Both authors (Tianzi Liu and Zebang Xu) acknowledge financial support from the Structural Transformation and Economic Growth (STEG) PhD Research Grant. The authors declare no conflict of interest.

The (Mis)Allocation Channel of Climate Change

Evidence from Global Firm-level Microdata *

Tianzi Liu †
Cornell University

Zebang Xu ‡
Cornell University

Job Market Paper, 2026-2027

December 2025, Updated Regularly

Abstract

Macroeconomic estimates of climate damages are large, yet firm-level micro estimates often appear modest. We reconcile this gap by uncovering a powerful channel—*the impact of extreme temperatures on capital misallocation*—as the major driver of aggregate climate damage. Using global firm-level microdata from 32 countries, we provide causal evidence that a day of extreme heat (> 30°C) increases the dispersion of returns to capital across firms by 0.29 log points, implying a 0.10% annual aggregate TFP loss for the average region–sector. The effect is more pronounced in hotter and more economically developed regions. To explain these mechanisms, we develop a tractable general-equilibrium model of firm dynamics with heterogeneous temperature sensitivities, endogenous adaptation, and time-to-build capital. The model shows that higher climate volatility increases regional dispersion in returns to capital, with amplification driven by heterogeneity in firms' temperature exposure, adaptation, and capital wedges. When estimated by the microdata, climate-induced misallocation lowers global TFP by 12.5% and accounts for 9% of cross-country productivity differences. Our results shift the focus of climate damage from direct physical productivity losses to *allocative efficiency losses in the cross-section*, highlighting firm heterogeneity as central for measuring aggregate damages and designing effective mitigation and adaptation policy.

*We thank Ryan Chahrour, Kristoffer Nimark, Ivan Rudik, Ezra Oberfield and Mathieu Taschereau-Dumouchel for their continued guidance and support. This paper has also benefited from discussions with Doug Almond, Levon Barseghyan, Johannes Boehm, Tamma Carleton, Julieta Caunedo, Luming Chen, Tomás Domínguez-Iino, Caroline Flammer, Todd Gerarden, Sampreet Singh Goraya, Cathy Kling, Kyle Meng, Pablo Ottonello, Will Rafey , Diego Restuccia, Jeffrey Shrader, and seminar participants at Columbia Climate School, Cornell Macro Lunch and Cornell SEERE Seminar, the 2024 Econometric Society NASM, CU Environmental & Resource Economics Workshop, SED 2025 Winter Meeting, the 2025 NBER Climate Finance PhD Workshop and ASSA 2026. We thank the Structural Transformation and Economic Growth (STEG) Research programme for supporting this research. First draft was circulated on February 16th, 2024.

†Department of Applied Economics and Management, Cornell University (email:tl567@cornell.edu).

‡Department of Economics, Cornell University (email:zx88@cornell.edu).

1 Introduction

Macroeconomic estimates of climate damages are large. By 2100, projected losses range from roughly 23 percent of global GDP in Burke, Hsiang, and Miguel (2015) to about 50 percent in Bilal and Känzig (2024). These costs are typically summarized as declines in aggregate productivity, a framing that dates to Nordhaus (1992). Yet direct estimates of firm- and plant-level productivity damage from temperature are much smaller. Addoum, Ng, and Ortiz-Bobea (2020) and Ponticelli, Xu, and Zeume (2023) estimate losses below 2 percent in U.S. plants, while the more recent global manufacturing evidence in Nath (2025) suggests average losses on the order of 1–2 percent once adaptation is taken into account. This gap raises a basic question: if average within-firm productivity damage is modest, what microeconomic channel generates large aggregate productivity losses?

This paper identifies and quantifies one important answer: climate-induced *across-firm* capital misallocation. Temperature shocks need not mainly operate by shifting down the average firm's production frontier. They also lower aggregate TFP by widening dispersion in firms' marginal revenue products of capital (MRPK), even within narrowly defined region–sector cells. The intuition is simple. Suppose one firm benefits from heat while another is harmed by it, either because their products differ in temperature sensitivity (for example, an ice cream parlor versus a ski resort) or because they differ in their endogenous adaptation choices. When an unexpected heat shock arrives, both firms enter the period with predetermined capital that takes time to build and cannot adjust on impact. The heat-favored firm becomes capital-scarce and its MRPK rises, while the heat-exposed firm becomes capital-abundant and its MRPK falls. Aggregate output is then lower than it would be under the counterfactual allocation in which capital had been installed at the firms with the higher realized returns. Climate damage, in this view, is partly allocative.

