in magnitude (mean -2.53% and -2.41%, respectively). The average IOCE computed via the upstream/downstream estimator is somehow erratic, but for most specifications this estimator understates the damages relative to our structural approach, as in the simulations.

6.3 The Effects of Warming between 2019 and 2100

The previous section presented estimates of the causal effects of warming between 1991 and 2019, a period over which global mean temperature rose by about 0.6°C. In this section we evaluate a broader range of counterfactuals, drawn from the IPCC's Coupled Model Intercomparison Project Phase 6 (CMIP6; Eyring et al., 2016). This allows us to trace out reduced-form damage curves similar to those used for integrated assessment of climate change (see, for instance, Nordhaus, 2017), and then estimate the global damage

elasticity of warming.

We construct our counterfactual climates using 32 climate-model realizations from CMIP6, a coordinated global modeling effort in which participating climate models simulate historical climate dynamics—typically from the mid-19th century to the present—using a common set of externally imposed forcings, representing greenhouse gas concentrations, aerosols, solar radiation, volcanic activity, and land-use changes. Each CMIP6 model also produces forward-looking projections under standardized scenarios, ensuring comparability across modeling centers. The CMIP6 simulations are internally consistent representations of future climates, in that each realization is made to respect physical constraints on energy balance and dynamics.

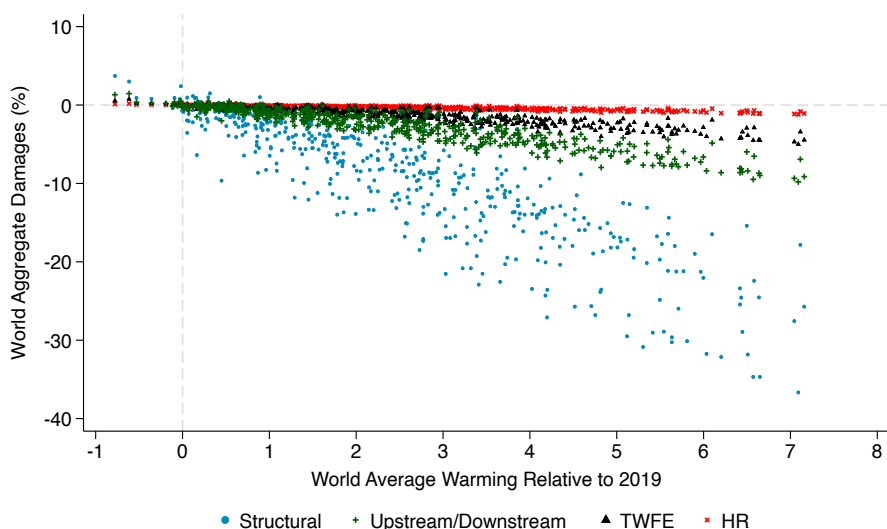
To construct a sample of climate counterfactuals with dense coverage in the high-warming part of the sample space and with a wide range of possible spatial patterns of warming, we extract one run of the SSP5-8.5 forcing scenario (i.e., high emissions) for each of the 32 CMIP6 models that conducted this simulation. Figure F.2 plots the trajectory of annual global mean temperature for each of the 32 models. We use every fifth year of the model outputs over the period 2019 to 2100, yielding 17 distinct climate draws per model, for a total of 544 counterfactual climate realizations across all 32 models.

Figure 7 plots the projected losses in global nominal output against the difference in global mean temperature compared to 2019, each point representing one of the 544 climate counterfactuals.⁴⁴ We also include historical counterfactuals for the years 1991, 2000, and 2010, computed as in the previous section. The three historical counterfactuals can be seen on the far left of Figure 7, as they are among the few that entail cooler climates than in 2019. Almost all of the counterfactuals drawn from CMIP6 entail a warmer climate.

We find that global climate damages tend to increase substantially in global mean temperature. A simple linear fit of the projected losses from our structural estimates (the blue dots) suggests that each additional 1°C causes about 4.6% loss in global nominal output on average. These damages are double or triple the elasticities that are usually cited in the literature, though only about one quarter as large as what Bilal & Känzig (2024) re-

⁴⁴We calculate the economic losses from climate change in two steps. We first solve for the change in outcomes that occurs when transitioning from the observed 2019 temperatures to the 2019 temperatures given by each CMIP6 model. This yields model-specific baseline values, allowing us to account for the variation in initial conditions seen in Figure F.2. As before, we also zero out trade balances in this step. Our second step is to solve for the further changes that occur when we transition from this model-specific baseline to a future climate realization from that same CMIP6 model run. The first step normalizes across CMIP6 models, while the second step yields a measure of the economic damage that would be caused by subjecting the 2019 world economy to a counterfactual climate.

Figure 7: Global Aggregate Damages For Different Warming Scenarios



Notes: Figure plots the percentage change in global nominal gross output under different counterfactual climate realizations, relative to the observed 2019 global nominal gross output. Each point represents the outcome associated with a single realization, including 3 historical climate realizations (1991, 2000, 2010) and 544 realizations based on projections from 32 different CMIP6 models. Structural estimates are plotted in blue with $\theta_{ag} = 1.5$ and $\theta_{manuf} = 4.6$, TWFE estimates in red, HR estimates in black, and Upstream/Downstream estimates in green. We include only the effect of extreme heat days.

port, using a global time series estimator. By contrast, using the TWFE (black triangles), heterogeneous-robust (red crosses), or upstream/downstream estimator (green crosses), we find elasticities of only -0.66% , -0.15% , and -1.40% , respectively.

Another interesting feature of Figure 7 is the variation in damages at a given temperature change. For example, if the world warms by 3°C on average, projected damages based on our structural estimates range between 8% and 20% of global output, depending on the spatial distribution of that warming. Estimates using TWFE and related methods appear to substantially understate uncertainty due to spatial warming patterns, and a fuller exploration of possible patterns would almost certainly extend our range further. This result emphasizes the need for incorporating substantial uncertainty about the damage elasticity in integrated assessment models with a single damage function. Exercises modeled on ours could help with calibrating the distribution of the damage function parameters.

Finally, we repeat the entire exercise for alternative climate shock specification and trade elasticities (see Figure J.6). Our structural estimator yields consistently larger damage elasticities than the TWFE and heterogeneous-robust estimators (mean elasticities of -5.8% , -0.76% , and -0.07% , respectively), while using trade elasticities from Boehm et al.

(2023) results in somewhat larger structural estimates (mean of -11.9%) and using those from Shapiro (2016) and Fontagné et al. (2023) results in slightly smaller estimates (mean -3.14% and -2.37%, respectively). Allowing for effects of cold days magnifies damages slightly, while interactions with baseline temperature generally do not affect the results much, with one exception—the upstream/downstream estimator yields a *positive* elasticity for this one specification, otherwise being very similar to the TWFE.

7 Conclusion

This paper proposes a structural approach for estimating climate damages. Most estimates in the literature are derived from the TWFE estimator or extensions thereof. We show that these estimators implicitly impose restrictive *ad hoc* constraints on the nature of cross-country spillovers. Our Monte Carlo experiments further demonstrate that these estimators do not reliably recover the effects of climate shocks when regions trade with each other.

We develop a structural estimator that draws on standard techniques from International Trade to model across-unit dependencies. The core of our contribution is a novel method for identifying the elasticity of productivity to climate. We rely on the same exclusion restrictions that are exploited in reduced-form estimation, but replace its implicit *ad hoc* assumptions with explicit assumptions that are grounded in theory and flexible enough to accommodate rich patterns of heterogeneity and spillovers. The main added data input our procedure requires, compared to reduced-form estimators, is the matrix of bilateral trade flows, information that is readily available at the international scale. Our procedure could also be used to quantify the effect of other observed country-level shock to productivity (conditional on the orthogonality assumption), beyond climate change, while accommodating a rich set of spillovers across countries.

In an empirical application, we find that extreme heat, as reflected by the number of days with maximum temperature over 30°C, mainly lowers agricultural productivity, while having a smaller effect on manufacturing. These productivity effects translate into losses of about 5% of global output and real wages in 2019, as compared to a counterfactual in which no warming had occurred since 1991. These estimates are at least 50% larger than the corresponding estimates based on the TWFE and extensions thereof. When we look out to the end of the century, we similarly project substantially larger climate damages, as well as much greater uncertainty about future damages driven by the spatial pattern of warming.

References

- Acemoglu, D., Akcigit, U., & Kerr, W. R. (2016). Networks and the macroeconomy: An empirical exploration. In M. Eichenbaum & J. A. Parker (Eds.), *NBER Macroeconomics Annual*, volume 30 of *NBER Chapters* (pp. 273–335).
- Adão, R., Arkolakis, C., & Esposito, F. (2020). General equilibrium effects in space: Theory and measurement. *National Bureau of Economic Research Working Paper*.
- Adão, R., Costinot, A., & Donaldson, D. (2024). Putting Quantitative Models to the Test: An Application to the US-China Trade War. *The Quarterly Journal of Economics*.
- Aguiar, A., Narayanan, B., & McDougall, R. (2016). An overview of the gtap 9 data base. *Journal of Global Economic Analysis*, 1(1), 181–208.
- Alves, G., Burton, W. H., & Fleitas, S. (2024). Difference-in-differences in equilibrium: Evidence from place-based policies. *Mimeo*.
- Anderson, J. E., Larch, M., & Yotov, Y. V. (2020). Transitional growth and trade with frictions: A structural estimation framework. *The Economic Journal*, 130(630), 1583–1607.
- Anderson, J. E. & van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1), 170–192.
- Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers*, 16(1), 159–178.
- Barrows, G., Calel, R., Jégard, M., & Ollivier, H. (2025). *Estimating Treatment Effects of Incomplete Regulation: A New Method with Application to the EU ETS*. Working paper, SSRN.
- Barrows, G., Ollivier, H., & Reshef, A. (2023). *Production Function Estimation with Multi-Destination Firms*. CESifo Working Paper Series 10716, CESifo.
- Bartelme, D., Lan, T., & Levchenko, A. A. (2024). Specialization, market access and real income. *Journal of International Economics*, 150, 103923.
- Bilal, A. & Känzig, D. R. (2024). The macroeconomic impact of climate change: Global vs. local temperature. *National Bureau of Economic Research Working Paper*, (32450).
- Boehm, C. E., Levchenko, A. A., & Pandalai-Nayar, N. (2023). The long and the short of trade elasticities. *Journal of Political Economy*, 131(7), 1771–1824.
- Borusyak, K., Dix-Carneiro, R., & Kovak, B. (2022). Understanding migration responses to local shocks. *Mimeo*.
- Borusyak, K. & Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6), 2155–2185.
- Broda, C. & Weinstein, D. E. (2006). Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2), 541–585.
- Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577), 235–239.
- Butts, K. (2021). Difference-in-differences estimation with spatial spillovers. *arXiv preprint arXiv:2105.03737*.
- Caliendo, L. & Parro, F. (2015). Estimates of the trade and welfare effects of NAFTA. *The*

- Review of Economic Studies*, 82(1), 1–44.
- Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Kopp, R. E., McCusker, K. E., Nath, I., Rising, J., Rode, A., Seo, H. K., Viaene, A., Yuan, J., & Zhang, A. T. (2022). Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*. *The Quarterly Journal of Economics*, 137(4), 2037–2105.
- CIESIN (2018). Gridded population of the world, version 4.11 (gpwv4): Population count. Revision 11. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC).
- Conte, M., Cotterlaz, P., Mayer, T., et al. (2022). The CEPII gravity database. *CEPII*.
- Costinot, A., Donaldson, D., & Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1), 205–248.
- Cruz, J.-L. & Rossi-Hansberg, E. (2024). The economic geography of global warming. *Review of Economic Studies*, 91(2), 899–939.
- Das, S., Magistretti, G., Pugacheva, E., & Wingender, P. (2022). Sectoral spillovers across space and time. *Journal of Macroeconomics*, 72(C), 103422.
- de Chaisemartin, C., D’Haultfoeuille, X., & Vazquez-Bare, G. (2024). Difference-in-difference estimators with continuous treatments and no stayers. *AEA Papers and Proceedings*, 114, 610–613.
- Dekle, R., Eaton, J., & Kortum, S. (2007). Unbalanced trade. *American Economic Review*, 97(2), 351–355.
- Dell, M., Jones, B., & Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3), 66–95.
- Deryugina, T. & Hsiang, S. M. (2014). Does the environment still matter? daily temperature and income in the United States. *Mimeo*.
- Deschênes, O. & Greenstone, M. (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American economic review*, 97(1), 354–385.
- Deschênes, O. & Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4), 152–185.
- Dingel, J. I. & Meng, K. C. (2025). Spatial correlation, trade, and inequality: Evidence from the global climate. *Mimeo*.
- Donaldson, D. (2015). The gains from market integration. *Annual Review of Economics*, 7(1), 619–647.
- Donaldson, D. & Hornbeck, R. (2016). Railroads and american economic growth: A “market access” approach. *The Quarterly Journal of Economics*, 131(2), 799–858.
- Eaton, J. & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5), 1741–1779.
- Egger, P. H. & Nigai, S. (2015). Structural gravity with dummies only: Constrained anova-

- type estimation of gravity models. *Journal of International Economics*, 97(1), 86–99.
- Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization. *Geoscientific Model Development*, 9(5), 1937–1958.
- Feng, A., Li, H., & Wang, Y. (2025). We are all in the same boat: Cross-border spillovers of climate shocks through international trade and supply chains. *The Economic Journal*, 135(669), 1433–1466.
- Fontagné, L., Guimbard, H., & Orefice, G. (2022). Tariff-based product-level trade elasticities. *Journal of International Economics*, 137, 103593.
- Fontagné, L., Lebrand, M. S. M., Murray, S., Santoni, G., & Ruta, M. (2023). Trade and infrastructure integration in africa. *Mimeo, Available at SSRN 4672520*.
- Fuglie, K. O. (2024). Agricultural productivity and climate mitigation. *Annual Review of Resource Economics*, 16, 1–22.
- Head, K. & Ries, J. (2008). Fdi as an outcome of the market for corporate control: Theory and evidence. *Journal of International Economics*, 74(1), 2–20.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Sabater, J. M., Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al. (2023). ERA5 hourly data on single levels from 1940 to present.
- Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, 8, 43–75.
- Hsieh, C.-T. & Ossa, R. (2016). A global view of productivity growth in china. *Journal of international Economics*, 102, 209–224.
- Hudgens, M. G. & Halloran, M. E. (2008). Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482), 832–842.
- Hultgren, A., Carleton, T., Delgado, M., Gergel, D. R., Greenstone, M., Houser, T., Hsiang, S., Jina, A., Kopp, R. E., Malevich, S. B., et al. (2025). Impacts of climate change on global agriculture accounting for adaptation. *Nature*, 642(8068), 644–652.
- IPCC (2021). *Climate Change 2021: The Physical Science Basis*. Cambridge Univ. Press.
- Jones, B., Moscona, J., von Dessauer, C., & Olken, B. (2025). With or without u? binning bias and the causal effects of temperature shocks. *Mimeo*.
- Leung, M. P. (2020). Treatment and spillover effects under network interference. *Review of Economics and Statistics*, 102(2), 368–380.
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3), 531–542.
- Mayer, T., Santoni, G., & Vicard, V. (2023). *The CEPII trade and production database*. CEPII.
- Nath, I. (2025). Climate change, the food problem, and the challenge of adaptation through sectoral reallocation. *Journal of Political Economy*, 133(6), 1705–1756.
- Nath, I. B., Ramey, V. A., & Klenow, P. J. (2024). How much will global warming cool global growth? *NBER Working Paper No. 32761*.
- Nordhaus, W. D. (1992). An optimal transition path for controlling greenhouse gases. *The*

- Energy Journal*, 13(1), 1–21.
- Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7), 1518–1523.
- Osberghaus, D. & Schenker, O. (2022). International trade and the transmission of temperature shocks. *ZEW-Centre for European Economic Research Discussion Paper*.
- Ossa, R. (2016). Quantitative models of commercial policy. In *Handbook of commercial policy*, volume 1 (pp. 207–259). Elsevier.
- Parsons, C. R., Skeldon, R., Walmsley, T. L., & Winters, L. A. (2007). *Quantifying International Migration: A Database of Bilateral Migrant Stocks*. Policy Research Working Paper 4165, World Bank.
- Redding, S. & Venables, A. J. (2004). Economic geography and international inequality. *Journal of international Economics*, 62(1), 53–82.
- Rudik, I., Lyn, G., Tan, W., & Ortiz-Bobea, A. (2022). The economic effects of climate change in dynamic spatial equilibrium. *Mimeo*.
- Schlenker, W. & Roberts, M. J. (2006). Nonlinear effects of weather on corn yields. *Review of agricultural economics*, 28(3), 391–398.
- Shapiro, J. & Walker, R. (2018). Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade. *American Economic Review*, 108(12), 3814–3854.
- Shapiro, J. S. (2016). Trade costs, co2, and the environment. *American Economic Journal: Economic Policy*, 8(4), 220–254.
- Silva, J. S. & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and statistics*, 88(4), 641–658.
- Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476), 1398–1407.
- Vazquez-Bare, G. (2023). Identification and estimation of spillover effects in randomized experiments. *Journal of Econometrics*, 237(1), 105237.
- Zappalà, G. (2024). Estimating sectoral climate impacts in a global production network. *Mimeo*.

Supplemental Appendix

A Theoretical Proofs

A.1 Proof of Proposition 1

First, we provide the proof for Proposition 1. Using equation (15), we can write:

$$\Delta \ln y_i = \Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T) + \Delta \vartheta_i$$

where $\Delta \vartheta_i = f_i(z_0, \varepsilon_T) - f_i(z_0, \varepsilon_0)$. Now, consider the slope parameter (10) of the TWFE model (8). We can rewrite it as

$$\begin{aligned} \check{\beta}^{FE} &= \frac{\sum_i (\Delta \ln y_i - \overline{\Delta \ln y}) (\Delta z_i - \overline{\Delta z})}{\sum_i (\Delta z_i - \overline{\Delta z}) (\Delta z_i - \overline{\Delta z})} \\ &= \frac{\sum_i (\Delta \ln y_i^\dagger + \Delta \vartheta_i - \overline{\Delta \ln y^\dagger} - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \\ &= \frac{\sum_i (\Delta \ln y_i^\dagger - \overline{\Delta \ln y^\dagger}) (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \\ &= \beta^{BLA}(z_T, z_0; \varepsilon_T) + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2}. \end{aligned}$$

One way to interpret this result is that the slope of the TWFE model equals the slope of the BLA plus a ratio of covariance between ϑ and z over the variance of z in the sample. As stated by the proposition, if and only if $Cov(\Delta \vartheta, \Delta z) = 0$, then $\mathbb{E} [\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})] = 0$, which implies that $\mathbb{E} [\check{\beta}^{FE}] = \mathbb{E} [\beta^{BLA}(z_T, z_0; \varepsilon_T)]$.

For the intercept, we use the definition from (11), i.e., $\check{\alpha}^{FE} = \overline{\Delta \ln y} - \check{\beta}^{FE} \overline{\Delta z}$, combined with $\overline{\Delta \ln y} = \overline{\Delta \ln y^\dagger} + \overline{\Delta \vartheta}$. This yields

$$\begin{aligned} \check{\alpha}^{FE} &= \overline{\Delta \ln y^\dagger} + \overline{\Delta \vartheta} - \check{\beta}^{FE} \overline{\Delta z} \\ &= \overline{\Delta \ln y^\dagger} + \overline{\Delta \vartheta} - \left(\beta^{BLA}(z_T, z_0; \varepsilon_T) + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \right) \overline{\Delta z} \\ &= \alpha^{BLA}(z_T, z_0; \varepsilon_T) + \overline{\Delta \vartheta} - \overline{\Delta z} \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \end{aligned}$$

As before, if and only if $Cov(\Delta \vartheta, \Delta z) = 0$, then $\mathbb{E} [\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})] = 0$, which implies that $\mathbb{E} [\check{\alpha}^{FE}] = \mathbb{E} [\alpha^{BLA}(z_T, z_0; \varepsilon_T)] + \mathbb{E} [\overline{\Delta \vartheta}]$. \square

Second, to gain intuition on the condition specified by Proposition 1, we proceed by taking a Taylor series expansion of $f_i(z_0, \boldsymbol{\varepsilon}_T) - f_i(z_0, \boldsymbol{\varepsilon}_0)$ around $(z_0, \boldsymbol{\varepsilon}_0)$. Thus,

$$\Delta \ln y_i = \Delta \ln y_i^\dagger(z_T, z_0; \boldsymbol{\varepsilon}_1) + \sum_{k=1}^{\infty} \frac{1}{k!} \sum_{j_1=1}^N \cdots \sum_{j_k=1}^N \frac{\partial^k f_i}{\partial \varepsilon_{j_1} \cdots \partial \varepsilon_{j_k}} \Big|_{(z_0, \boldsymbol{\varepsilon}_0)} \prod_{r=1}^k \Delta \varepsilon_{j_r} \quad (\text{A.1})$$

Taking expectations over the vector $\Delta \boldsymbol{\varepsilon}$, we obtain

Lemma 1. *Assume $\mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\prod_{r=1}^k \Delta \varepsilon_{j_r} \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] = m(k)$ for all $\{j_1, \dots, j_k\} \subseteq \{1, \dots, N\}$, where $m(k) \in \mathbb{R}$ is a constant that does not depend on the choice of the indices, but only on k . Then*

$$\mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\check{\beta}^{FE} \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] = \mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\sum_i \zeta_i \frac{\Delta \ln y_i^\dagger(z_T, z_0; \boldsymbol{\varepsilon}_T)}{\Delta z_i} \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] + \sum_i \zeta_i \frac{\Delta \xi_i}{\Delta z_i},$$

where $\Delta \xi_i \equiv \sum_{k=1}^{\infty} \frac{m(k)}{k!} \sum_{j_1=1}^N \cdots \sum_{j_k=1}^N \frac{\partial^k f_i}{\partial \varepsilon_{j_1} \cdots \partial \varepsilon_{j_k}} \Big|_{(z_0, \boldsymbol{\varepsilon}_0)}$.