We begin by developing a climate–TFP accounting framework that decomposes region–sector-level aggregate TFP into an efficient frontier component (technology) and a loss component due to capital misallocation (allocative inefficiency). We then show formally that climate-induced changes in within-region-sector MRPK dispersion provide a sufficient statistic for the misallocation channel. Guided by this framework, we estimate the causal effect of temperature shocks on capital misallocation using firm-level data from 30 European countries, China, and India, matched to ERA5-Land weather data. Identification comes from year-to-year changes in local temperature distributions within region-, relative to other regions within the same country-year.

Our reduced-form estimates show a clear U-shaped relationship between temperature and capital misallocation: both extreme heat and extreme cold raise MRPK dispersion. Quantitatively, an additional day above 30°C relative to a day in the 5°C–10°C range increases MRPK dispersion by about 0.29 log points, implying a 0.10 percent decline in annual aggregate TFP. By contrast, the direct technology effect is much smaller: for heat shocks, the implied reduction in physical productivity is only about one-fifth as large as the loss operating through misallocation. The main empirical message is therefore not simply that temperature lowers average productivity, but that it also makes the cross-sectional allocation of capital less efficient. This distinction clarifies our relationship to the existing literature: previous *micro* estimates isolate, at best, the technology channel and miss the misallocation channel entirely, while previous *macro* estimates measure the sum of

the two.

A second empirical result is that the misallocation effect of heat is substantially larger in hotter and more economically developed regions. This pattern is important for interpreting climate damages over the long run. It implies that the allocative consequences of warming do not alleviate with development, and it suggests that market-based adaptation may protect the average firm without eliminating cross-firm differences in exposure and adjustment. These heterogeneous effects also help bridge the micro–macro disconnect in the climate damages literature: once the misallocation channel is incorporated, our micro-based estimates line up closely with corresponding macro estimates across countries, whereas without misallocation the micro estimates are essentially flat.

We use our causal estimates to project misallocation-induced TFP losses under the SSP3-4.5 scenario, a pathway of moderate mitigation. The same heterogeneity leads to large increases in losses as more regions become both hotter and richer over time: relative to current income and climate conditions, the global cost of climate-induced misallocation reaches 36.41% of aggregate TFP by the end of the century.

To explain these facts, we build a tractable firm-dynamics model with heterogeneous firms, time-to-build capital, pre-existing capital wedges, and costly endogenous adaptation. Firms differ persistently in productivity, product exposure to climate, and capital wedges, and they choose capital and adaptation under risk before temperature is realized. We model adaptation in reduced form as a quasi-fixed cost that lowers the loading of TFP on climate, implying that adaptation is complementary to scale: firms with a larger expected profit base or a lower risk premium adapt more and invest more capital. Because firms differ in exposure, the same temperature surprise generates heterogeneous realized productivities and returns to capital.

The model highlights two mechanisms. A *level* effect implies that climates farther from the bliss point magnify idiosyncratic climate damages (e.g., a sudden blizzard or fire at the plant) and hence the scope for misallocation. A *forecast-and-risk* effect implies that forecast errors and forecast uncertainty translate into cross-firm dispersion in realized returns because firms differ in effective exposure to temperature. As a result, hotter and richer economies may adapt more on average, yet still exhibit greater dispersion in adaptation across firms, which widens the distribution of effective exposure and makes the misallocation channel stronger.

The data support these mechanisms. Using Indian plant-level air-conditioning adoption as a proxy for adaptation, we show that firms with more capital and higher lagged MRPK adapt more. Using global data, we also show that heat shocks shift MRPK differentially across firms depending on size and distortions, that TFP volatility is U-shaped in long-run temperature, and that larger forecast errors and forecast uncertainty also lead to greater MRPK dispersion. These patterns line up with the model's central implication that common temperature realizations generate dispersed returns because firms differ in exposure, adaptation, and adjustment capacity.

We then estimate the model on micro data from the EU, China, and India and use it to quantify the aggregate importance of the channel. The estimated model implies that climate-induced misallocation lowers global TFP by about 12.5 percent in an average year over 1981–2019, and that this channel accounts for about 9 percent of cross-country TFP differences. Most of this global cost comes from persistent climate differences rather than from short-run forecast errors. The broader implication is that measuring climate damages through average firm-level productivity

losses alone misses an important aggregate margin. Climate policy and adaptation policy affect allocative efficiency not only by protecting the representative firm, but also by shaping the dispersion of firms' exposure and adjustment across the cross section.

Contributions to the Literature. This paper contributes to three literatures. First, we reconcile aggregate and micro estimates of climate damages. Macro estimates imply large losses from warming (Dell, Jones, and Olken 2012; Burke, Hsiang, and Miguel 2015; Lemoine 2018; Carleton et al. 2022; Nath, Ramey, and Klenow 2023; Bilal and Känzig 2024), whereas firm-level estimates often find more modest direct productivity effects (Addoum, Ng, and Ortiz-Bobea 2020; Zhang et al. 2018; Somanathan et al. 2021; Tarsia 2023; Ponticelli, Xu, and Zeume 2023). We show that part of the micro-macro gap reflects the missing allocative margin: temperature shocks widen within-region dispersion in the marginal revenue product of capital and thereby lower aggregate TFP even when average within-firm productivity effects are modest. Once it is incorporated, our micro-based estimates align closely with corresponding macro estimates across countries.

Second, we contribute to the recent climate-macro literature (Nath 2025; Cruz and Rossi-Hansberg 2023; Bakkensen and Barrage 2021; Casey, Fried, and Gibson 2022; Rudik et al. 2021), where models typically abstract from within-region distortions and therefore from climate-induced factor misallocation. We contribute a sufficient-statistics framework that measures this channel in the data and a tractable dynamic model with endogenous adaptation that rationalizes the reduced-form evidence. Closest to our paper is Caggese et al. (2023). Relative to that work, we provide direct causal evidence on within-region dispersion in marginal products using historical temperature variation and firm-level data from 32 countries.

Third, the paper contributes to the literature on misallocation (Restuccia and Rogerson 2008; Hsieh and Klenow 2009; Gopinath et al. 2017; David and Zeke 2021; Asker, Collard-Wexler, and De Loecker 2014; David and Venkateswaran 2019; Baqaee and Farhi 2019) and to the growing literature that uses quasi-experimental variation to identify its drivers (Sraer and Thesmar 2023; Bau and Matray 2023). We add environmental shocks—and climate change more broadly—to the set of forces that can generate economically meaningful capital misallocation.

The rest of the paper proceeds as follows. Section 2 develops the climate–TFP accounting framework. Sections 3 and 4 describe the data and present the reduced-form evidence. Sections 5, 6, and 7 develop, test, and estimate the dynamic model. Section 8 concludes.

2 A Framework for Climate TFP Accounting

This section develops a measurement framework that links firm-level climate exposure to region-sector aggregate TFP. We extend Hsieh and Klenow (2009)'s distorted closed-economy model of heterogeneous firms to allow climate to affect firm-specific productivity, demand, and wedges. In this framework, climate affects aggregate productivity through two conceptually distinct channels: changes in firm-level productivity, which we call *technology*, and changes in the dispersion of marginal products induced by wedges, which we call *misallocation*. Beyond the demand and production functional forms, we allow for *arbitrary* climate-impact heterogeneity. This flexibility lets the empirical strategy capture the total contribution of all mechanisms through which climate af-

fects aggregate TFP without committing to a specific microfoundation for climate-related wedges. Our goal here is not to explain the origin of those wedges but to derive the sufficient statistics that guide the empirical strategy in Section 4. Section 5 then turns to the question of *why* climate changes these objects.

2.1 Model preliminaries

Consider an economy with many regions r and sectors s. We study the aggregation of firm-level activity within a region-sector pair $n = (r, s)$. We allow the fundamentals of firm i in market n to be arbitrary (smooth) functions of current and past regional climate conditions, $\tilde{\mathbf{T}}_{rt}$, aggregate states, $\tilde{\mathbf{X}}_{nt}$, and idiosyncratic firm states, $\tilde{\mathbf{Z}}_{nit}$.[1] This general formulation accommodates a wide range of structural models within our accounting framework.