Proof: Combining (10) with (A.1), given the definition of ζ_i , yields

$$\begin{aligned} \mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\check{\beta}^{FE} \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] &= \mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \boldsymbol{\varepsilon}_T)}{\Delta z_i} \right) \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] \\ &\quad + \mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\sum_i \zeta_i \left(\frac{\sum_{k=1}^{\infty} \frac{1}{k!} \sum_{j_1=1}^N \cdots \sum_{j_k=1}^N \frac{\partial^k f_i}{\partial \varepsilon_{j_1} \cdots \partial \varepsilon_{j_k}} \Big|_{(z_0, \boldsymbol{\varepsilon}_0)} \prod_{r=1}^k \Delta \varepsilon_{j_r}}{\Delta z_i} \right) \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] \\ &= \mathbb{E}_{\Delta \boldsymbol{\varepsilon}} \left[\sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \boldsymbol{\varepsilon}_T)}{\Delta z_i} \right) \mid z_T, z_0, \boldsymbol{\varepsilon}_0 \right] + \sum_i \frac{\zeta_i}{\Delta z_i} \left(\sum_{k=1}^{\infty} \frac{m(k)}{k!} \sum_{j_1=1}^N \cdots \sum_{j_k=1}^N \frac{\partial^k f_i}{\partial \varepsilon_{j_1} \cdots \partial \varepsilon_{j_k}} \Big|_{(z_0, \boldsymbol{\varepsilon}_0)} \right) \end{aligned}$$

The last equality follows from the law of iterated expectations and the condition from Lemma 1. \square

The condition from Lemma 1 corresponds to imposing stationarity in the k th moment for unobserved shocks. For $k = 1$, it corresponds to the usual conditional mean independence assumption. For $k > 1$, the joint distribution of $\Delta \boldsymbol{\varepsilon}_{j_r}$ must be such that the expectation of the product of draws does not depend on which units are considered. A sufficient condition for this to hold is that $\Delta \boldsymbol{\varepsilon}_{j_r}$ are i.i.d. In this case, $m(k) = (\mathbb{E}[\Delta \boldsymbol{\varepsilon}])^k$. A weaker condition requires that, in expectation, the product of unobserved factor draws is the same for any index sets.

Lemma 1 reveals that, taking expectations only over the vector of changes in the unobserved shocks, and under the specified condition, the TWFE estimator can be decomposed into two components: first, the expected β^{BLA} conditional on z_1, z_0 , and $\boldsymbol{\varepsilon}_0$; second, a weighted average of the slope of the individual overall sensitivity to all unobserved shocks

over each unit's own observed shock. The second component does not depend on the realization of the changes in the unobserved shocks, as each unit's overall sensitivity is determined in the first period.

Furthermore,

Lemma 2. *If the assumption from Lemma 1 holds, and if*

$$\mathbb{E} \left[\sum_i \left(\Delta \xi_i - \overline{\Delta \xi} \right) \left(\frac{(\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \right) \right] = 0 \text{ with } \overline{\Delta \xi} \equiv \frac{1}{N} \sum_k \Delta \xi_k,$$

$$\text{then } \mathbb{E} [\check{\beta}^{FE}] = \mathbb{E} [\beta^{BLA}] \text{ and } \mathbb{E} [\check{\alpha}^{FE}] = \mathbb{E} [\alpha^{BLA}] + \mathbb{E} [\overline{\Delta \xi}].$$

Proof: Using the law of iterated expectations (LIE), we have

$$\begin{aligned} \mathbb{E} [\check{\beta}^{FE}] &= \mathbb{E}_{z_T, z_0, \varepsilon_0} \left[\mathbb{E}_{\Delta \varepsilon} [\check{\beta}^{FE} | z_T, z_0, \varepsilon_0] \right] \\ &= \mathbb{E}_{z_T, z_0, \varepsilon_0} \left[\mathbb{E}_{\Delta \varepsilon} \left[\sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T)}{\Delta z_i} \right) \middle| z_T, z_0, \varepsilon_0 \right] + \sum_i \zeta_i \frac{\Delta \xi_i}{\Delta z_i} \right] \\ &= \mathbb{E} \left[\sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T)}{\Delta z_i} \right) \right] + \mathbb{E}_{z_T, z_0, \varepsilon_0} \left[\sum_i \frac{(\Delta z_i - \overline{\Delta z})}{\sum_k \Delta z_k (\Delta z_k - \overline{\Delta z})} \Delta \xi_i \right] \\ &= \mathbb{E} \left[\sum_i \zeta_i \left(\frac{\Delta \ln y_i^\dagger(z_T, z_0; \varepsilon_T)}{\Delta z_i} \right) \right] = \mathbb{E} [\beta^{BLA}(z_T, z_0; \varepsilon_T)]. \end{aligned}$$

The second equality follows from Lemma 1. The third equality follows from the LIE and the definition of ζ_i . The fourth equality follows from the condition specified in Lemma 2.

Moreover,

$$\begin{aligned} \mathbb{E} [\check{\alpha}^{FE}] &= \mathbb{E}_{z_T, z_0, \varepsilon_0} \left[\mathbb{E}_{\Delta \varepsilon} [\check{\alpha}^{FE} | z_T, z_0, \varepsilon_0] \right] \\ &= \mathbb{E}_{z_T, z_0, \varepsilon_0} \left[\mathbb{E}_{\Delta \varepsilon} \left[\left(\overline{\Delta \ln y^\dagger} + \frac{1}{N} \sum_i \sum_{k=1}^{\infty} \frac{1}{k!} \sum_{j_1=1}^N \cdots \sum_{j_k=1}^N \frac{\partial^k f_i}{\partial \varepsilon_{j_1} \cdots \partial \varepsilon_{j_k}} \bigg|_{(z_0, \varepsilon_0)} \prod_{r=1}^k \Delta \varepsilon_{j_r} - \check{\beta}^{FE} \overline{\Delta z} \right) \middle| z_T, z_0, \varepsilon_0 \right] \right] \\ &= \mathbb{E} [\alpha^{BLA}(z_T, z_0; \varepsilon_T)] + \mathbb{E} [\overline{\Delta \xi}]. \end{aligned}$$

The first equality follows from the LIE. The second equality follows from (11) and (A.1). The last equality follows from Lemma 1 and from the result for $\mathbb{E} [\check{\beta}^{FE}]$. \square

The condition specified in Lemma 2, in addition to the condition on stationarity in the

k th moment for unobserved shocks, is an orthogonality requirement: in expectation the covariance between the change in individual overall sensitivity to all unobserved shocks and the normalized deviations of the observed shocks must be zero. The two conditions spelled out in Lemmas 1 and 2 decompose the condition from Proposition 1 into a component that depends on the realization of unobserved shocks and a component that depends on country-level sensitivity to these unobserved shocks prior to any realizations. This decomposition helps build our intuition, even though the conditions they rely on are less general than Proposition 1.

A.2 Extension to Multiple Period Model

In practice, most TWFE regressions use year-on-year variation in outcomes and shocks, instead of relying on a single long difference between $t = 0$ and $t = T$. Hence, researchers exploit variation from the intermediate periods $t \in \{1, T - 1\}$ and use the following model in first differences (i.e., $\Delta x_{it} \equiv x_{it} - x_{it-1}$ for any variable x):

$$\Delta \ln y_{it} = \beta^{FE,p} \Delta z_{it} + \alpha_t^{FE,p} + \Delta \xi_{it},$$

where $\alpha_t^{FE,p}$ denotes time fixed effects and $\Delta \xi_{it}$ denotes the error term. This regression corresponds to model (9), reproduced here for convenience and denoted by "p" for panel version. Each observation therefore represents the change in outcomes and shocks between $t - 1$ and t . Relative to a more standard TWFE specification in levels, our first-difference formulation is slightly more flexible with respect to the evolution of outcomes over time, as it allows for arbitrary unit-specific level differences and removes persistent unit-level components through differencing.

The slope and constant parameters can be computed in the usual way:

$$\begin{aligned} \check{\beta}^{FE,p} &= \frac{\sum_t \sum_i \Delta \ln y_{it} (\Delta z_{it} - \overline{\Delta z_t})}{\sum_t \sum_i \Delta z_{it} (\Delta z_{it} - \overline{\Delta z_t})} \\ \check{\alpha}_t^{FE,p} &= \overline{\Delta \ln y_t} - \check{\beta}^{FE,p} \overline{\Delta z_t}, \end{aligned}$$

where $\overline{\Delta \ln y_t}$ ($\overline{\Delta z_t}$) indicates the average value of $\Delta \ln y_{it}$ (Δz_{it}) in the economy at time t .

To make progress in comparing the best linear approximation to the IOCEs with the slope $\beta^{FE,p}$, we must consider a different target parameter, namely year-on-year IOCEs:

$$\Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \varepsilon_t) \equiv f_i(z_t, \varepsilon_t) - f_i(z_{t-1}, \varepsilon_t). \quad (\text{A.2})$$

Given (15), we obtain the following decomposition:

$$\Delta \ln y_{it} = \Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \boldsymbol{\varepsilon}_t) + \Delta \vartheta_{it},$$

where $\Delta \vartheta_{it} = f_i(z_t, \boldsymbol{\varepsilon}_t) - f_i(z_{t-1}, \boldsymbol{\varepsilon}_{t-1})$.

We obtain the BLA to the panel IOCEs (A.2) as follows:

$$(\alpha^{BLA,p}, \beta^{BLA,p}) \equiv \arg \min_{a,b} \sum_{t=0}^T \sum_{i=1}^N \left(\Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \boldsymbol{\varepsilon}_t) - a - b \Delta z_{it} \right)^2. \quad (\text{A.3})$$

Solving (A.3) yields

$$\beta^{BLA,p} = \sum_t \sum_i \zeta_{it} \left(\frac{\Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \boldsymbol{\varepsilon}_t)}{\Delta z_{it}} \right) \quad (\text{A.4})$$

$$\alpha^{BLA,p} = \overline{\Delta \ln y_{it}^\dagger} - \beta^{BLA,p} \overline{\Delta z_{it}} \quad (\text{A.5})$$

with $\zeta_{it} \equiv \frac{\Delta z_{it} (\Delta z_{it} - \overline{\Delta z_t})}{\sum_k \Delta z_{kt} (\Delta z_{kt} - \overline{\Delta z_t})}$ and $\sum_i \zeta_{it} = 1$, and where $\overline{\Delta \ln y_{it}^\dagger}$ and $\overline{\Delta z_{it}}$ indicate the average values of $\Delta \ln y_{it}^\dagger$ and Δz_{it} across the entire sample.

We obtain

Proposition 1. *Using a TWFE panel regression (9), we obtain the following decomposition:*

$$\begin{aligned} \check{\beta}^{FE,p} &= \beta^{BLA}(z_T, z_0; \boldsymbol{\varepsilon}_T) + Dev^{\beta^{BLA}} + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta})(\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2} + Dev^{Cov} \\ \check{\alpha}_t^{FE,p} &= \alpha^{BLA}(z_T, z_0; \boldsymbol{\varepsilon}_T) + Dev^{\alpha^{BLA}} + \overline{\Delta \vartheta_t} - \overline{\Delta z_t} \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta})(\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2} - \overline{\Delta z_t} Dev^{Cov}, \end{aligned}$$

where $Dev^{XBLA} \equiv X^{BLA,p} - X^{BLA}(z_T, z_0; \boldsymbol{\varepsilon}_T)$ and $Dev^{Cov} \equiv \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t})(\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} - \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta})(\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2}$.

PROOF: First, we decompose the TWFE panel slope estimate into

$$\begin{aligned}
\check{\beta}^{FE,p} &= \frac{\sum_t \sum_i (\Delta \ln y_{it} - \overline{\Delta \ln y_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} \\
&= \frac{\sum_t \sum_i \left(\Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \varepsilon_t) + \Delta \vartheta_{it} - \overline{\Delta \ln y_t^\dagger} - \overline{\Delta \vartheta_t} \right) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} \\
&= \frac{\sum_t \sum_i (\Delta \ln y_{it}^\dagger(z_t, z_{t-1}; \varepsilon_t) - \overline{\Delta \ln y_t^\dagger}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} + \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} \\
&= \beta^{BLA,p} + \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})}.
\end{aligned}$$

Thus, we obtain the following decomposition:

$$\begin{aligned}
\check{\beta}^{FE,p} &= \beta^{BLA}(z_T, z_0; \varepsilon_T) + \left(\beta^{BLA,p} - \beta^{BLA}(z_T, z_0; \varepsilon_T) \right) + \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2} \\
&+ \left(\frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} - \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2} \right),
\end{aligned}$$

where we note that $\Delta x_i \equiv x_{iT} - x_{i0}$ is written in long difference, whereas $\Delta x_{it} \equiv x_{it} - x_{it-1}$. This expression simplifies into the one from Proposition 1 if we denote $\text{Dev}^{\beta^{BLA}} \equiv \beta^{BLA,p} - \beta^{BLA}(z_T, z_0; \varepsilon_T)$ and $\text{Dev}^{Cov} \equiv \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} - \frac{\sum_i (\Delta \vartheta_i - \overline{\Delta \vartheta}) (\Delta z_i - \overline{\Delta z})}{\sum_k (\Delta z_k - \overline{\Delta z})^2}$.

For the intercept, we have

$$\begin{aligned}
\check{\alpha}^{FE,p} &= \overline{\Delta \ln y_t^\dagger} + \overline{\Delta \vartheta_t} - \check{\beta}^{FE,p} \overline{\Delta z_t} \\
&= \overline{\Delta \ln y_t^\dagger} + \overline{\Delta \vartheta_t} - \left(\beta^{BLA,p} + \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})} \right) \overline{\Delta z_t} \\
&= \alpha^{BLA,p} + \overline{\Delta \vartheta_t} - \overline{\Delta z_t} \frac{\sum_t \sum_i (\Delta \vartheta_{it} - \overline{\Delta \vartheta_t}) (\Delta z_{it} - \overline{\Delta z_t})}{\sum_s \sum_k \Delta z_{ks} (\Delta z_{ks} - \overline{\Delta z_s})}.
\end{aligned}$$

Therefore, the expression from Proposition 1 is obtained by denoting $\text{Dev}^{\alpha^{BLA}} \equiv \alpha^{BLA,p} - \alpha^{BLA}(z_T, z_0; \varepsilon_T)$ and using the definition of Dev^{Cov} . \square

Relative to Proposition 1, Proposition 1 decomposes the slope and constant parameters obtained in the panel version of the TWFE model into terms that are similar to the ones from Proposition 1 and additional terms that identify deviations in the parameters in the panel version relative to the long difference version, and similar deviations for the empirical covariance terms. This extension therefore adds extra terms that could induce the

TWFE model parameters to deviate even more (or less) from the BLA to the IOCEs defined between $t = 0$ and $t = T$.

B Heterogeneous-Robust Estimator of de Chaisemartin et al. (2024)

de Chaisemartin et al. (2024) build an estimator for continuous treatment and without stayers, which matches our empirical application.

In de Chaisemartin et al. (2024), but with our notation, they describe a setting with two periods, $t \in \{0, T\}$, and a continuous treatment z_t on the support \mathcal{Z}_t . For any $(z_0, z_T) \in \mathcal{Z}$, let $y_{it}(z_0, z_T)$ denote the unit i 's potential outcome at t with treatment z , and let y_{it} denote the observed outcomes. They impose the following assumptions: (i) static model for potential outcomes $y_{it}(z_t)$; (ii) parallel trends $\forall x \in \mathcal{Z}_0, E(\Delta y | z_0 = x, z_T) = E(\Delta y | z_0 = x)$; (iii) \mathcal{Z}_t is a bounded subset of \mathbb{R} and potential outcomes are Lipschitz bounded. Finally, they also assume that there are no stayers, but there are quasi-stayers: $P(\Delta z = 0) = 0$ and $P(|\Delta z| \leq \eta) > 0, \forall \eta > 0$. The estimand is the following:

$$\theta_0 = E \left(\frac{\text{sign}(\Delta z) \Delta y_i^\dagger(z_T, z_0; \epsilon_1)}{E(|\Delta z|)} \right), \quad (\text{B.1})$$

which is a weighted average of the slopes of units' potential-outcome functions, from their period-one to their period-two treatment, which they call the WAOSS (weighted average of switchers' slopes). Their Theorem 1 establishes that without stayers, θ_0 is identified by the limit of a difference-in-difference comparing Δy of all units and of quasi-stayers.

By definition, we have $\Delta y_i^\dagger(z_T, z_0; \epsilon_T) = \Delta y_i - [f_i(z_0, \epsilon_T) - f_i(z_0, \epsilon_0)]$. Hence, θ_0 is equivalent to the expression we consider in the main text of the paper:

$$\check{\beta}^{HR} = \frac{\frac{1}{N} \sum_{i=1}^N \text{sign}(\Delta z_i) \times \left(\Delta \ln y_i - \check{E}[f(z_0, \epsilon_T) - f(z_0, \epsilon_0) | z_0] \right)}{\frac{1}{N} \sum_{i=1}^N |\Delta z_i|}, \quad (\text{B.2})$$

which can be computed once $E[f(z_0, \epsilon_1) - f(z_0, \epsilon_0) | z_0]$ has been estimated.

For implementation, we follow de Chaisemartin et al. (2024) in considering a simple two-step estimator based on the following parametric restriction: $g(z_0, \delta) = g_{\lambda_0}(z_0, \delta)$, where the family $(g_\lambda)_{\lambda \in \mathcal{R}^p}$ is known, but λ_0 not. We adopt the following assumptions:

$$g(z_0, \delta) = E[f(z_0, \epsilon_T) - f(z_0, \epsilon_0) | z_0] + \delta E \left[\frac{f(z_0 + \delta, \epsilon_T) - f(z_0, \epsilon_T)}{\delta} \Big| z_0, \Delta z = \delta \right] \quad (\text{B.3})$$

where

$$E[f(z_0, \varepsilon_T) - f(z_0, \varepsilon_0)|z_0] = \lambda_{0,1} + \lambda_{0,2}z_0 + \lambda_{0,3}\Delta X \quad (\text{B.4})$$

$$E\left[\frac{f(z_0 + \delta, \varepsilon_T) - f(z_0, \varepsilon_T)}{\delta} \Big| z_0, \Delta z = \delta\right] = \lambda_{0,4} + \lambda_{0,5}z_0 + \lambda_{0,6}\delta, \quad (\text{B.5})$$

where X denote several control variables (e.g., labor endowment, precipitation). Thus the first step consists in estimating λ_0 and then computing $g_{\hat{\lambda}_0}(z_0, 0)$, which equals $\check{E}[f(z_0, \varepsilon_T) - f(z_0, \varepsilon_0)|z_0]$. In a second step, we compute $\check{\beta}^{HR}$.

Additionally, we extend this implementation to the panel structure with multiple periods. Consider $\Delta x_{it} \equiv x_{it} - x_{it-1}$ for any variable x , we run the following regression for each period t :

$$\Delta y_{it} = \lambda_{t,1} + \lambda_{t,2}z_{i,t-1} + \lambda_{t,3}\Delta L_{ist} + \lambda_{t,4}\Delta z_{it} + \lambda_{t,5}z_{i,t-1}\Delta z_{it} + \lambda_{t,6}\Delta z_{it}\Delta z_{it} + \varepsilon_{it},$$

where the outcome y denotes either the sectoral gross output or the total gross output. We then build

$$\check{\beta}^{HR} = \frac{\frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N \text{sign}(\Delta z_{it}) \times \left(\Delta \ln y_{it} - \check{\lambda}_{t,1} - \check{\lambda}_{t,2}z_{i,t-1} - \check{\lambda}_{t,3}\Delta L_{ist} \right)}{\frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N |\Delta z_{it}|}, \quad (\text{B.6})$$

for each outcome. To fit the one-treatment-variable setting, we only consider the number of days with maximum temperature reaching at least 30°C as variable z .

Standard errors for $\check{\beta}^{HR}$ are computed via a cluster-level bootstrap, resampling countries with their full time-series of outcomes, weather realizations, and controls, 50 times, with replacement. The estimator has a two-step structure: the first step recovers period-specific coefficients $(\check{\lambda}_{t,1}, \check{\lambda}_{t,2}, \check{\lambda}_{t,3})$ from T separate OLS regressions, which are then used to construct generated residuals entering the second-step average. Treating these first-step estimates as fixed would understate the true variance by neglecting sampling uncertainty from the first step. A delta-method correction would require characterizing the full covariance structure of $T \times 3$ first-step estimators—a high-dimensional object that is difficult to estimate reliably in finite samples, particularly given the cross-period correlation induced by the panel structure. Moreover, the inverse-magnitude weighting $|\Delta z_{it}|^{-1}$ creates heteroskedastic amplification that depends on the same variable driving the first-step regressions, further complicating analytical variance estimation. The cluster-level bootstrap naturally accounts for all these sources of uncertainty—including within-country serial corre-

lation, heteroskedastic weighting, and first-step sampling error—without requiring explicit modeling of any of them, and is asymptotically valid under the regularity conditions of de Chaisemartin et al. (2024).

C Upstream/Downstream Estimator: Multi-Period Extension and Model-based Approach

While we present the typical upstream/downstream estimator (13) in a two-period setting in the main paper, in practice, researchers use the following regression:

$$\Delta \ln y_{it} = \beta^{Own} \Delta z_{it} + \beta^{Upstream} \sum_{j \neq i} h_{ij,0} \Delta z_{jt} + \beta^{Downstream} \sum_{k \neq i} h_{ki,0} \Delta z_{kt} + \alpha_i^{UD} + \xi_{it}, \quad (\text{C.1})$$

in pooled first difference with multiple periods, where $h_{ij,0}$ indicates the upstream exposure of unit i to unit j (j sells to i) measured at time $t = 0$, and $h_{ki,0}$ indicates the downstream exposure of unit i to unit k (k buys from i) measured at time $t = 0$. This is also the version that we use in our Monte Carlo simulations and in our empirical application.

Das et al. (2022), Feng et al. (2025), and Zappalà (2024) use an “upstream/downstream” estimator to estimate spillovers across countries. None of these papers derive the estimation equation explicitly from a model, nor does any other paper we are aware of; but Zappalà (2024) discusses an open-economy extension to Acemoglu et al. (2016) to motivate measuring exposure with the Leontief inverse matrix of trade shares. Nevertheless, without a formal model, it is hard to know how to measure exposure.

In this appendix, we extend Acemoglu et al. (2016) to an open-economy setting to see if the model can inform how to construct the exposure measures. We start with a restrictive case that imposes Cobb-Douglas functions for supply and demand. We then extend to general functional forms.

C.1 Open Economy extension to Acemoglu et al. (2016)

Environment. There are $i = 1, \dots, N$ countries and $s = 1, \dots, S$ sectors. We treat each country-sector (is) as a node. All countries produce and consume in all sectors. There is a representative firm in each country-sector node. Competition is perfect.

Labor is the only primary factor. Country i has exogenous labor supply L_i . Labor is mobile across the sectors, implying a common wage w_i . Labor markets clear $L_i = \sum_s L_{is}$.

Demand. In each country i , the representative household has Cobb-Douglas preferences over the rm goods:

$$U_i = \prod_{rm \in \Theta} (C_i^{rm})^{\beta_i^{rm}} \quad (\text{C.2})$$

with C_i^{rm} indicating the quantity of good rm consumed by country i , and $\sum_{rm \in \Theta} \beta_i^{rm} = 1$, and Θ is the set of all country-sector nodes. Given nominal income $Y_i = w_i L_i$, final good expenditure of country i on good rm is

$$X_i^{rm} = \beta_i^{rm} Y_i = \beta_i^{rm} w_i L_i. \quad (\text{C.3})$$

Consumers in i pay price p_i^{rm} for good produced in rm . We assume iceberg trade costs, so $p_i^{rm} = p_{rm} \tau_i^{rm}$ where p_{rm} is the factory-gate price of good rm and $\tau_i^{rm} \geq 1$ is the iceberg cost.

Production. Output produced in country i in sector s (i.e., at node is) is Cobb-Douglas in labor and the material composite with constant returns to scale:

$$Q_{is} = A_{is} L_{is}^{\eta_{is}} \prod_{rm \in \Theta} (M_{is}^{rm})^{\alpha_{is}^{rm}} \quad (\text{C.4})$$

where $0 < \eta_{is} < 1$ is the labor share, A_{is} is the Hicks-neutral productivity, and L_{is} is the labor allocation, M_{is}^{rm} are the physical quantities of country-sector rm purchased by country-sector is , and $\eta_{is} + \sum_{rm \in \Theta} \alpha_{is}^{rm} = 1$.

Market clearing. Market clearing for each good is implies total nominal output equals the sum of final demand in all countries and intermediate demand from all nodes:

$$Y_{is} = \sum_j X_j^{is} + \sum_{rm \in \Theta} p_i^{rm} M_{is}^{rm} = \sum_j \beta_j^{is} w_j L_j + \sum_{rm \in \Theta} \alpha_{rm}^{is} Y_{rm} \quad (\text{C.5})$$

where the second equality uses the first order conditions under Cobb-Douglas functions.

Nominal Revenue Response. In this Cobb-Douglas model, we can solve for nominal revenue of each node explicitly. It turns out that nominal revenue in each node is fixed, i.e. unrelated to productivity. This in turn implies that w_i and L_{is} are also fixed. To see this, substitute $w_j L_j = w_j \sum_{h=1}^S L_{jh} = \sum_{h=1}^S \eta_{jh} Y_{jh}$ into (C.5)

$$Y_{is} = \sum_j \beta_j^{is} \sum_{h=1}^S \eta_{jh} Y_{jh} + \sum_{rm \in \Theta} \alpha_{rm}^{is} Y_{rm} \quad (\text{C.6})$$

which can be written in matrix form:

$$\mathbf{Y} = \mathbf{B}\mathbf{Y} + \mathbf{M}\mathbf{Y} \quad (\text{C.7})$$

with

$$\mathbf{Y} \equiv \begin{pmatrix} Y_{11} \\ \vdots \\ Y_{1S} \\ Y_{21} \\ \vdots \\ Y_{NS} \end{pmatrix} \in \mathbb{R}^{NS}, \quad \mathbf{B} = \begin{pmatrix} B_{11} & B_{12} & \cdots & B_{1N} \\ B_{21} & B_{22} & \cdots & B_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ B_{N1} & B_{N2} & \cdots & B_{NN} \end{pmatrix}, \quad B_{ij} \in \mathbb{R}^{S \times S},$$

where each block B_{ij} has entries $[B_{ij}]_{s,h} = \beta_j^{is} \eta_{jh}$, for $s, h = 1, \dots, S$ and

$$\mathbf{M} = \begin{pmatrix} M_{11} & M_{12} & \cdots & M_{1N} \\ M_{21} & M_{22} & \cdots & M_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ M_{N1} & M_{N2} & \cdots & M_{NN} \end{pmatrix}, \quad \mathbf{M}_{ir} \in \mathbb{R}^{S \times S},$$

with each block $[M_{ir}]_{s,m} = \alpha_{rm}^{is}$, for $s, m = 1, \dots, S$. Equation (C.7) defines a homogeneous linear system. Because every column of $\mathbf{M} + \mathbf{B}$ sums to one, the matrix $\mathbf{M} + \mathbf{B}$ is column-stochastic, and therefore has a Perron–Frobenius eigenvalue equal to one. Equation (C.7) therefore characterizes \mathbf{Y} as a right eigenvector of the matrix $\mathbf{M} + \mathbf{B}$ associated with eigenvalue 1. Such an eigenvector exists, is unique up to scale, and has strictly positive components. Crucially, both \mathbf{M} and \mathbf{B} depend only on the Cobb–Douglas expenditure and production parameters; they do *not* depend on the productivity levels A_{is} . Hence the Perron–Frobenius eigenvector of $\mathbf{M} + \mathbf{B}$ —and therefore the entire pattern of nominal outputs \mathbf{Y} —is independent of productivity. Productivity shocks cannot change the relative solution to (C.7).

To determine the level of nominal output (i.e. to fix the scale of the eigenvector), a normalization is required, such as specifying world nominal income or fixing one country’s nominal GDP. But this normalization is likewise independent of productivity in the Cobb–Douglas model. Therefore, both the pattern and the level of nominal outputs are invariant to changes in A_{is} . As a consequence, $w_i L_{is}$ must be fixed because $w_i L_{is} = \eta_{is} Y_{is}$. Thus, we

also have w_i is fixed because $\sum_s w_i L_{is} = w_i L_i$. If the LHS is fixed, and L_i is fixed, then w_i must be fixed. Finally, if w_i is fixed and Y_{is} , then L_{is} is also fixed. These results will be used to determine the real output effects.

Real Output Response Totally differentiating the production function and substituting in with first order conditions for M_{is}^{rm} and L_{is} yields:

$$d \ln Q_{is} = d \ln A_{is} + \eta_{is} \left(d \ln Q_{is} + d \ln P_{is} - d \ln w_i \right) + \sum_{rm \in \Theta} \alpha_{is}^{rm} \left(d \ln Q_{is} + d \ln P_{is} - d \ln P_{rm} \right)$$

and combining terms:

$$-d \ln P_{is} = d \ln A_{is} - \eta_{is} d \ln w_i + \sum_{rm \in \Theta} \alpha_{is}^{rm} \left(-d \ln P_{rm} \right)$$

We have, for every node, $d \ln Y_{rm} = d \ln Q_{rm} + d \ln P_{rm}$. Since nominal output Y_{rm} is fixed, then $d \ln Q_{rm} = -d \ln P_{rm}$. Also, $d \ln w_i = 0$. Hence,

$$d \ln Q_{is} = d \ln A_{is} + \sum_{rm \in \Theta} \alpha_{is}^{rm} \left(d \ln Q_{rm} \right), \quad (\text{C.8})$$

which implies

$$d \ln \mathbf{Q} = \left(\mathbf{I} - \mathbf{M} \right)^{-1} d \ln \mathbf{A}. \quad (\text{C.9})$$

Since the spectral radius of matrix \mathbf{M} is less than one, we can invert $\mathbf{I} - \mathbf{M}$ and denoting by $\mathbf{L} \equiv (\mathbf{I} - \mathbf{M})^{-1}$, we observe that $\mathbf{L} = \sum_{k=0}^{\infty} \mathbf{M}^k$, which reflects not only the direct input-output linkage, but also the indirect ones due to the Leontieff inverse. This is the same result as in [Acemoglu et al. \(2016\)](#), extended to an open economy setting. We note that productivity shocks only affect real output, while nominal output, wages and incomes are fixed, in this setting. Hence, this model cannot accommodate cross-unit interdependencies that relate to income shocks, competition or sectoral reallocation. By construction, these mechanisms are shut down.

C.2 General Model

In this section, we keep the same environment (constant returns to scale technologies and perfect competition), but leave the utility function and the production function general (instead of Cobb-Douglas).

Production. Output at node is is Hicks-neutral in productivity, but general in labor and the materials inputs:

$$Q_{is} = A_{is}F\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right) \quad (\text{C.10})$$

Firms are price takers. Profit maximization implies the firm solves the program:

$$\max_{L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}} \Pi_{is} = p_{is}A_{is}F\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right) - w_i L_{is} - \sum_{rm \in \Theta} p_{rm} M_{is}^{rm}. \quad (\text{C.11})$$

Given the first-order condition for labor, $p_{is}A_{is} \frac{\partial F_{is}}{\partial L_{is}}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right) = w_i$, multiplying both sides by $L_{is}/(p_{is}Q_{is})$ gives

$$\lambda_{is}^L \equiv \frac{w_i L_{is}}{p_{is} Q_{is}} = \frac{L_{is} \frac{\partial F_{is}}{\partial L_{is}}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right)}{F_{is}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right)} \equiv \varepsilon_{FL}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right) \quad (\text{C.12})$$

where ε_{FL} is the output elasticity of F_{is} with respect to labor, evaluated at the observed input vector $L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}$, and λ_{is}^L is the cost share in revenues of L , observable in the data.

Given first order conditions for each input (r, m) , $p_{is}A_{is} \frac{\partial F_{is}}{\partial M_{is}^{rm}}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right) = p_{rm}$, the cost-share identity is

$$\lambda_{is}^{rm} \equiv \frac{p_{rm} M_{is}^{rm}}{p_{is} Q_{is}} = \frac{M_{is}^{rm} \frac{\partial F_{is}}{\partial M_{is}^{rm}}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right)}{F_{is}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right)} \equiv \varepsilon_{F,rm}\left(L_{is}, M_{is}^{11}, \dots, M_{is}^{NS}\right), \quad (\text{C.13})$$

and λ_{is}^{rm} is the cost share in revenues of rm , observable in the data.

Taking the first order approximation to the production function yields

$$d \log Q_{is} = d \log A_{is} + \varepsilon_{FL} d \log L_{is} + \sum_{rm \in \Theta} \varepsilon_{F,rm} d \log M_{is}^{rm} \quad (\text{C.14})$$

and substituting with cost shares yields

$$d \log Q_{is} = d \log A_{is} + \lambda_{is}^L d \log L_{is} + \sum_{rm \in \Theta} \lambda_{is}^{rm} d \log M_{is}^{rm}. \quad (\text{C.15})$$

Market Clearing Market clearing implies

$$\begin{aligned} Y_{is} &= \sum_j X_j^{is} + \sum_{rm \in \Theta} p_r^{is} M_{rm}^{is} \\ &= \sum_j \pi_j^{Fis} \sum_{h=1}^S w_j L_{jh} + \sum_{rm \in \Theta} p_r^{is} M_{rm}^{is}, \end{aligned} \quad (\text{C.16})$$

where $\pi_j^{Fis} = X_j^{is}/w_j L_j$ is the share of expenditures on final consumption in country j on node is . Using the first order conditions for labor and materials yields

$$Y_{is} = \sum_j \pi_j^{Fis} \sum_{h=1}^S \lambda_{jh}^L Y_{jh} + \sum_{rm \in \Omega} \lambda_{rm}^{is} Y_{rm}. \quad (\text{C.17})$$

Taking a first order approximation generates

$$d \log Y_{is} = \sum_j \gamma_{is}^{Fj} \left(d \log \pi_j^{Fis} + \sum_{h=1}^S \left(d \log \lambda_{jh}^L + d \log Y_{jh} \right) \right) + \sum_{rm \in \Theta} \gamma_{is}^{rm} \left(d \log \lambda_{is}^{rm} + d \log Y_{rm} \right) \quad (\text{C.18})$$

where $\gamma_{is}^{Fj} \equiv \frac{X_j^{is}}{Y_{is}}$ is the export share of is for final consumption in j , and $\gamma_{is}^{rm} \equiv \frac{X_{rm}^{is}}{Y_{is}}$ is the export share of is for intermediate inputs used by rm . Equation (C.18) defines a homogeneous linear system, as in the Cobb-Douglas case. But in this formulation, the variables are equilibrium outcomes, not exogenous parameters, and hence depend on the realization of productivity.

Real Output Response. Rewriting (C.15),

$$\begin{aligned} d \log Q_{is} &= d \log A_{is} + \lambda_{is}^L d \log v_{is} + \sum_{rm \in \Theta} \lambda_{is}^{rm} \left(d \log X_{is}^{rm} - d \log P_i^{rm} \right) \\ &= d \log A_{is} + \lambda_{is}^L d \log v_{is} + \sum_{rm \in \Theta} \lambda_{is}^{rm} \left(d \log Q_{rm} + d \log \tau_i^{rm} + d \log \gamma_{is}^{rm} \right) \end{aligned} \quad (\text{C.19})$$

where $v_{is} \equiv \frac{L_{is}}{L_i}$, i.e. the share of labor in i working in sector s and γ_{is}^{rm} is defined above.

Solving and putting in matrix notation yields

$$d \ln \mathbf{Q} = (\mathbf{I} - \mathbf{M})^{-1} d \ln \mathbf{A} + (\mathbf{I} - \mathbf{M})^{-1} \left(\boldsymbol{\lambda}^L d \ln \mathbf{v} + \mathbf{M} d \ln \boldsymbol{\Gamma} + \mathbf{M} d \ln \boldsymbol{\tau} \right) \quad (\text{C.20})$$

where $\boldsymbol{\Gamma}(\boldsymbol{\tau})$ is the matrix of export shares (iceberg cost) with elements γ_{rm}^{is} (τ_i^{rm}), and \mathbf{M} is the matrix of import shares, same as above for the Cobb-Douglas case.

The change in real output depends on the change in productivity (first term on right hand side) and a change in allocation (second term on right hand side). The change in the allocation is endogenous. The model cannot be solved in terms of just exogenous components without more structure.

Nominal Output Response. Substituting in Y_{is} into (C.15)

$$d \log Y_{is} = d \log P_{is} + d \log A_{is} + \lambda_{is}^L d \log v_{is} + \sum_{rm \in \Theta} \lambda_{is}^{rm} \left(d \log Y_{rm} - d \log P_i^{rm} + d \log \gamma_{is}^{rm} \right) \quad (\text{C.21})$$

Now totally differentiating the unit cost function, and substituting with the first order condition for labor yields

$$d \log P_{is} - \sum_{rm \in \Theta} \lambda_{is}^{rm} d \log P_i^{rm} = -d \log A_{is} + \lambda_{is}^L \left(d \log Y_{is} + d \log \lambda_{is}^L - d \log v_{is} \right) \quad (\text{C.22})$$

Substituting this expression into (C.21) yields

$$d \log Y_{is} = \lambda_{is}^L d \log Y_{is} + \lambda_{is}^L d \log \lambda_{is}^L + \sum_{rm \in \Theta} \lambda_{is}^{rm} d \log Y_{rm} + \sum_{rm \in \Theta} \lambda_{is}^{rm} d \log \gamma_{is}^{rm} \quad (\text{C.23})$$

rearranging and writing in matrix form gives

$$d \ln \mathbf{Y} = \left(\mathbf{I} - \mathbf{M} \right)^{-1} \left(\frac{\boldsymbol{\lambda}^L}{1 - \boldsymbol{\lambda}^L} d \log \boldsymbol{\lambda}^L + \frac{1}{1 - \boldsymbol{\lambda}^L} \mathbf{M} d \ln \boldsymbol{\Gamma} \right) \quad (\text{C.24})$$

In this more general model that maintains the constant returns to scale assumption, we have an effect on nominal output through reallocation.

C.3 Summary

We can extend the framework from [Acemoglu et al. \(2016\)](#) to open economies. However, to do so, we have to treat each country-sector as a node. This yields a very restrictive assumption compared to standard trade models, in which specific goods enter the production function without regard for origin ([Eaton & Kortum, 2002](#); [Caliendo & Parro, 2015](#)). Nevertheless, adopting such a framework, imposing Cobb-Douglas production and consumption in nodes, extends the results from [Acemoglu et al. \(2016\)](#) to open economies: real output can be modeled as a function of the Leontief inverse matrix of input-output linkages and the productivity shocks in each country-sector; and nominal outputs (as well as income and wages) are invariant to productivity shocks.

In a slightly more general model in which country-sectors are still treated as nodes

but supply and demand functions are left general, the results from [Acemoglu et al. \(2016\)](#) do not hold. Both real output and nominal output in each country-sector depend on the allocation of labor across sectors and on the distribution of consumption across suppliers. These allocation terms are endogenous to productivity shocks. To express real or nominal output as a function of these productivity shocks, more structure is required. As a result, in the general model, there is no guidance with respect to how to construct exposure in the upstream/downstream estimator.