2.2 Environment

Region-sector output is a CES aggregate of differentiated firm varieties of measure J_n:[2]

$$Y_{nt} = \left(\int_0^{J_n} B_{nit}^{\frac{1}{\sigma_n}} Y_{nit}^{\frac{\sigma_n - 1}{\sigma_n}} \, di \right)^{\frac{\sigma_n}{\sigma_n - 1}}, \tag{1}$$

where $\sigma_n > 1$ is the elasticity of substitution across varieties. The implied inverse demand for firm i is $Y_{nit} = B_{nit} Y_{nt} \left(\frac{P_{nit}}{P_{nt}} \right)^{-\sigma_n}$, where $P_{nt} = \left(\int_0^{J_n} B_{nit} P_{nit}^{1-\sigma_n} \, di \right)^{\frac{1}{1-\sigma_n}}$ is the region-sector price index. The demand shifter $B_{nit} \equiv B_{ni}(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}, \tilde{\mathbf{Z}}_{nit})$ is a firm-specific function of climate and economic conditions, capturing how certain goods or services may be less desirable in hotter climates.[3]

Each firm produces with a constant-returns-to-scale Cobb-Douglas technology,

$$Y_{nit} = A_{nit} K_{nit}^{\alpha_{Kn}} L_{nit}^{\alpha_{Ln}}, \tag{2}$$

where $\alpha_{Kn} + \alpha_{Ln} = 1$. Physical productivity is firm specific and given by $A_{nit} \equiv A_{ni}(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}, \tilde{\mathbf{Z}}_{nit})$, allowing climate conditions to affect firms differently within and across region-sectors. This heterogeneity may reflect differences in production technologies and climate adaptation, such as rainfed versus irrigated farms in agriculture, or outdoor versus indoor air-conditioned workplaces in manufacturing and services.[4]

1. For concreteness, $\mathbf{T}_{rt}$ may include realizations of daily temperature, precipitation, and other extreme-weather events. The vectors $\mathbf{X}_{nt}$ and $\mathbf{Z}_{nit}$ may be interpreted as other aggregate and firm-specific productivity or demand shocks. Their joint data-generating process may be stochastic. Throughout, a tilde denotes the history of realizations up to date t. We do not impose any restriction on whether $\mathbf{X}_{nt}$ and $\mathbf{Z}_{nit}$ depend on $\mathbf{T}_{rt}$.

2. In theory, J_n, the local variety, could also vary over time to capture how climate might affect firm entry and exit. However, accurately measuring these dynamics within a granular region-sector pair over time is challenging due to data limitations, particularly for cross-country analysis. Therefore, the model does not address this aspect.

3. This is particularly evident in service industries where climate can sharply affect consumer behavior. For example, Zivin and Neidell (2014) find that heat shocks shift Americans from outdoor to indoor recreational activities (e.g., away from recreational fishing as found by Dundas and Haefen 2020). Anecdotal examples in food services include ice-cream parlors versus hot tea shops.

4. In agriculture, rainfed farms are more exposed to heat stress than irrigated farms (Piao et al. 2010). In manufacturing, workers in plants with air conditioning are less vulnerable to heat than those in plants without it (Somanathan et al. 2021). For evidence on cross-region heterogeneity in productivity damages due to adaptation, see Nath (2025).

Each firm faces a variety of frictions, including climate-related ones, which we represent as reduced-form, tax-like wedges. Firm i chooses its price and inputs to solve

$$\begin{aligned} \max_{P_{nit},K_{nit},L_{nit}} \quad & (1-\tau^Y_{nit})P_{nit}A_{nit}K_{nit}^{\alpha_{Kn}}L_{nit}^{\alpha_{Ln}} - (1+\tau^K_{nit})R_{nt}K_{nit} - (1+\tau^L_{nit})W_{nt}L_{nit} \\ \text{s.t.} \quad & Y_{nit} = B_{nit}Y_{nt}\left(\frac{P_{nit}}{P_{nt}}\right)^{-\sigma_n}. \end{aligned} \tag{3}$$

Here, R_{nt} is the user cost of capital and W_{nt} is the wage. The output wedge, $\tau^Y_{nit} \equiv \tau^Y_{ni}(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}, \tilde{\mathbf{Z}}_{nit})$, captures firm-specific distortions to effective output prices, including revenue taxes and markups relative to the Dixit-Stiglitz benchmark. These distortions may vary with climate conditions, for example if temperature shocks affect local market structure (Ponticelli, Xu, and Zeume 2023). Similarly, the input wedges, $\tau^F_{nit} \equiv \tau^F_{ni}(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}, \tilde{\mathbf{Z}}_{nit})$ for $F \in \{K, L\}$, capture distortions to effective input costs. These may reflect, for instance, tighter financing frictions for firms more exposed to adverse climate conditions (Ginglinger and Moreau 2023). For dynamic input choices with adjustment frictions, such as capital, they may also capture ex post investment mistakes generated by unanticipated temperature-related productivity shocks, in the spirit of Asker, Collard-Wexler, and De Loecker (2014) and David, Hopenhayn, and Venkateswaran (2016).

Within each period, total factor supplies are taken as given: $K_{nt} = \sum_i K_{nit}$ and $L_{nt} = \sum_i L_{nit}$. We call the allocation with wedges the *distorted equilibrium*, and the allocation that sets all wedges to zero while holding $\{A_{nit}, B_{nit}, K_{nt}, L_{nt}\}$ fixed the *efficient equilibrium*.[5]

Firm optimality implies that, for any input $F \in \{K, L\}$, its marginal revenue product (MRPF) satisfies:

$$\text{MRPF}_{nit} = \alpha_{Fn}\frac{\sigma_n - 1}{\sigma_n}\frac{P_{nit}Y_{nit}}{F_{nit}} = \frac{1+\tau^F_{nit}}{1-\tau^Y_{nit}}P^F_{nt}, \tag{4}$$

where $P^K_{nt} = R_{nt}$ and $P^L_{nt} = W_{nt}$. This is the key measurement result. In the efficient equilibrium, marginal revenue products are equalized within a region-sector. With wedges, cross-firm dispersion in marginal revenue products reveals input misallocation. A higher input wedge raises MRPF by discouraging input use, while a lower output wedge lowers MRPF by encouraging input use.

To focus on the main margin in the paper, the next lemma specializes to how capital wedges affect equilibrium (mis-)allocation. Let $\theta_{nit} \equiv \frac{P_{nit}Y_{nit}}{\int_0^{J_n} P_{njt}Y_{njt}\,dj}$ denote firm i's sales share, and let θ^*_{nit} denote its efficient-equilibrium counterpart.

Lemma 1 (Equilibrium (mis-)allocation of capital) *With capital distortions τ^K_{nit} only, the log ratio of firm i's distorted and efficient capital allocations satisfies*

$$\log\left(\frac{K_{nit}}{K^*_{nit}}\right) = -\underbrace{\log\left(\frac{1+\tau^K_{nit}}{1+\tau^K_{nt}}\right)}_{\textit{relative wedge effect}} + \underbrace{\log\left(\frac{\theta_{nit}}{\theta^*_{nit}}\right)}_{\textit{size effect}}, \tag{5}$$

where the efficient capital allocation is pinned down entirely by demand-adjusted productivity:

$$\frac{K^*_{nit}}{K_{nt}} = \theta^*_{nit} = \frac{B_{nit}A_{nit}^{\sigma_n - 1}}{\int_0^{J_n} B_{njt}A_{njt}^{\sigma_n-1}\,dj}.$$

5. Efficiency here is static and conditional on the period-t aggregate factor supplies.

Proof. See Appendix A.2. ■

Lemma 1 shows that wedges distort capital allocation through two margins. The first is a *relative wedge effect*: firms facing a higher capital wedge than the region-sector average hold too little capital relative to the efficient benchmark. The second is a *size effect*: wedges can also make a firm inefficiently small, reducing its capital demand relative to its efficient counterfactual. By contrast, in the efficient equilibrium, capital is allocated entirely according to firms' relative demand-adjusted productivity. This benchmark is implicit in much of the empirical literature, but it is precisely what climate-related distortions can break. We next derive the aggregate implications of these distortions.