D Structural Model: Equilibrium Conditions and Counterfactuals

This appendix derives the equilibrium system reported in equations (20)-(24) in the main text, and then shows that the solution to the exact hat system (32)-(35) is invariant to normalization.

D.1 Deriving Equilibrium Conditions

Labor allocation across sectors (22). Firms minimize costs, which implies that labor payments are a constant share of revenue:

$$w_{it}L_{ist} = \eta_{is}Y_{ist}. \quad (\text{D.1})$$

Sectoral wage equalization within country i then yields the equation (22) in the main text:

$$\frac{L_{ist}}{L_{it}} = \frac{\eta_{is}Y_{ist}}{\sum_{h=1}^S \eta_{ih}Y_{iht}}. \quad (\text{D.2})$$

Sectoral expenditures (21). Let X_{ist}^F denote final consumption spending by country i on sector- s goods. Cobb–Douglas preferences across sectors imply

$$X_{ist}^F = \alpha_{is}^C \left(w_{it}L_{it} + \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} X_{ikh} \right), \quad (\text{D.3})$$

where household income equals labor income plus rebated tariff revenue. Under perfect competition and constant returns to scale, profits are zero, so labor income equals the sum of wage bills across sectors: $w_{it}L_{it} = \sum_{h=1}^S w_{it}L_{iht}$. Using the constant labor cost share condition,

$$w_{it}L_{iht} = \eta_{ih}Y_{iht} \quad \Rightarrow \quad w_{it}L_{it} = \sum_{h=1}^S \eta_{ih}Y_{iht}, \quad (\text{D.4})$$

we obtain

$$X_{ist}^F = \alpha_{is}^C \left(\sum_{h=1}^S \eta_{ih} Y_{iht} + \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} X_{ikh} \right). \quad (\text{D.5})$$

Intermediate demand is derived as follows. Sector h in country i spends a share $(1 - \eta_{ih})$ of its revenue on intermediate inputs, allocated across sectors with shares α_{ihs}^M . Hence intermediate spending by country i on sector s equals

$$X_{ist}^M = \sum_{h=1}^S \alpha_{ihs}^M (1 - \eta_{ih}) Y_{iht}. \quad (\text{D.6})$$

Total expenditures in country i on sector s are therefore

$$X_{ist} = X_{ist}^F + X_{ist}^M = \sum_{h=1}^S \left(\eta_{ih} \alpha_{is}^C + (1 - \eta_{ih}) \alpha_{ihs}^M \right) Y_{iht} + \alpha_{is}^C \sum_{h=1}^S \sum_{k=1}^N \text{tariff}_{ikh} X_{ikh}, \quad (\text{D.7})$$

which is equation (21) in the main text.

Trade shares. Given Cobb–Douglas production, unit cost in country i , sector s , at time t is

$$c_{ist} \equiv w_{it}^{\eta_{is}} \prod_{h=1}^S p_{iht}^{\alpha_{ish}^M (1 - \eta_{is})}. \quad (\text{D.8})$$

Under iceberg trade costs, the delivered unit price in destination n of a variety sourced from i is

$$p_{nist}(j) = \frac{c_{ist} \tau_{nist}}{v_{ist}(j)}, \quad (\text{D.9})$$

where $v_{ist}(j)$ is an i.i.d. Fréchet draw with distribution $F_{ist}(v) = \exp\{-A_{ist} v^{-\theta_s}\}$.

Standard Eaton–Kortum results imply that the sector- s price distribution in importer n has cdf $G_{nst}(p) = 1 - \exp\{-\Phi_{nst} p^{\theta_s}\}$, where

$$\Phi_{nst} = \sum_{k=1}^N A_{knt} (c_{knt} \tau_{nknt})^{-\theta_s}. \quad (\text{D.10})$$

The CES price index satisfies

$$p_{nst} = \gamma_s \Phi_{nst}^{-1/\theta_s}, \quad \gamma_s \equiv \left[\Gamma \left(\frac{\theta_s + 1 - \sigma_s}{\theta_s} \right) \right]^{\frac{1}{1 - \sigma_s}}. \quad (\text{D.11})$$

Finally, the probability that country i is the lowest-cost supplier to n in sector s equals the

trade share

$$\pi_{nist} \equiv \frac{X_{nist}}{X_{nst}} = \frac{A_{ist} c_{ist}^{-\theta_s} \phi_{nist}}{\Phi_{nst}}, \quad (\text{D.12})$$

where bilateral accessibility is defined as $\phi_{nist} \equiv \tau_{nist}^{-\theta_s}$.

Structural gravity and multilateral resistance (23)-(24). Using (D.12), bilateral trade flows satisfy

$$X_{nist} = \pi_{nist} X_{nst} = \frac{A_{ist} c_{ist}^{-\theta_s}}{\Phi_{nst}} \phi_{nist} X_{nst}. \quad (\text{D.13})$$

Summing over destinations yields market clearing for exporter (i, s):

$$Y_{ist} = \sum_{n=1}^N X_{nist} = A_{ist} c_{ist}^{-\theta_s} \underbrace{\sum_{k=1}^N \frac{X_{kst}}{\Phi_{kst}} \phi_{kist}}_{\equiv \Omega_{ist}}. \quad (\text{D.14})$$

Substituting into (D.13) yields the gravity equation, which is equation (23) in the main text:

$$X_{nist} = \frac{Y_{ist}}{\Omega_{ist}} \frac{X_{nst}}{\Phi_{nst}} \phi_{nist}. \quad (\text{D.15})$$

Adding up import shares implies

$$\Phi_{nst} = \sum_{k=1}^N \frac{Y_{kst}}{\Omega_{kst}} \phi_{nkst}, \quad (\text{D.16})$$

while exporter market clearing implies

$$\Omega_{ist} = \sum_{k=1}^N \frac{X_{kst}}{\Phi_{kst}} \phi_{kist}. \quad (\text{D.17})$$

Equations (D.16) and (D.17) jointly correspond to (24). Multilateral resistance terms are defined up to normalization; throughout we adopt the normalization implicit in these equations.

Revenue equation (20). Substituting unit cost and price indices into (D.14) yields

$$Y_{ist} = A_{ist} w_{it}^{-\theta_s \eta_{is}} \left(\prod_{h=1}^S \gamma_h^{-\theta_s} \alpha_{ish}^M (1 - \eta_{is}) \right) \prod_{h=1}^S \Phi_{iht}^{\alpha_{ish}^M \frac{\theta_s}{\theta_h} (1 - \eta_{is})} \Omega_{ist}. \quad (\text{D.18})$$

Using $w_{it}L_{ist} = \eta_{is}Y_{ist}$ and collecting terms yields

$$Y_{ist}^{1+\theta_s\eta_{is}} = \left[\eta_{is}^{-\theta_s\eta_{is}} \prod_{h=1}^S \gamma_h^{-\theta_s\alpha_{ish}^M(1-\eta_{is})} \right] A_{ist} L_{ist}^{\theta_s\eta_{is}} \prod_{h=1}^S \Phi_{iht}^{\alpha_{ish}^M \frac{\theta_s}{\theta_h}(1-\eta_{is})} \Omega_{ist}. \quad (\text{D.19})$$

Taking both sides to the power $1/(1+\theta_s\eta_{is})$ and substituting the parametric form for A_{ist} given in (17) yields equation (24) in the main text, with constants absorbed into the fixed and unobserved components.

D.2 Exact hat algebra with normalized trade frictions

This appendix shows that the normalizations used to identify trade frictions and multilateral resistance terms do not affect counterfactual outcomes. In particular, we show that the exact hat algebra system used to compute counterfactuals, equations (32)-(35) in the main text, is invariant to the (unknown) normalization terms that enter the estimated objects $\check{\phi}_{nist}$, $\check{\Phi}_{nst}$, and $\check{\Omega}_{ist}$.

Step 1: Normalized objects and their relation to the true ones. From the gravity estimation step, bilateral trade frictions are identified only up to normalization, and we work with the normalized frictions $\check{\phi}_{nist}$. The associated multilateral resistance terms $\check{\Phi}_{nst}$ and $\check{\Omega}_{ist}$ solve the system (23)-(24) when ϕ is replaced by $\check{\phi}$. Under the normalization described in the text, the estimated (multilateral resistance) objects relate to the true ones according to

$$\check{\Phi}_{nst} = \frac{\Phi_{nst}}{\Phi_{1st}} \cdot \frac{\phi_{nst}^e}{\phi_{1st}^e} \cdot \frac{\phi_{11st}}{\phi_{nmst}}, \quad \check{\Omega}_{ist} = \frac{\Omega_{ist}}{\Phi_{1st}} \cdot \frac{\phi_{11st}}{\phi_{1st}^e} \cdot \frac{1}{\phi_{ist}^e}, \quad (\text{D.20})$$

where $\phi_{nst}^e \equiv \exp(\frac{1}{N} \sum_k \ln \phi_{knst})$ and similarly for ϕ_{ist}^e .

The key observation is that (D.20) is *multiplicative* and can be written compactly as

$$\check{\Phi}_{nst} = a_{st} b_{nst} \Phi_{nst}, \quad \check{\Omega}_{ist} = a_{st} f_{ist} \Omega_{ist}, \quad (\text{D.21})$$

for some (unknown) positive scalars a_{st} common to all countries in sector s at time t , and positive scalars b_{nst} and f_{ist} that are importer- and exporter-specific, respectively. (Their exact expressions are implied by (D.20) and are not needed below.)

Step 2: Hat algebra for multilateral resistance terms is invariant. Let $\hat{x} \equiv x'/x$ denote proportional changes from the baseline equilibrium to a counterfactual. The exact hat algebra system for the inward multilateral resistance term (32) can be derived from equation

(24) as follows. Starting from

$$\Phi_{nst} = \sum_k \frac{Y_{kst}}{\Omega_{kst}} \phi_{nkst},$$

take the ratio between counterfactual and baseline:

$$\widehat{\Phi}_{nst} = \frac{\sum_k \frac{Y'_{kst}}{\Omega'_{kst}} \phi_{nkst}}{\sum_k \frac{Y_{kst}}{\Omega_{kst}} \phi_{nkst}} = \frac{\sum_k \widehat{Y}_{kst} \frac{Y_{kst}}{\Omega_{kst}} \cdot \widehat{\Omega}_{kst}^{-1} \cdot \phi_{nkst}}{\sum_k \frac{Y_{kst}}{\Omega_{kst}} \phi_{nkst}}. \quad (\text{D.22})$$

Now replace all unobserved objects by their normalized counterparts. Using (D.21), we have

$$\frac{Y_{kst}}{\check{\Omega}_{kst}} = \frac{Y_{kst}}{a_{st} f_{kst} \Omega_{kst}} = \frac{1}{a_{st}} \cdot \frac{1}{d_{kst}} \cdot \frac{Y_{kst}}{\Omega_{kst}}.$$

Substituting this into (D.22), the common factor $1/a_{st}$ cancels between the numerator and denominator, leaving an expression that depends only on *ratios* and therefore does not depend on the normalization. After rearrangement, we obtain the operational expression reported in equation (32):

$$\widehat{\Phi}_{nst} = \frac{\sum_k \widehat{Y}_{kst} \widehat{\Omega}_{kst}^{-1} \left(\frac{Y_{kst}}{\check{\Omega}_{kst}} \right) \check{\phi}_{nkst}}{\sum_k \left(\frac{Y_{kst}}{\check{\Omega}_{kst}} \right) \check{\phi}_{nkst}} = \frac{\sum_k \left(\widehat{Y}_{kst} \widehat{\Omega}_{kst}^{-1} \right) \left(\frac{Y_{kst}}{\check{\Omega}_{kst}} \right) \check{\phi}_{nkst}}{\sum_k \left(\frac{Y_{kst}}{\check{\Omega}_{kst}} \right) \check{\phi}_{nkst}}, \quad (\text{D.23})$$

which matches the first part of (32).

The same argument applies to outward multilateral resistance. Starting from

$$\Omega_{ist} = \sum_k \frac{X_{kst}}{\Phi_{kst}} \phi_{kist},$$

taking ratios yields

$$\widehat{\Omega}_{ist} = \frac{\sum_k \widehat{X}_{kst} \widehat{\Phi}_{kst}^{-1} \left(\frac{X_{kst}}{\check{\Phi}_{kst}} \right) \check{\phi}_{kist}}{\sum_k \left(\frac{X_{kst}}{\check{\Phi}_{kst}} \right) \check{\phi}_{kist}}, \quad (\text{D.24})$$

which is the second part of (32). Again, all normalization factors cancel because they enter multiplicatively and appear in both the numerator and denominator.

Step 3: Hat algebra for outcomes (33)–(35). Given $\widehat{\Phi}$ and $\widehat{\Omega}$, the remaining counterfactual objects are obtained by taking ratios of the equilibrium conditions (20)–(22). Because those conditions are written in multiplicative form, any remaining constants (including the γ_h terms and the MR normalizations absorbed into fixed effects in estimation) cancel out in

proportional changes.

Revenues. Taking the ratio of (20) between counterfactual and baseline yields equation (33):

$$\widehat{Y}_{ist} = \exp\left(\sum_{v=1}^V \frac{\mu_s^v}{1 + \theta_s \eta_{is}} ((z_{it}^v)' - z_{it}^v)\right) \widehat{L}_{ist}^{\frac{\theta_s \eta_{is}}{1 + \theta_s \eta_{is}}} \widehat{\Omega}_{ist}^{\frac{1}{1 + \theta_s \eta_{is}}} \prod_{h=1}^S \widehat{\Phi}_{iht}^{\alpha_{ish}^M \frac{\theta_s}{\theta_h} \frac{(1 - \eta_{is})}{1 + \theta_s \eta_{is}}}. \quad (\text{D.25})$$

Expenditures. Taking ratios of (21) yields equation (34):

$$\widehat{X}_{ist} = \sum_{h=1}^S \left(\alpha_{is}^C \eta_{ih} + (1 - \eta_{ih}) \alpha_{ihs}^M \right) \frac{Y_{iht}}{X_{ist}} \widehat{Y}_{iht} + \alpha_{is}^C \sum_{h=1}^S \sum_{k=1}^N \frac{\text{tariff}_{ikh} X_{ikht}}{X_{ist}} \widehat{X}_{ikht}, \quad (\text{D.26})$$

Labor allocation. Taking ratios of (22) yields equation (35):

$$\widehat{L}_{ist} = \widehat{L}_{it} \cdot \widehat{Y}_{ist} \cdot \frac{\sum_h \eta_{ih} Y_{iht}}{\sum_h \eta_{ih} Y'_{iht}} = \widehat{L}_{it} \cdot \frac{\eta_{is} Y_{ist}}{\sum_h \eta_{ih} Y_{iht}} \frac{\widehat{Y}_{ist}}{\sum_h \eta_{ih} Y_{iht} \widehat{Y}_{iht} / \sum_h \eta_{ih} Y_{iht}}. \quad (\text{D.27})$$

E Alternative procedure to estimate trade elasticities via GMM

The estimation procedure described in section 3.2 relies on tariff data to identify the trade elasticities, which are then used as inputs into the estimation of μ_s^v . Although the matrix of bilateral tariff rates is far from complete, there are sufficient observations and variation to identify trade elasticities over the last 30-40 years.

If, however, one wanted to quantify the effects of observed factors for a time period for which tariff data was not available, or for a region of the world for which tariff information does not exist or are all 0, it is possible to amend the procedure to identify the trade elasticities purely from variation in the observed factor, e.g., climate. This, naturally, places greater demands on the observed factor to offer sufficient variation to identify this additional parameter.

For a given guess of $\boldsymbol{\theta}^* = \{\theta_1^*, \theta_2^*, \dots, \theta_S^*\}$, $\ln \xi_{ist}(\boldsymbol{\theta}^*)$ can be computed from (30):

$$\begin{aligned} \ln \xi_{ist}(\boldsymbol{\theta}^*) &= \ln Y_{ist} - \frac{(\theta_s)^* \eta_{i,s}}{1 + (\theta_s)^* \eta_{i,s}} \ln L_{ist} - \frac{1}{1 + (\theta_s)^* \eta_{i,s}} \ln \check{\Omega}_{ist} \\ &\quad - \frac{(\theta_s)^* (1 - \eta_{i,s})}{1 + (\theta_s)^* \eta_{i,s}} \sum_{h=1}^S \frac{\alpha_{i,s,h}^M}{(\theta_h)^*} \ln \check{\Phi}_{iht} \end{aligned} \quad (\text{E.1})$$

Taking first differences, from (31) we have

$$\Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) = \delta_i + \delta_s^z \Delta z_{it} + \sum_h (\eta_{i,s} - 1) \alpha_{i,s,h}^M \times \delta_{st} + \Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) \quad (\text{E.2})$$

where δ_{st} absorbs the normalization term $\ln\left(\frac{\Phi_{1ht}^e \phi_{1ht}^e}{\Phi_{11st}}\right)$ and $\Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) = \Delta \omega_{ist} - \Delta \ln \phi_{ist}^e + \sum_h \theta_s (1 - \eta_{i,s}) \alpha_{i,s,h}^M \left(\Delta \ln\left(\frac{\phi_{iht}^e}{\phi_{iist}^e}\right)\right)$. Estimating (E.2) by OLS, including controls for expected Δz_{it} , we can extract $\Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*)$.

At the true values of $\boldsymbol{\theta}$, $\Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*)$ is orthogonal to Δz_{jt} for $j \neq i$. Hence, the moment conditions

$$E \left\{ \begin{pmatrix} \Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) \times Z_{ist}^\Omega \\ \Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) \times Z_{i1t}^\Phi \\ \Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) \times Z_{i2t}^\Phi \\ \cdot \\ \cdot \\ \cdot \\ \Delta \widetilde{\xi}_{ist}(\boldsymbol{\theta}^*) \times Z_{iSt}^\Phi \end{pmatrix} \right\} = 0$$

can be used to identify $\boldsymbol{\theta}$, where instruments $Z_{ist}^\Omega, Z_{i1t}^\Phi, \dots, Z_{iSt}^\Phi$ are weighted averages of trading partner shocks Δz_{jt} , with weights equal to $t - 1$ period sector-specific trade shares.

In our application, weather shocks do not offer sufficient variation for this GMM estimation procedure to converge. We therefore rely on a different procedure that uses tariff data to estimate $\boldsymbol{\theta}$. We offer this GMM estimator for researchers who may have reason to use it in lieu of tariff data, perhaps for studying other types of shocks or a different time period.

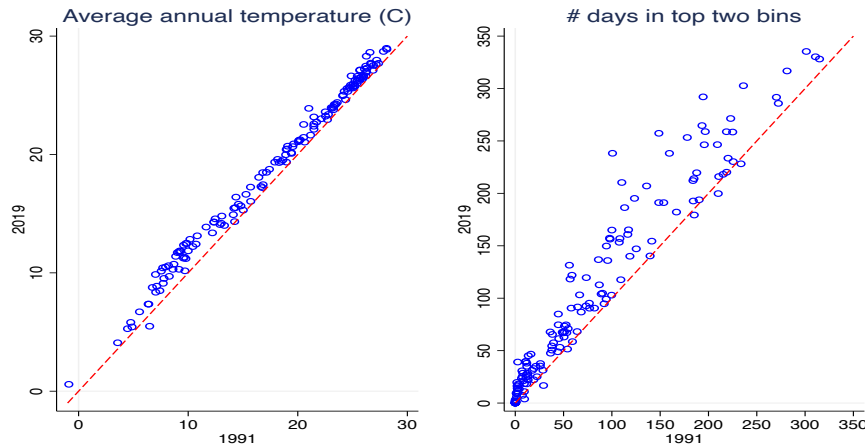
F Data and Descriptive Statistics

In addition to the details provided in the main text, below we document some data construction steps that are not crucial to the understanding of the paper and we report descriptive statistics.

Trade and Production Data. Mayer et al. (2023) and Fontagné et al. (2023) compile square matrices of bilateral trade flows for each sector, including self-trade. Domestic sales are calculated as the difference between the value of production and total exports. When the sum of total exports exceeds the value of production, they set the value to missing and extrapolate. This affects less than 1% of observations.

Weather Data. We use temperature data from the global reanalysis ERA-5 dataset compiled by the European Centre for Medium-Range Weather Forecasts (Hersbach et al., 2023), aggregated at the annual level by country. Figure F.1 depicts average annual temperature in degrees Celsius (left), and the number of days with maximum temperature above 30°C (right), at the beginning (x-axis) and at the end of the sample (y-axis). Almost all countries warmed between 1991 and 2019. The average growth in annual temperature was 1.2°C , and the average increase in the number of days in the top two bins was 22—a substantial rightward shift in the climate distribution over the period.

Figure F.1: Warming between 1991-2019



Notes: Figure plots average temperature in $^{\circ}\text{C}$ (left) and the number of days with maximum temperature above 30°C (right) in the year 2019 (y-axis) against the year 1991 (x-axis). Each dot corresponds to a country.

Table F.1 shows significant variation in weather variables year-to-year over the period. The extreme temperature bins seem to be the most volatile: two thirds of the observations deviate more than 75% from the country mean for the $(-\infty, 0^{\circ}\text{C}]$ bin, and more than 10% of the observations for the $(30^{\circ}\text{C}, +\infty)$ bin. Conversely, the middle bin $(0^{\circ}\text{C}, 30^{\circ}\text{C}]$ is more stable over time for a given country. As is standard in this literature, we exploit these deviations from country averages over time to identify effects.

Gravity Data. The gravity variables come from the Gravity database developed by CEPII (Conte et al., 2022), which we downloaded in the version V202211.

Tariff Data. We use bilateral tariff information from the MAcMap-HS6 dataset (Fontagné et al., 2022), which provides applied *ad valorem* equivalent tariffs (including non-tariff barriers) at the HS6 level for importer-exporter pairs in benchmark years 2001, 2004, 2007, 2010, 2013 and 2016. The MAcMap dataset imputes missing values using a reference

Table F.1: Observed variation in annual temperature measured with daily bins (1991-2019)

Proportion of country-year observations with number of days in each temperature bin [...] % above/below country mean					
	1%	10%	25%	50%	75%
$-\infty-0^{\circ}\text{C}$	0.99	0.90	0.76	0.67	0.61
$0-30^{\circ}\text{C}$	0.76	0.14	0.03	0.00	0.00
$30-+\infty^{\circ}\text{C}$	0.95	0.57	0.32	0.18	0.12

Reading: for the $(30^{\circ}\text{C}, +\infty^{\circ}\text{C}]$ variable, 95% of observations deviate more than 1% from the country mean while only 12% of observations deviate more than 75% from the country mean.

groups of countries. When, for a given importer–exporter–HS6 product triplet, either unit values or trade flows are missing, the bilateral value is replaced with the median unit-value or aggregate trade-flow of a reference group of “similar” countries (e.g., by income/trade openness classification). We aggregate the HS6-level tariffs to one manufacturing sector and one agricultural sector, using trade expenditures as weights.

Furthermore, to build counterfactuals, we need to know the tariff rates for all bilateral pairs and all goods in 2019, but MAcMap includes all countries from the BACI dataset, except for 10 destination countries and 2 origin countries, which are excluded due to infrequent reporting. For the missing information, we impute tariff rates as follows: For Fiji (FJI), Turkmenistan (TKM), Mongolia (MNG), Angola (AGO), Gambia (GMB), Cabo Verde (CPV), Singapore (SGP), and Serbia (SRB) we use values from TRAINS when possible. When the value from TRAINS is missing, we impute with the average import tariff for the specified importer-sector. For Palestine (PSE), we set the tariff rate to the tariff rate for Israel (ISR). For Iraq (IRQ), we set all import tariffs to 5%, following the approach used in World Bank trade briefs ([World Bank, 2004](#)).

GTAP Data. We compute production and consumption shares from the Global Trade Analysis Project (GTAP) Data Base version 9 ([Aguiar et al., 2016](#)). We consider domestic and import expenditures at purchaser’s price for the reference year 2011. The database contains information on 57 commodities for 116 countries and 24 aggregate regions. Table [F.2](#) reports the mapping between these commodities and three broad sectors: agriculture, manufacturing, and other. Table [F.3](#) reports the mapping between aggregate regions and their individual constitutive countries.

Labor Data. Table [F.4](#) reports the mapping between more disaggregated sectors from ILOEST and three broad sectors: agriculture, manufacturing, and other.

Table F.2: Aggregation of GTAP commodities into sectors

GTAP Sector	Aggregated Sector
Paddy rice, Wheat, Cereal grains nec	Agriculture
Vegetables, fruit, nuts	Agriculture
Oil seeds	Agriculture
Sugar cane, sugar beet	Agriculture
Plant-based fibers	Agriculture
Crops nec	Agriculture
Bovine cattle, sheep, goats, horses	Agriculture
Animal products nec	Agriculture
Raw milk	Agriculture
Wool, silk-worm cocoons	Agriculture
Forestry, Fishing	Agriculture
Coal, Oil, Gas, Minerals nec	Manufacturing
Meat: cattle, sheep, goats, horse	Manufacturing
Meat products nec	Manufacturing
Vegetable oils and fats	Manufacturing
Dairy products, Processed rice, Sugar	Manufacturing
Food products nec	Manufacturing
Beverages and tobacco products	Manufacturing
Textiles, Wearing apparel, Leather products	Manufacturing
Wood products	Manufacturing
Paper products, publishing	Manufacturing
Petroleum, coal products	Manufacturing
Chemical, rubber, plastic products	Manufacturing
Mineral products nec	Manufacturing
Ferrous metals, Metals nec, Metal products	Manufacturing
Motor vehicles and parts, Transport equipment nec	Manufacturing
Electronic equipment, Machinery and equipment nec	Manufacturing
Manufactures nec	Manufacturing
Electricity, Gas manufacture, distribution, Water	Other
Construction, Trade	Other
Transport nec, Sea transport, Air transport	Other
Communication, Financial services nec	Other
Insurance, Business services nec	Other
Recreation and other services	Other
Public Administration, Defense, Education, Health	Other
Dwellings	Other

Table F.3: Attribution of aggregate regions values to missing countries in GTAP

individual country	GTAP 9 Composite Region
Barbados	Rest of Caribbean
Central African Republic	Rest of Central Africa
Republic of the Congo	Rest of Central Africa
Belize	Rest of Central America
Burundi	Rest of Eastern Africa
Eritrea	Rest of Eastern Africa
Moldova	Rest of Eastern Europe
Bosnia and Herzegovina	Rest of Europe
Iceland	Rest of Europe
North Macedonia	Rest of Europe
Serbia	Rest of Europe
Tajikistan	Rest of Former Soviet Union
Turkmenistan	Rest of Former Soviet Union
Iraq	Rest of Western Asia
Lebanon	Rest of Western Asia
Palestine	Rest of Western Asia
Yemen	Rest of Western Asia
Algeria	Rest of North Africa
Fiji	Rest of Oceania
Suriname	Rest of South America
Maldives	Rest of South Asia
Angola	South Central Africa
Cape Verde	Rest of Western Africa
Gambia	Rest of Western Africa
Niger	Rest of Western Africa

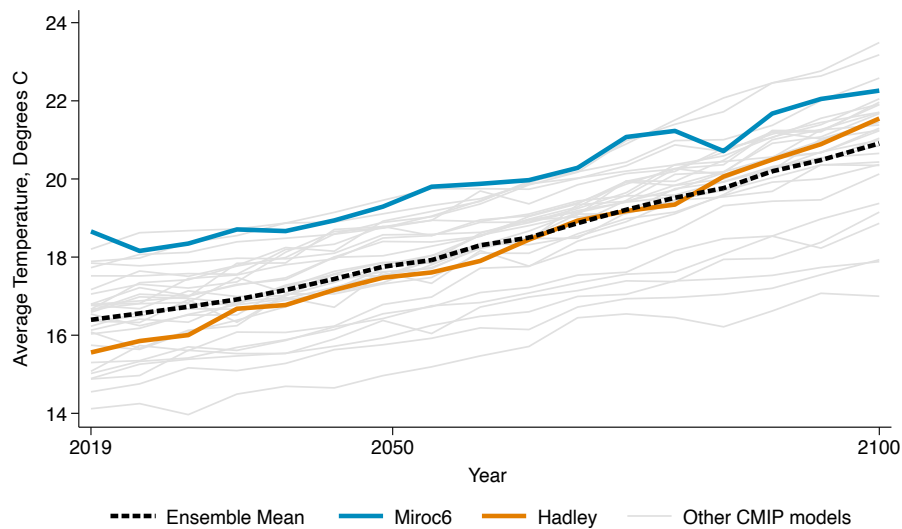
Table F.4: Aggregation of ILO industries into sectors

ILO industry	sector
Agriculture	Agriculture
Manufacturing	Manufacturing
Mining and quarrying	Manufacturing
Electricity, gas and water supply	Manufacturing
Construction	Other
Trade, Transportation, Accommodation and Food	Other
Business and Administrative Services	Other
Public Administration, Community, Social	Other
Other Services and Activities	Other

Descriptive statistics Table F.5 provides an overview of the key variables in our dataset. At the top of the table, the mean across countries of the population-weighted annual temperature was nearly 18°C over our study period, which is somewhat higher than the global mean temperature. The reason for this is that land is generally warmer than the oceans, which are naturally excluded from our calculation of country-level temperature.

Further down in Table F.5 we see that, although the value of gross agricultural output is generally much smaller than manufacturing output (11.67 vs. 208.13), it constitutes 15% of the intermediate input for the manufacturing sector on average ($\alpha_{Manuf,Ag}^M$), and more than 35% for one in twenty countries. Even if climate change is mainly expected to hit agricultural production in the first instance, these effects are likely to spill over to the far bigger manufacturing sector due to these intersectoral linkages.

Figure F.2: Average Global Warming Across CMIP6 Models



Notes: Figure plots global average annual temperature for 32 different models from the CMIP6. Global average temperature is computed as the simple average of pixels on dry land. Black dashed line indicates the simple average across all 32 models.

Averaging across all models, global average annual temperature increases from about 15.5°C to almost 20°C by 2100, or 4.5°C of global warming. There is significant amount of heterogeneity. In MIROC6 (in blue), global average annual temperature increases only 4°C, whereas the Hadley-3 (in orange) projects an increase of nearly 6°C.

Note that, because each CMIP model simulates the entire time path of a potential climate trajectory starting in mid-19th century, they do not agree even on past global mean temperatures. A model like MIROC6 exhibits less warming but a higher current temperature.

Table F.5: Descriptive Statistics

	Levels				Changes between 1991 and 2019 (%)			
	Mean	p5	p50	p95	Mean	p5	p50	p95
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Country-Level Variables</i>								
Mean temp. (°C)	17.78	6.02	19.05	27.22	9.05	1.52	5.51	25.59
Days above 30°C	85.06	0.00	53.48	253.68	57.98	0.00	35.33	203.74
Days below 0°C	12.42	0.00	0.00	68.45	-85.85	-196.75	-86.50	15.73
<i>L</i>	9.84	0.04	1.33	26.00	10.11	-67.01	4.11	87.33
<i>Y</i> _{Tot}	219.69	0.66	16.41	905.81	139.17	29.75	131.64	322.08
<i>Y</i> _{Ag}	11.67	0.04	1.72	39.04	105.45	-21.14	100.85	245.46
<i>Y</i> _{Manuf}	208.13	0.31	14.07	847.20	149.45	23.00	133.11	325.68
<i>X</i> _{Tot}	217.58	1.02	19.87	914.90	143.05	43.32	139.71	273.68
<i>X</i> _{Ag}	11.67	0.08	1.71	37.64	111.33	-10.75	107.64	258.24
<i>X</i> _{Manuf}	206.03	0.59	17.35	881.42	153.43	41.26	143.89	297.53
α_{Ag}^C	0.16	0.03	0.11	0.41				
α_{Manuf}^C	0.84	0.59	0.89	0.97				
η_{Ag}	0.30	0.16	0.28	0.47				
η_{Manuf}	0.13	0.05	0.12	0.22				
$\alpha_{Ag,Ag}^M$	0.39	0.12	0.39	0.66				
$\alpha_{Ag,Manuf}^M$	0.61	0.34	0.61	0.88				
$\alpha_{Manuf,Ag}^M$	0.15	0.01	0.11	0.36				
$\alpha_{Manuf,Manuf}^M$	0.85	0.64	0.89	0.99				
<i>Bilateral Variables</i>								
Trade (total)	1.604	0.000	0.001	1.208	1.609	-0.002	0.001	1.223
Trade in Agriculture	0.085	0.000	0.000	0.018	0.103	0.000	0.000	0.030
Trade in Manufacturing	1.518	0.000	0.001	1.145	1.506	-0.002	0.001	1.152
τ_{Ag}	0.101	0.000	0.048	0.361	-0.036	-0.298	-0.007	0.138
τ_{Manuf}	0.064	0.000	0.038	0.204	-0.024	-0.157	-0.008	0.076
Distance	7.227	1.036	6.812	15.597				
Common colonizer	0.080	0.000	0.000	1.000				
Common language	0.112	0.000	0.000	1.000				
Contiguous	0.029	0.000	0.000	0.000				
Cross-border flow	0.992	1.000	1.000	1.000				

Notes: We have 4,495 observations for the country-level variables and 692,230 observations for the bilateral trade variables for 1991-2019. Mean temperature is recorded in degree Celsius, while days above 30°C (below 0°C) correspond to the number of days within a year where the maximum temperature reached the threshold. *L* reports active population in millions of people. Gross output *Y* and expenditures *X* are reported in trillions of current US dollars. Trade flows are reported in tens of billions of current US dollars.

G Gravity Estimates

Estimates of trade elasticities θ_s can be recovered from gravity regressions or by GMM using the procedure outlined in Appendix E. In this section, we present estimates of sector-specific gravity regressions. We use the point estimates on bilateral tariff data to generate simulated data for the Monte Carlo experiments, and to compute counterfactuals in the empirical application.

Table G.6: Gravity Estimates

	Agriculture		Manufacturing	
	OLS	IV	OLS	IV
$\ln(1 + \text{tariff}_{nist})$	-1.004 (0.590)	-1.473 (0.778)	-2.358 (0.452)	-4.591 (0.955)
$\ln(\text{distance}_{ni})$	-1.448 (0.104)	-1.501 (0.106)	-1.353 (0.070)	-1.331 (0.072)
Border_{ni}	-3.742 (0.397)	-3.635 (0.411)	-1.233 (0.241)	-1.104 (0.240)
Contiguous_{ni}	0.880 (0.162)	0.795 (0.167)	0.698 (0.173)	0.680 (0.170)
$\text{CommonLanguage}_{ni}$	0.514 (0.145)	0.564 (0.148)	0.776 (0.104)	0.758 (0.103)
ColonialLink_{ni}	0.259 (0.198)	0.204 (0.201)	0.845 (0.201)	0.849 (0.202)
# Obs	39,843	27,447	88,098	59,967
# Importers	122	122	122	122
# Exporters	130	129	130	130
# Importers \times Exporters	8,682	7,623	14,373	13,399

Note: Dependent variable is $\ln \check{\phi}_{nist}$. All regressions include importer-year and exporter-year fixed effects. Standard errors are clustered by importer and exporter and reported in parentheses. Sample includes the years 2001, 2004, 2007, 2010, 2013, and 2016 – the years available in the MAcMap dataset.

In the merged dataset, we have bilateral trade and tariff data for 132 countries over the period 1991 - 2019, along with gravity measures. We first solve for trade frictions ϕ_{nist} following equation (28), and then estimate equation (29). We pool all years together and estimate separately for each sector both by OLS, and by instrumental variables, instrumenting current tariffs with 3-year lag of tariffs.

Table G.6 shows the estimated trade elasticities. Across both sectors and specifications, we find that distance, tariffs, and crossing an international border all lower trade flows, while sharing a contiguous border, common language, and colonial link all increase trade flows, as expected. The point estimates on $\ln(1 + \text{tariff}_{nist})$ appear lower in magnitude than values that are usually reported in the literature, but when we instrument with 3-year lagged tariffs, we find $\theta_{ag} = 1.47$ and $\theta_{manuf} = 4.59$, which are within the range found in the literature. In particular, they are higher than estimates from [Boehm et al. \(2023\)](#): $\theta_{ag} = 0.8$, $\theta_{manuf} = 1.1$, but lower than estimates from [Shapiro \(2016\)](#): $\theta_{ag} = 3.3$, $\theta_{manuf} = 8.5$ and [Fontagné et al. \(2022\)](#): $\theta_{ag} = 6.9$, $\theta_{manuf} = 8.3$. In the empirical application, we test the sensitivity of our results to these alternative measures from the literature, but treat the estimates from Table G.6 as our preferred specification.

H Re-centering procedure from [Jones et al. \(2025\)](#)

We follow [Jones et al. \(2025\)](#) to build a statistical model of changes in *expected* exposure to each temperature bin. The procedure uses daily maximum temperature observations from ERA5 ([Hersbach et al., 2023](#)) from 1960 to 2020, aggregated from grid-cells to countries with population weights. It goes as follows:

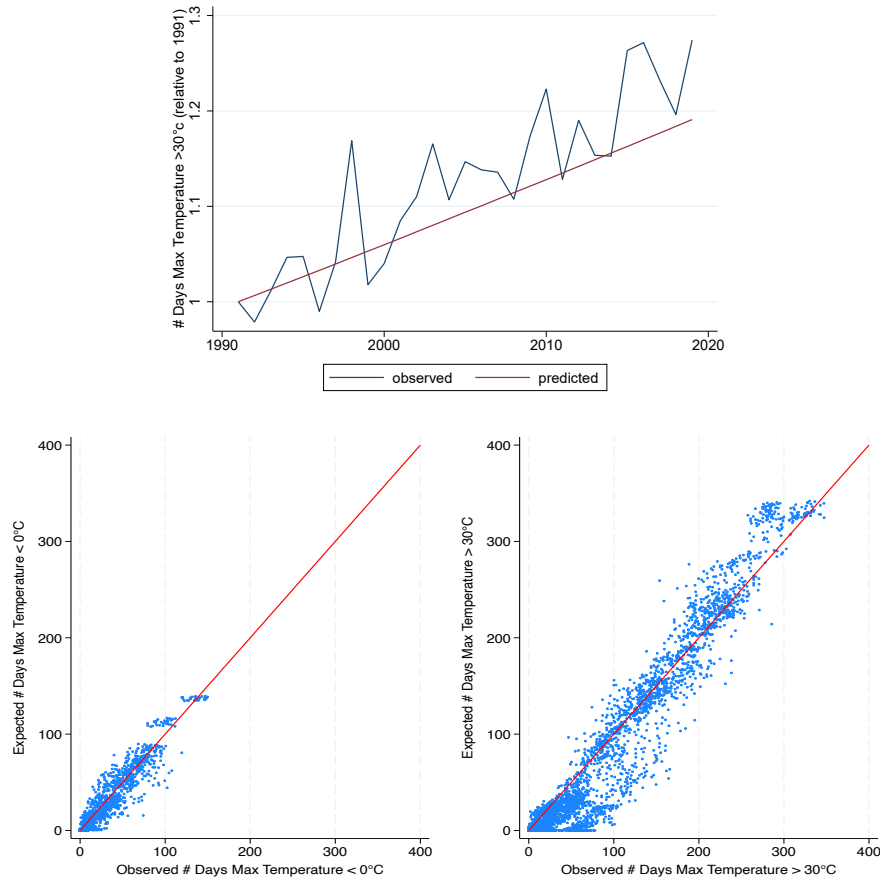
1. Define z_{idmt} as the observed temperature in country i , day d , month m and year t .
2. Compute the country-month-year temperature averages over days, \bar{z}_{imt} .
3. Estimate linear time trends in average country-month temperature across all years, by running the following regression for each country-month pair (i, m) : $\bar{z}_{imt} = \mu_{im} + \gamma_{im} \times t + \varepsilon_{imt}$, where $\hat{\gamma}_{im}$ is the estimated linear trend coefficient for month m in country i . We further transform this estimated coefficient using a Bayesian method to shrink it towards the average estimate: $\tilde{\gamma}_{im} = \frac{V(\hat{\gamma})}{V(\hat{\gamma}) + \sigma(\gamma_{im})^2} \hat{\gamma}_{im} + \left(1 - \frac{V(\hat{\gamma})}{V(\hat{\gamma}) + \sigma(\gamma_{im})^2}\right) \bar{\gamma}$ where $\sigma(\gamma_{im})^2$ is the variance of the estimated coefficient $\hat{\gamma}_{im}$, $\bar{\gamma}$ is the average estimated coefficient over all (i, m) pairs, $V(\hat{\gamma})$ is the variance of the estimated coefficient over all (i, m) pairs.
4. De-trend the original daily temperature observations z_{idmt} by subtracting off the country-month mean shift that has taken place up to time t : $z_{idm}^{\text{detrend}} = z_{idmt} - \tilde{\gamma}_{im} \times t$.
5. Combine all de-trended temperature realizations in each country-month into a single distribution, denoted $F_{im}(z)$.
6. Project forward the country-month distributions of temperature realizations, using the country-month specific time trend $\tilde{\gamma}_{im}$: $\check{z}_{idmt} = z_{idm}^{\text{detrend}} + \tilde{\gamma}_{im} \times t$. This generates

an empirical distribution of potential temperature realizations for each country-year, denoted $F_{imt}(z)$.

- Combine $F_{imt}(z)$ across months within each country-year: this generates the counterfactual temperature distribution at the county-year level, $F_{it}(z)$. Then compute the share of observations in each country-year that fall the different temperature bins, and normalize by the number of days in the year: this gives the estimate of non-random exposure.

Figure H.1 shows the comparison between observed and predicted non-random bin variables, built using this method.

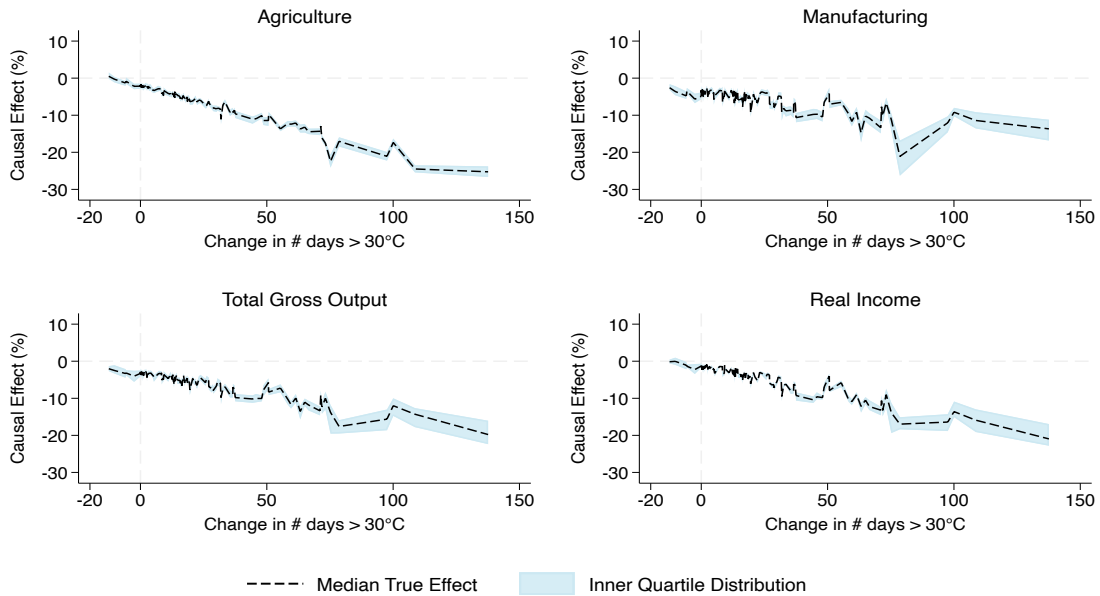
Figure H.1: Observed vs Predicted Weather Variables



Notes: The upper figure plots the observed and predicted count of hot days, averaged across countries, whereas the bottom figure plots the observed and predicted count of cold and hot days, by country-year. The *observed* counts of days are computed from realized ERA5 temperature observations; the *predicted* counts of days are computed following Jones et al. (2025) procedure.

I Additional Monte Carlo Results

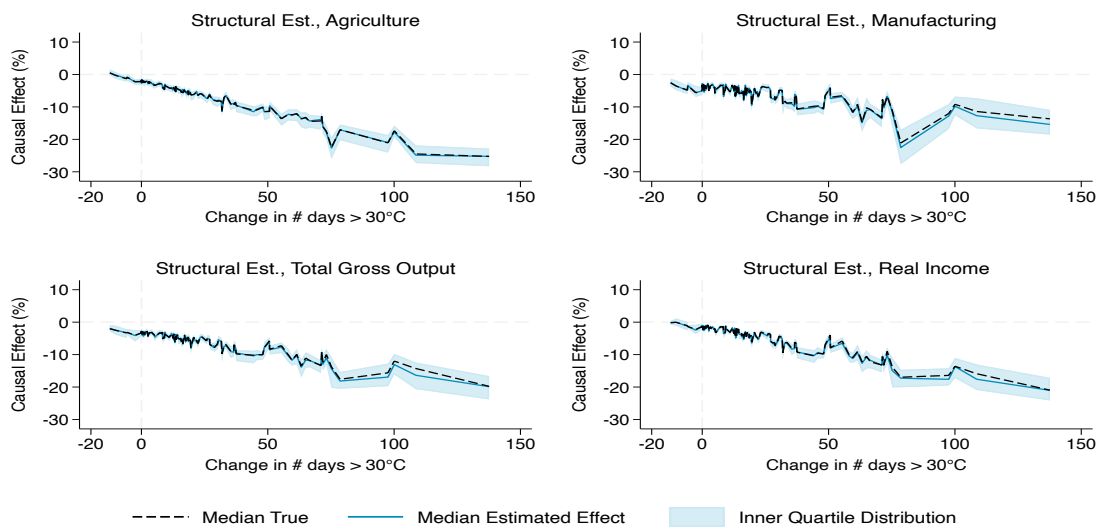
Figure I.1: True Treatment Effects Across Replications



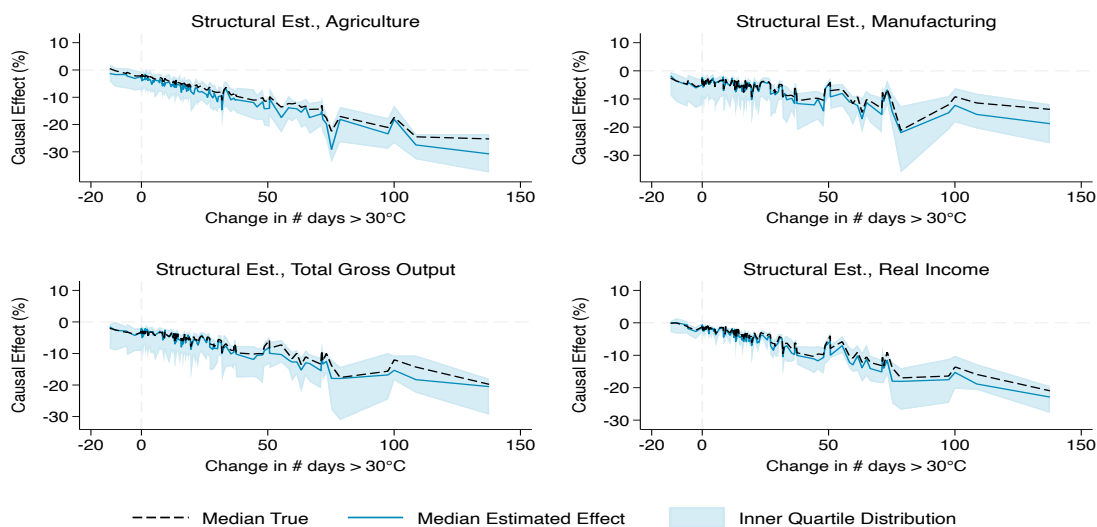
Notes: Figure plots the Individual Overall Causal Effect in % on the y-axis against the change in the # of days with maximum temperature above 30°C between 1991 and 2019, i.e., the change in the observed factor we wish to quantify. The black dashed line plots the median IOCE for agriculture (left panel) manufacturing (middle panel) and gross output (right panel) across 200 replications. The blue shaded region depicts the interquartile range of the distribution of these effects across replications. Simulations include 132 countries and 29 time periods and two sectors.

Figure I.2: Estimated vs True Treatment Effects, Structural and GMM Estimators

(a) Structural Estimator



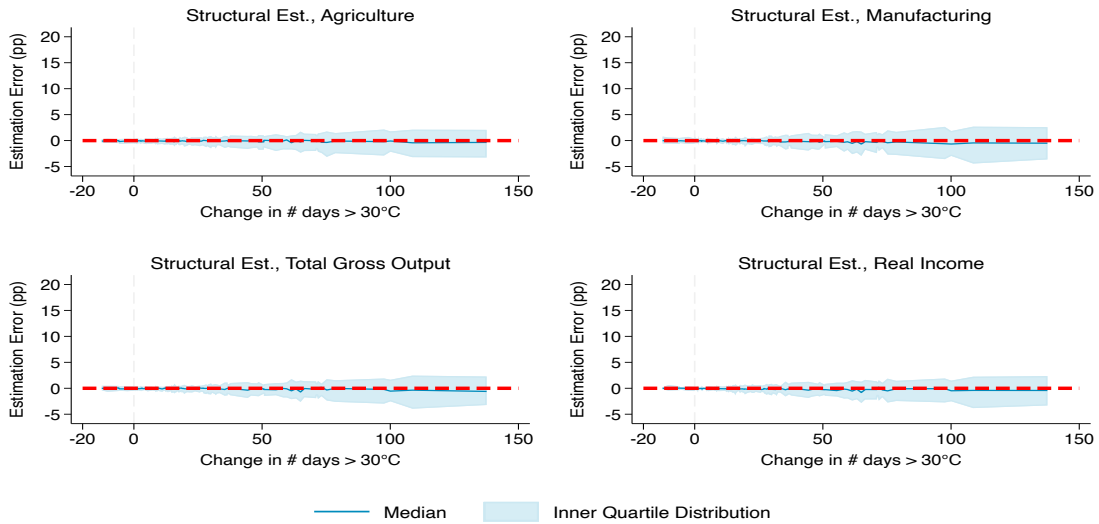
(b) GMM Estimator



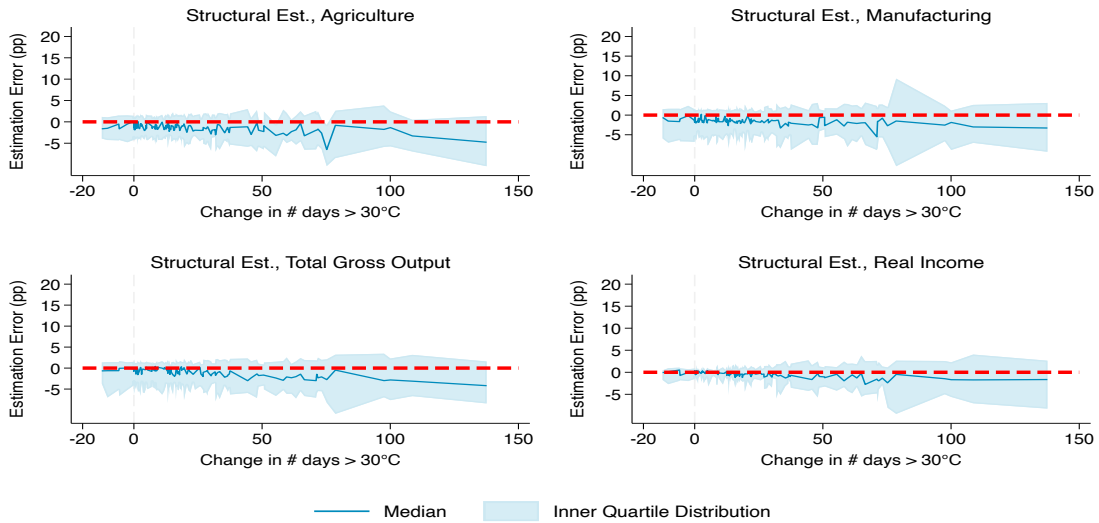
Notes: Figure plots the Individual Overall Causal Effect in % on the y-axis against the change in the # of days with maximum temperature above 30°C between 1991 and 2019. The black dashed line plots the median IOCE for agriculture (left panel) manufacturing (middle panel) and gross output (right panel) across 200 replications. The blue solid line presents the median estimate from the structural estimator (panel a) and the GMM estimator (panel b), and the blue shaded region depicts the interquartile range of the distribution of these estimates. Simulations include 132 countries and 29 time periods and two sectors.

Figure I.3: Estimation Errors, Structural and GMM Estimators

(a) Structural Estimator



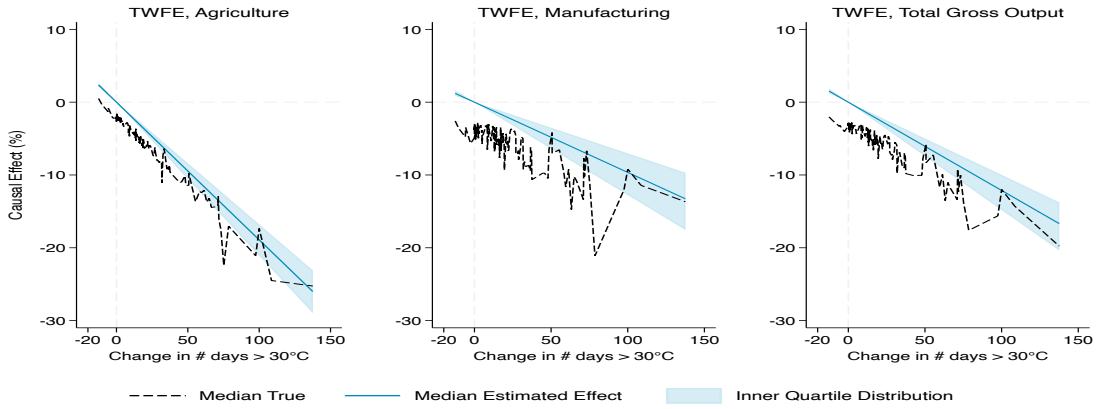
(b) GMM Estimator



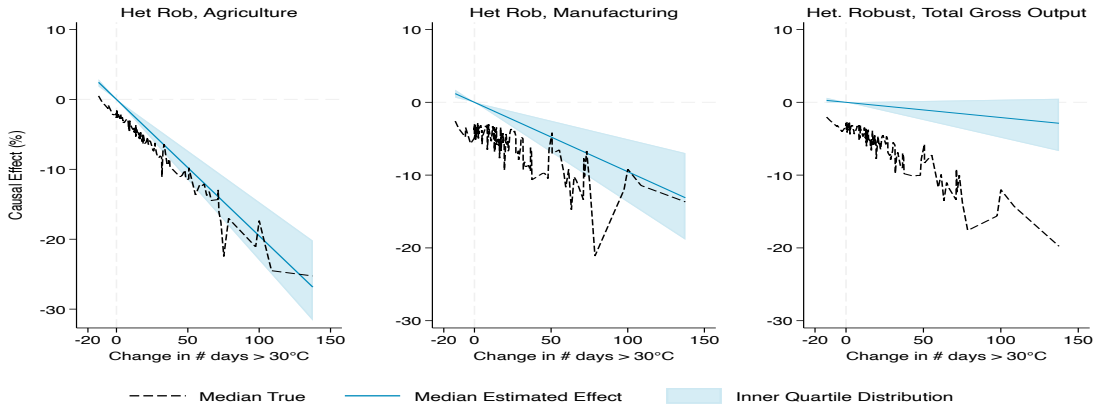
Notes: Figure plots the difference between estimated IOCE and true IOCE, i.e., the estimation error by replication, against the change in the # of days with maximum temperature above 30°C between 1991 and 2019. The blue solid line presents the median estimation error for agriculture (left panel) manufacturing (middle panel) and gross output (right panel) across 200 replications based on the respective estimator. The blue shaded region depicts the interquartile range of the distribution of these estimation errors. Simulations include 132 countries and 29 time periods and two sectors.

Figure I.4: Estimated vs True Treatment Effects: Comparison of Estimators

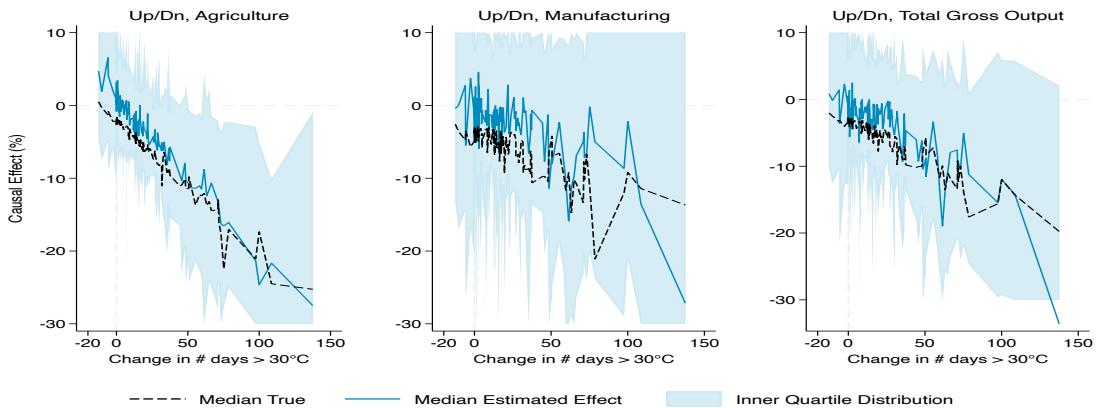
(a) TWFE Estimator



(b) Heterogeneous-Robust Estimator



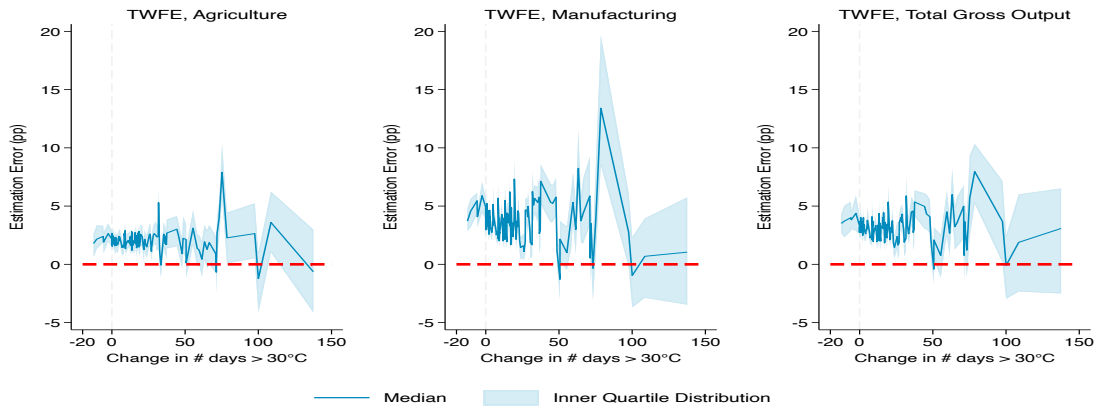
(c) Upstream/Downstream Estimator



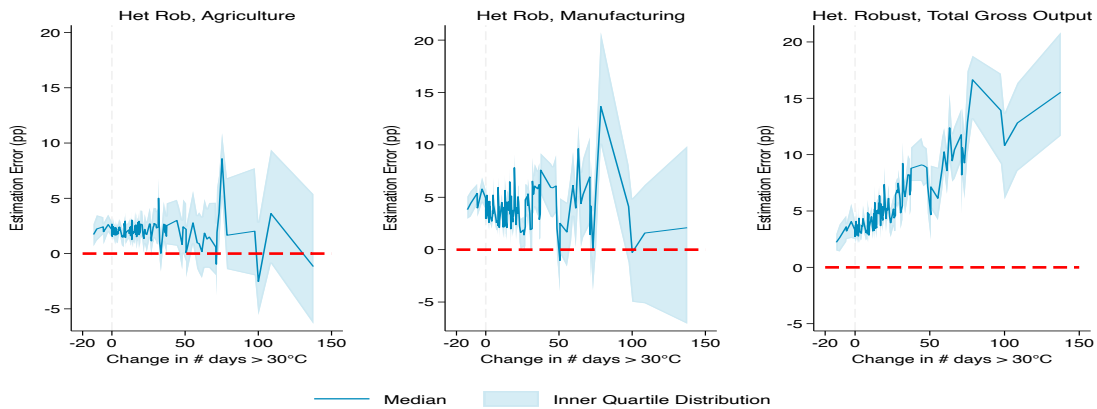
Notes: Each panel plots the Individual Overall Causal Effect in % on the y-axis against the change in the # of days with maximum temperature above 30°C between 1991 and 2019. The black dashed line plots the median IOCE for agriculture (left panel), manufacturing (middle panel), and gross output (right panel) across 200 replications. The blue solid line presents the median estimate from the respective estimator, and the blue shaded region depicts the interquartile range of the distribution of these estimates. Simulations include 132 countries and 29 time periods and two sectors.

Figure I.5: Estimation Errors: Comparison of Estimators

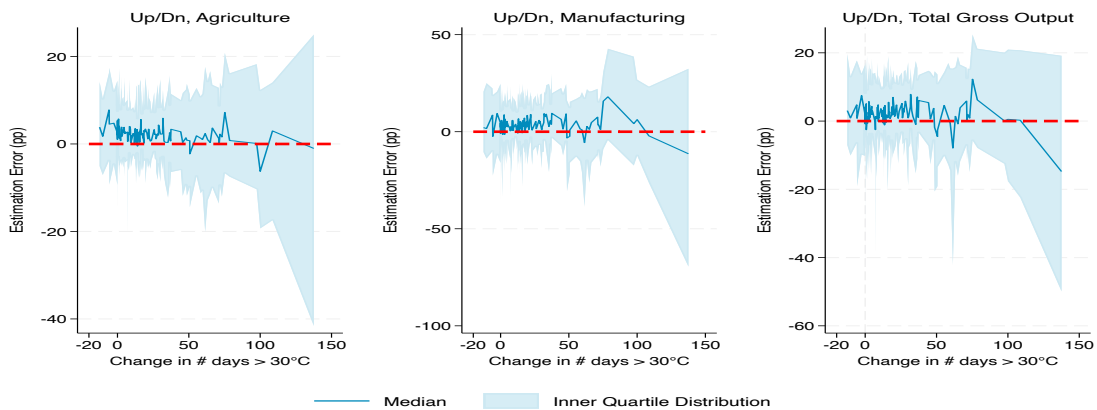
(a) TWFE Estimator



(b) Heterogeneous-Robust Estimator




(c) Upstream/Downstream Estimator



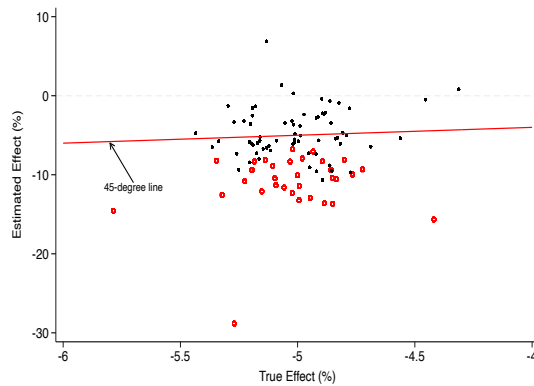
Notes: Figure plots the difference between estimated IOCE and true IOCE, i.e., the estimation error by replication, against the change in the # of days with maximum temperature above 30°C between 1991 and 2019. The blue solid line presents the median estimation error for agriculture (left panel) manufacturing (middle panel) and gross output (right panel) across 200 replications from the respective estimator. The blue shaded region depicts the interquartile range of the distribution of these estimation errors. Simulations include 132 countries and 29 time periods and two sectors.

Figure I.6: Estimated vs True Treatment Effects, Global Estimator


(a) Baseline Variance




(b) Low Variance



(c) Extended-Period Panel



(d) Extended-Period Panel, distributions

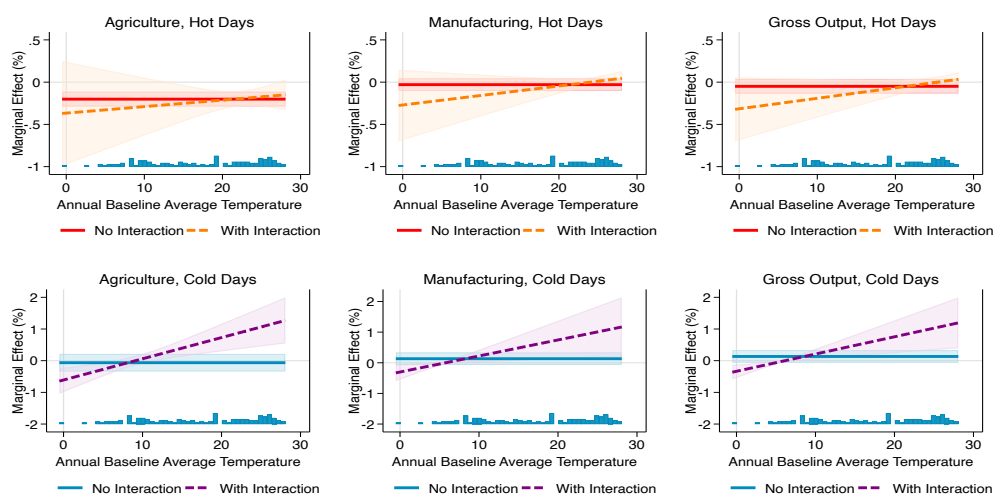


Notes: Figure plots the estimated treatment effect on global gross output in % on the y-axis against the true treatment effect in % on the x-axis, where the estimates are computed via the local projection approach from [Bilal & Känzig \(2024\)](#). The 45-degree line is plotted in red. Red dots indicate replications for which the point estimate on contemporaneous temperature shocks are statistically significant at the 5% level. In panels (a) and (b), simulations include 132 countries and 29 time periods and the variance of productivity draws is 0.3 (0.02) for panel (a) (panel b). In panel (c) and (d), simulations include 58 periods, as each initial year was split into 2, with the same weather realizations but different productivity draws (high variance). Panel (d) plots the distribution of estimated and true effects from panel (c).

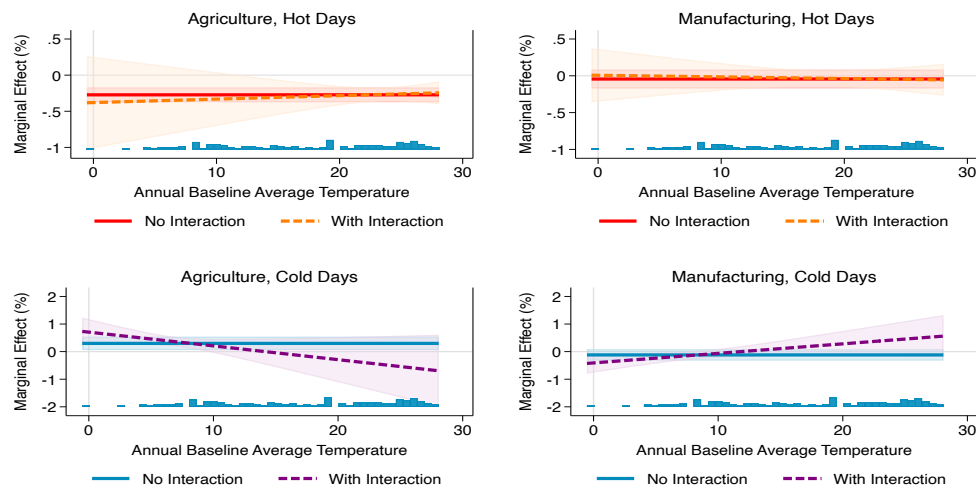
J Further Empirical Results

Figure J.1: Marginal Effects, # Hot Days vs # Cold Days

(a) TWFE



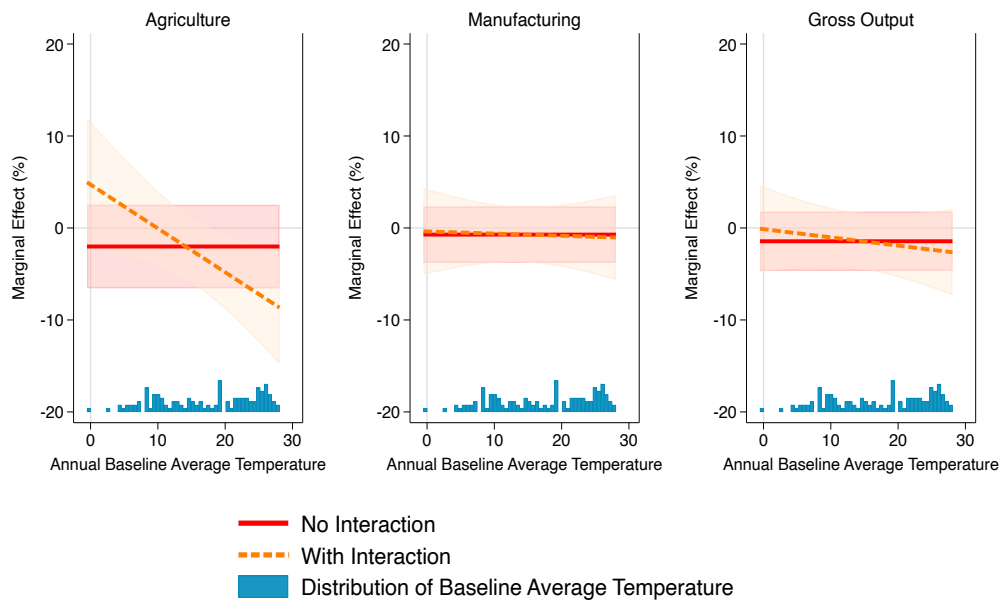
(b) Structural



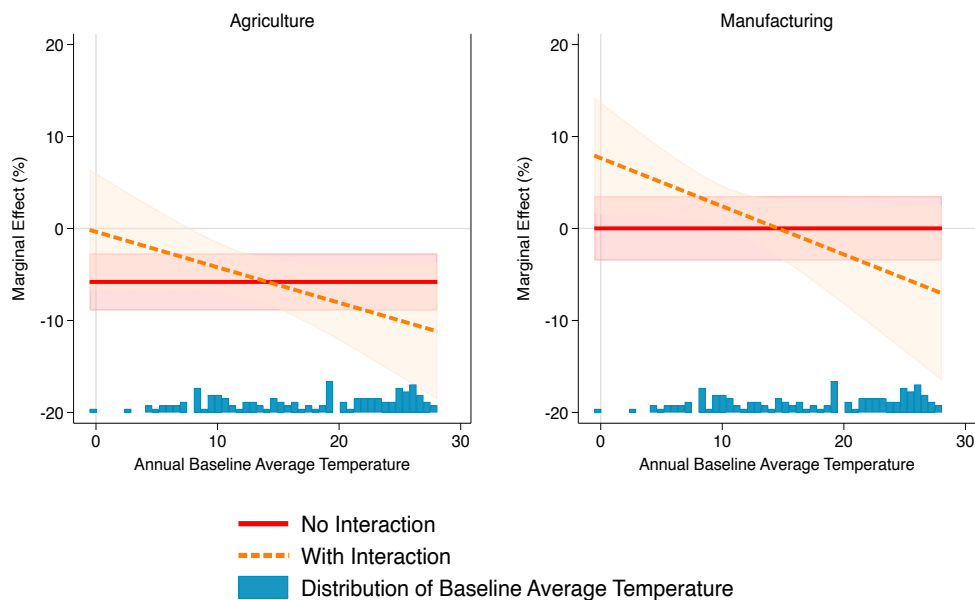
Notes: Figure plots marginal effects (in percent) of an additional hot day ($> 30^{\circ}\text{C}$) or cold day ($< 0^{\circ}\text{C}$) on sectoral nominal output, estimated using TWFE (panel a) or structural (panel b) estimators. Marginal effects are shown as a function of baseline annual average temperature (country-level average over first three years of sample in $^{\circ}\text{C}$). “With interaction” specifications allow cold/hot days to interact with baseline average temperature; “No interaction” specifications impose constant effects. Shaded areas denote 95% confidence intervals. Histogram plots the distribution of baseline average temperature in blue.

Figure J.2: Marginal Effects, Average Annual Temperature

(a) TWFE



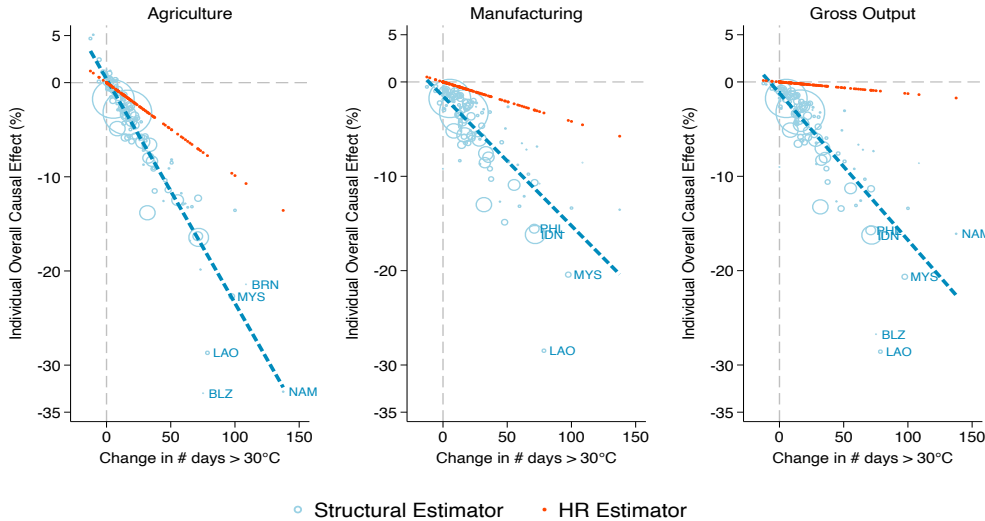
(b) Structural



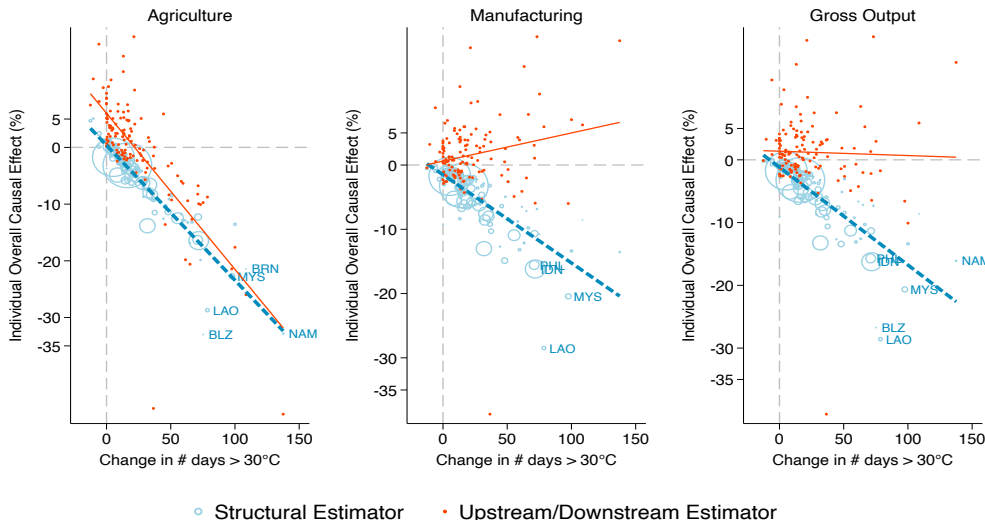
Notes: Figure plots marginal effects (in percent) of an increase in annual average temperature ($^{\circ}\text{C}$) on sectoral nominal output, estimated using TWFE (panel a) and structural (panel b) estimators. Panels report results for agriculture, manufacturing, and aggregate gross output. Marginal effects are shown as a function of baseline annual average temperature ($^{\circ}\text{C}$). “With interaction” specifications allow temperature shocks to interact with baseline average temperature; “No interaction” specifications impose constant effects. Shaded areas denote 95% confidence intervals. Histogram plots the distribution of baseline average temperature in blue.

Figure J.3: Comparing Structural Estimator with Reduced-Form Estimators

(a) Heterogeneous-Robust Estimator



(b) Upstream/Downstream Estimator



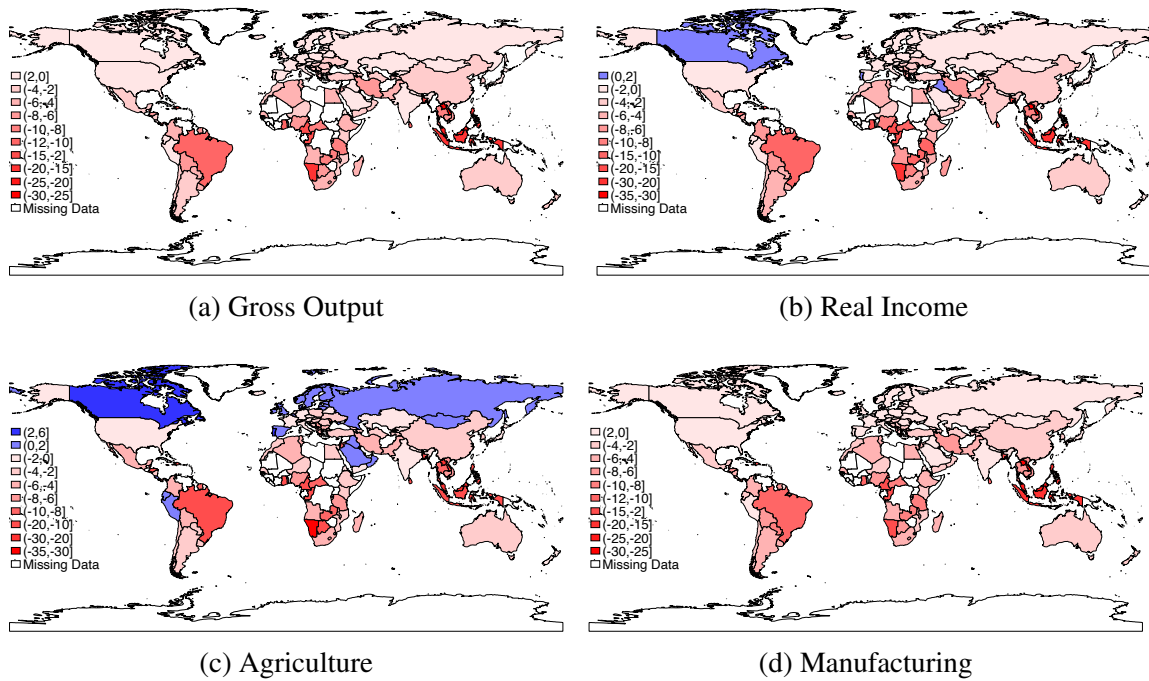
Notes: Figure plots the country-specific % change in agriculture gross output, manufacturing gross output, and aggregate gross output resulting from the change from 1991 weather to 2019 weather against observed change in # extreme heat days. Structural estimates are depicted with blue circles and scaled by the size of the workforce in 1991. Alternative estimators are depicted with red dots, in panel (a) for the heterogeneous-robust estimator and panel (b) for the upstream/downstream estimators.

Table J.1: Upstream/Downstream Point Estimates, all specifications

	Agriculture				Manufacturing			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\beta^{>30}$	-0.0063 (0.0023)	-0.0128 (0.0035)	-0.0062 (0.0023)	-0.0128 (0.0037)	0.0013 (0.0008)	0.0013 (0.0017)	0.0015 (0.0008)	0.0016 (0.0017)
$\beta_{UP,AG}^{>30}$	0.0019 (0.0026)	0.0098 (0.0102)	0.0017 (0.0026)	0.0104 (0.0108)	0.0003 (0.0019)	0.0089 (0.0145)	0.0002 (0.0019)	0.0104 (0.0154)
$\beta_{UP,Man}^{>30}$	0.0131 (0.0058)	0.0074 (0.0158)	0.0131 (0.0060)	0.0080 (0.0165)	0.0079 (0.0055)	0.0017 (0.0218)	0.0088 (0.0056)	0.0036 (0.0221)
$\beta_{DN,AG}^{>30}$	0.0011 (0.0007)	-0.0029 (0.0046)	0.0011 (0.0007)	-0.0023 (0.0048)	-0.0005 (0.0011)	-0.0065 (0.0049)	-0.0005 (0.0012)	-0.0063 (0.0049)
$\beta_{DN,Man}^{>30}$	-0.0105 (0.0060)	0.0014 (0.0199)	-0.0105 (0.0062)	-0.0007 (0.0202)	-0.0102 (0.0055)	-0.0132 (0.0216)	-0.0110 (0.0057)	-0.0167 (0.0227)
$\beta^{>30} \times \bar{z}_i$		0.0006 (0.0002)		0.0006 (0.0002)		0.0000 (0.0001)		0.0000 (0.0001)
$\beta_{UP,AG}^{>30} \times \bar{z}_i$		-0.0007 (0.0005)		-0.0007 (0.0006)		-0.0007 (0.0007)		-0.0007 (0.0007)
$\beta_{UP,Man}^{>30} \times \bar{z}_i$		0.0004 (0.0008)		0.0004 (0.0008)		0.0003 (0.0011)		0.0003 (0.0011)
$\beta_{DN,AG}^{>30} \times \bar{z}_i$		0.0003 (0.0004)		0.0003 (0.0004)		0.0005 (0.0004)		0.0005 (0.0004)
$\beta_{DN,Man}^{>30} \times \bar{z}_i$		-0.0007 (0.0010)		-0.0007 (0.0010)		0.0001 (0.0011)		0.0003 (0.0011)
$\beta^{<0}$			-0.0018 (0.0014)	-0.0084 (0.0019)			0.0001 (0.0006)	-0.0013 (0.0013)
$\beta_{UP,AG}^{<0}$			-0.0038 (0.0044)	0.0206 (0.0210)			0.0046 (0.0055)	0.0194 (0.0212)
$\beta_{UP,Man}^{<0}$			0.0058 (0.0054)	-0.0343 (0.0317)			-0.0029 (0.0044)	0.0099 (0.0352)
$\beta_{DN,AG}^{<0}$			-0.0001 (0.0004)	-0.0039 (0.0031)			0.0002 (0.0003)	0.0005 (0.0034)
$\beta_{DN,Man}^{<0}$			0.0027 (0.0064)	0.0210 (0.0325)			0.0004 (0.0055)	-0.0368 (0.0311)
$\beta^{<0 \times \bar{z}_i}$				0.0010 (0.0003)				0.0003 (0.0002)
$\beta_{UP,AG}^{<0} \times \bar{z}_i$				-0.0030 (0.0019)				-0.0020 (0.0023)
$\beta_{UP,Man}^{<0} \times \bar{z}_i$				0.0041 (0.0030)				-0.0011 (0.0037)
$\beta_{DN,AG}^{<0} \times \bar{z}_i$				0.0004 (0.0004)				0.0000 (0.0004)
$\beta_{DN,Man}^{<0} \times \bar{z}_i$				-0.0018 (0.0032)				0.0039 (0.0033)

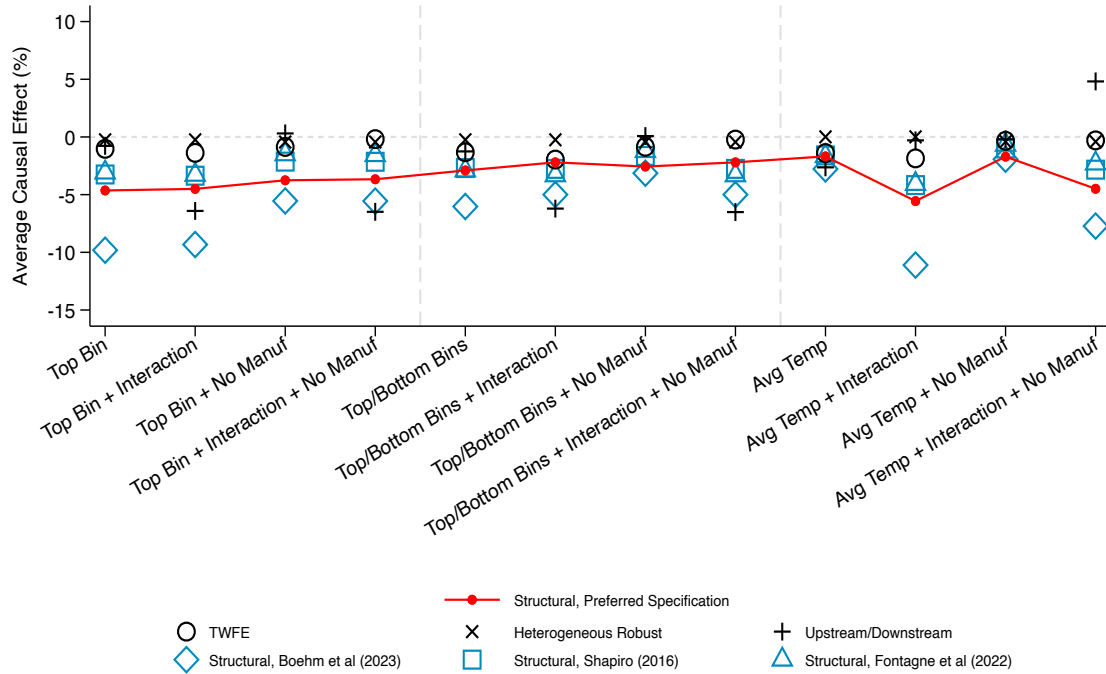
Note: Dependent variable is year-on-year growth in agricultural (columns 1-4) and manufacturing gross outputs (columns 5-8). In all regressions, we control for flexible continent-by-year and country fixed effects, as well as country-level variables $(1 - \eta_{is}) \alpha_{ish}^M$ interacted with year fixed effects. We also control for the expected year-on-year change in the number of days in the year with a maximum temperature above 30°C, computed via the procedure from Jones et al. (2025). All regressions are weighted by baseline population. Baseline climate \bar{z}_i is the average annual temperature in °C in the first three years of the sample. Standard errors are clustered by country. Sample includes 132 countries over the period 1991-2019. Point estimates statistically distinguishable from zero are presented in bold font.

Figure J.4: Structural Estimates of Effects by Sector and in Aggregate



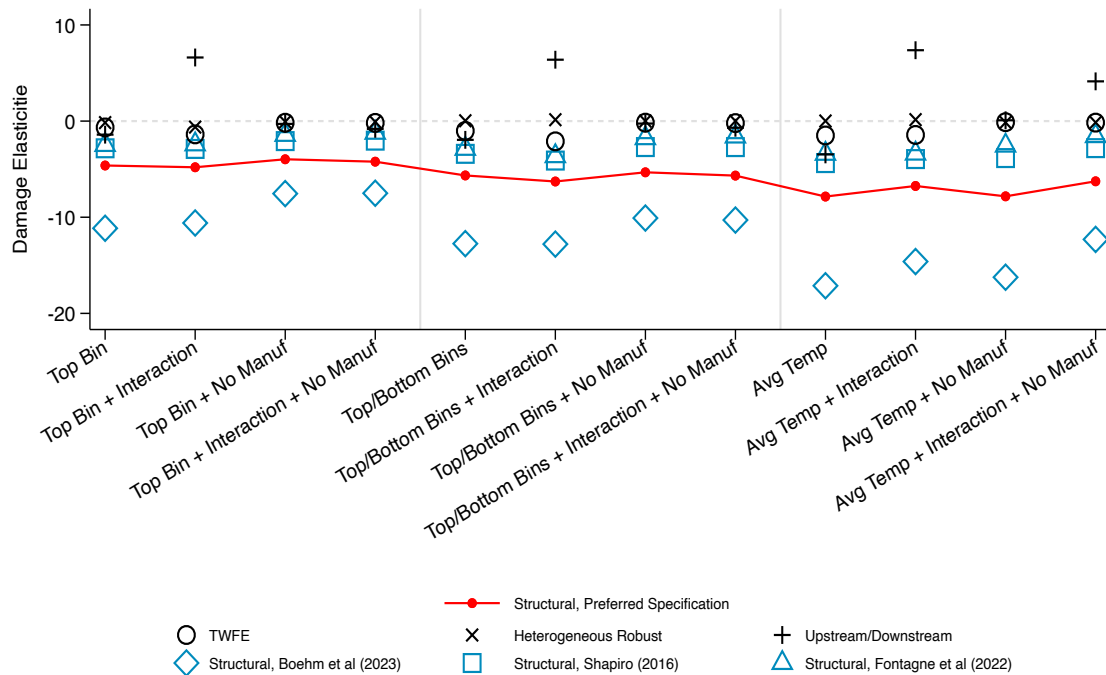
Notes: Figure plots the % change in 2019 aggregate gross output (panel a), real wage (panel b), agriculture gross output (panel c), and manufacturing gross output (panel d) resulting from the change in # days with maximum temperature above 30°C between 1991 and 2019 estimated using our structural procedure. We set $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$

Figure J.5: Average Causal Effects from warming between 1991 and 2019 by Specification



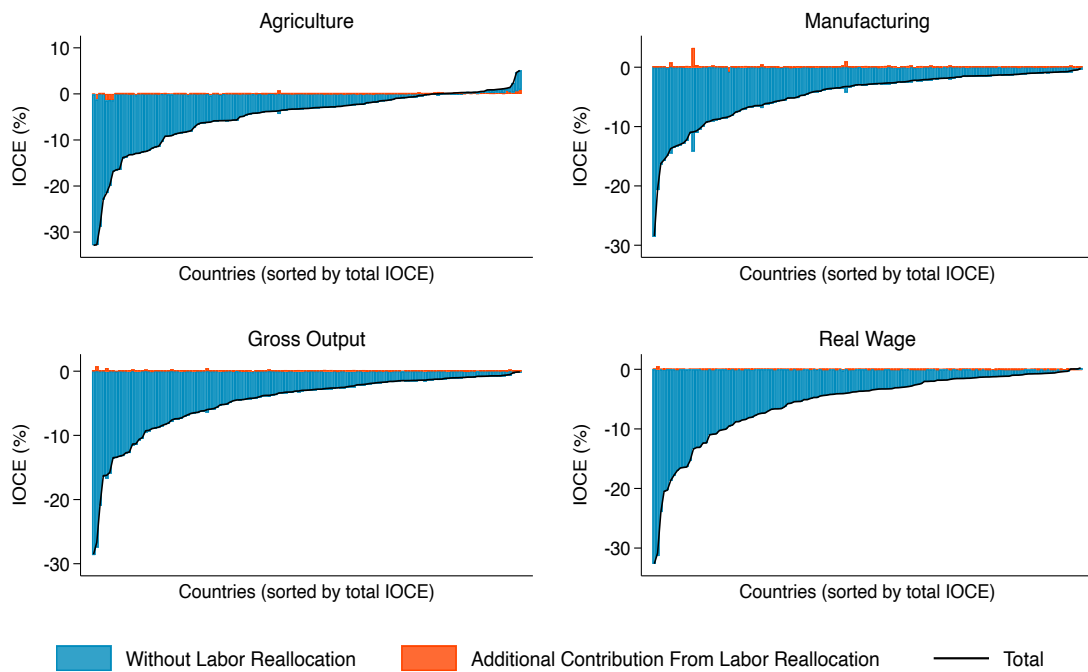
Notes: Figure plots across-country average causal effect from warming between 1991 - 2019 by estimator and environmental specification. Red dots indicate our structural approach using our preferred estimates $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$. Blue markers indicate results from our structural approach using estimates of θ_s from [Boehm et al. \(2023\)](#) ($\theta^{Ag} = 0.8$ and $\theta^{Manuf} = 1.1$), [Shapiro \(2016\)](#) ($\theta^{Ag} = 3.3$ and $\theta^{Manuf} = 8.5$) and [Fontagné et al. \(2022\)](#) ($\theta^{Ag} = 6.9$ and $\theta^{Manuf} = 8.3$). First 4 columns present result from specifications that include the number of days with maximum temperature above 30°C. Next 4 columns present result from specifications that include both the number of days with maximum temperature above 30°C and the number of days with maximum temperature below 0°C. Last 4 columns present result from specifications that include annual average temperature in °C. Even-numbered columns include interactions with baseline annual average temperature in °C (“+ interaction”). Columns 3-4, 7-8, 11-12 impose zero effects to productivity in the manufacturing sector.

Figure J.6: Estimated Elasticity of Global Gross Output Loss to Global Average Annual Temperature Growth ($^{\circ}\text{C}$), by Specification



Notes: Figure plots estimated elasticity of global gross output loss to global average annual temperature growth ($^{\circ}\text{C}$) by estimator and environmental specification. Elasticity is computed as the point estimate of a regression with global gross output loss as the dependent variable and the global average annual temperature growth as the independent variable. Red dots indicate our structural approach using our preferred estimates $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$. Blue markers indicate results from our structural approach using estimates of θ_s from Boehm et al. (2023) ($\theta^{Ag} = 0.8$ and $\theta^{Manuf} = 1.1$), Shapiro (2016) ($\theta^{Ag} = 3.3$ and $\theta^{Manuf} = 8.5$) and Fontagné et al. (2022) ($\theta^{Ag} = 6.9$ and $\theta^{Manuf} = 8.3$). First 4 columns present result from specifications that include the number of days with maximum temperature above 30°C . Next 4 columns present result from specifications that include both the number of days with maximum temperature above 30°C and the number of days with maximum temperature below 0°C . Last 4 columns present result from specifications that include annual average temperature in $^{\circ}\text{C}$. Even-numbered columns include interactions with baseline annual average temperature in $^{\circ}\text{C}$ (“+ interaction”). Columns 3-4, 7-8, 11-12 impose zero effects to productivity in the manufacturing sector.

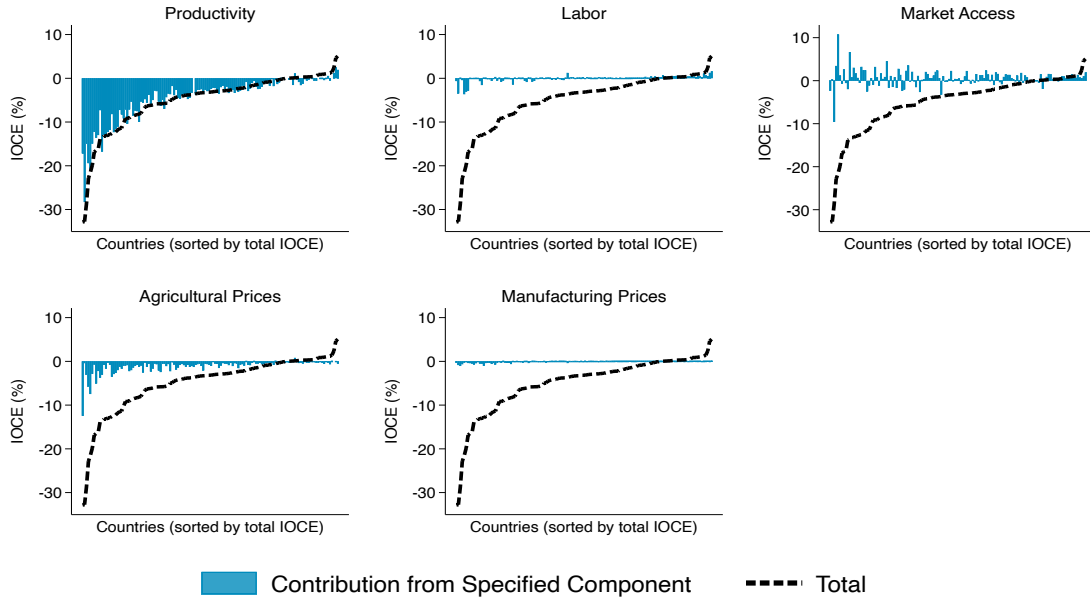
Figure J.7: Contribution to IOCE from Labor Reallocation



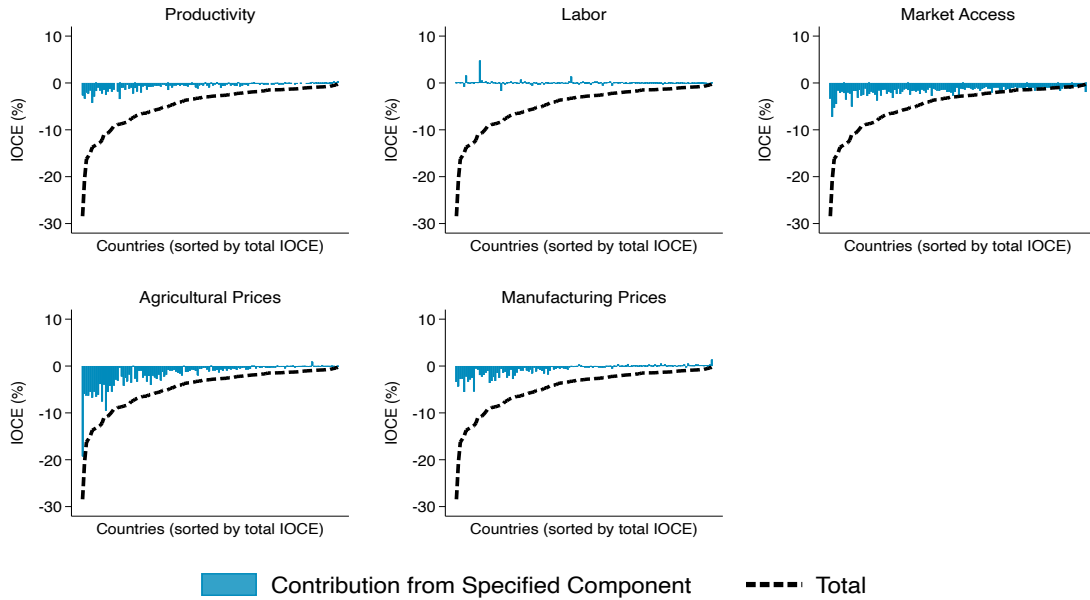
Notes: Figure plots the % change in 2019 agriculture gross output, manufacturing gross output, total gross output, and real wage, resulting from the change in # days with maximum temperature above 30°C between 1991 and 2019 estimated using our structural procedure (black line). We set $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$. In blue, we plot the IOCE imposing no labor reallocation. The difference between this counterfactual and the total IOCE yields the contribution to the IOCE from labor reallocation, plotted in orange. Countries are sorted on the x-axis by outcome-specific IOCE.

Figure J.8: Decomposition of the Different Channels, by Sector

(a) Agricultural Gross Output



(b) Manufacturing Gross Output



Notes: Figure plots the % change in 2019 agriculture (panel a) or manufacturing (panel b) gross output resulting from the change in # days with maximum temperature above 30°C between 1991 and 2019 estimated using our structural procedure (black line), with $\theta^{Ag} = 1.5$ and $\theta^{Manuf} = 4.6$. In blue, we plot contribution to the IOCE from each specified component. Decomposition results from log-linearizing equilibrium changes in gross revenues (33), and then computing each individual component. Countries are sorted on the x-axis by IOCE.

References

- Acemoglu, D., Akcigit, U., & Kerr, W. R. (2016). Networks and the macroeconomy: An empirical exploration. In M. Eichenbaum & J. A. Parker (Eds.), *NBER Macroeconomics Annual*, volume 30 of *NBER Chapters* (pp. 273–335).
- Aguiar, A., Narayanan, B., & McDougall, R. (2016). An overview of the gtap 9 data base. *Journal of Global Economic Analysis*, 1(1), 181–208.
- Bilal, A. & Känzig, D. R. (2024). The macroeconomic impact of climate change: Global vs. local temperature. *National Bureau of Economic Research Working Paper*, (32450).
- Boehm, C. E., Levchenko, A. A., & Pandalai-Nayar, N. (2023). The long and the short of trade elasticities. *Journal of Political Economy*, 131(7), 1771–1824.
- Caliendo, L. & Parro, F. (2015). Estimates of the trade and welfare effects of NAFTA. *The Review of Economic Studies*, 82(1), 1–44.
- Conte, M., Cotterlaz, P., Mayer, T., et al. (2022). The CEPII gravity database. *CEPII*.
- Das, S., Magistretti, G., Pugacheva, E., & Wingender, P. (2022). Sectoral spillovers across space and time. *Journal of Macroeconomics*, 72(C), 103422.
- de Chaisemartin, C., D'Haultfoeuille, X., & Vazquez-Bare, G. (2024). Difference-in-difference estimators with continuous treatments and no stayers. *AEA Papers and Proceedings*, 114, 610–613.
- Eaton, J. & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5), 1741–1779.
- Feng, A., Li, H., & Wang, Y. (2025). We are all in the same boat: Cross-border spillovers of climate shocks through international trade and supply chains. *The Economic Journal*, 135(669), 1433–1466.
- Fontagné, L., Guimbard, H., & Orefice, G. (2022). Tariff-based product-level trade elasticities. *Journal of International Economics*, 137, 103593.
- Fontagné, L., Lebrand, M. S. M., Murray, S., Santoni, G., & Ruta, M. (2023). Trade and infrastructure integration in Africa. *Mimeo*, Available at SSRN 4672520.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Sabater, J. M., Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al. (2023). ERA5 hourly data on single levels from 1940 to present.
- Jones, B., Moscona, J., von Dessauer, C., & Olken, B. (2025). With or without U? binning bias and the causal effects of temperature shocks. *Mimeo*.
- Mayer, T., Santoni, G., & Vicard, V. (2023). *The CEPII trade and production database*. CEPII.
- Shapiro, J. S. (2016). Trade costs, CO₂, and the environment. *American Economic Journal: Economic Policy*, 8(4), 220–254.
- World Bank (2004). *Iraq Trade Brief*. Technical report, The World Bank. Imposed a 5% tariff on most imported goods as part of the reconstruction tariff regime (with exceptions for food, medicine, medical equipment, clothing, books, and humanitarian aid).
- Zappalà, G. (2024). Estimating sectoral climate impacts in a global production network. *Mimeo*.