2.3 Aggregation and TFP decomposition

We now aggregate across firms. With a continuum of firms in each region-sector, we adopt the aggregation notion of Krusell and Smith (1998) and treat the cross-sectional distribution of firm states $\tilde{\mathbf{Z}}_{nit}$ as summarized by a finite set of moments contained in $\tilde{\mathbf{X}}_{nt}$. Using a second-order approximation, we show how dispersion in micro wedges maps into aggregate productivity losses through a set of sufficient statistics.

Proposition 1 (Aggregation and TFP Decomposition.) *Each region-sector n admits an aggregate production function of the form*

$$Y_{nt} = TFP_{nt} K_{nt}^{\alpha_{Kn}} L_{nt}^{\alpha_{Ln}}, \tag{6}$$

where, up to a second-order approximation, the region-sectoral aggregate total factor productivity $TFP_{nt} := TFP_n\left(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}\right)$ can be decomposed as follows:

$$\log TFP_n(\tilde{\mathbf{T}}_{rt}, \cdot) = \underbrace{\frac{1}{\sigma_n - 1} \log \left[\mathbb{E}_i \left[B_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot) \left(A_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot) \right)^{\sigma_n - 1} \right] \right]}_{\text{Technology}(\log TFP^E_{nt})} + \frac{1}{\sigma_n - 1} \log J_n$$

$$\text{Misallocation Loss} \left\{ \begin{array}{l} \underbrace{-\frac{\sigma_n}{2} \operatorname{var}_{\log\left(1-\tau^Y_{ni}\right)}(\tilde{\mathbf{T}}_{rt}, \cdot)}_{\text{Output Wedge Dispersion}} \underbrace{- \sum_{F \in \{K,L\}} \frac{\alpha_{Fn} + \alpha_{Fn}^2(\sigma_n - 1)}{2} \operatorname{var}_{\log\left(1+\tau^F_{ni}\right)}(\tilde{\mathbf{T}}_{rt}, \cdot)}_{\text{Factor Wedge Dispersion}} \\ \underbrace{+ \sigma_n \sum_{F \in \{K,L\}} \alpha_{Fn} \operatorname{cov}_{\log\left(1-\tau^Y_{ni}\right), \log\left(1+\tau^F_{ni}\right)} \left(\tilde{\mathbf{T}}_{rt}, \cdot \right)}_{\text{Output-Factor Mixed Distortion}} \\ \underbrace{- (\sigma_n - 1) \alpha_{Kn} \alpha_{Ln} \operatorname{cov}_{\log\left(1+\tau^K_{ni}\right), \log\left(1+\tau^L_{ni}\right)} \left(\tilde{\mathbf{T}}_{rt}, \cdot \right)}_{\text{Factor Mixed Distortion}}. \end{array} \right. \tag{7}$$

Here, expectations, variances, and covariances are taken across firms within region-sector n at date t. Holding variety fixed, Proposition 1 decomposes aggregate TFP into two terms: *technology* and *misallocation loss*. The technology component, $\log \text{TFP}^E_{nt}$, represents the production possibility frontier. It is a CES aggregation of demand-adjusted physical productivity across firms. We label this term *technology* in the spirit of Basu and Fernald (2002) and Baqaee and Farhi (2019).

The remaining variance and covariance terms summarize the *misallocation loss*, that is, the TFP cost of wedges. These moments are themselves smooth functions of climate. Greater dispersion in output or factor wedges, captured by $\text{var}_{\log(1-\tau_{ni}^{Y})}$ and $\text{var}_{\log(1+\tau_{ni}^{F})}$, lowers TFP by widening marginal products across firms and thereby increasing misallocation. The resulting losses are larger when products are more substitutable (higher σ_n) and when the distorted factor has a larger cost share (higher α_{Fn}). The covariance terms capture whether different wedges reinforce or offset each other. In particular, productivity losses are larger when (i) firms facing higher taxes on output also face higher input prices, so that $\text{cov}_{\log(1-\tau_{ni}^{Y}),\log(1+\tau_{ni}^{F})} < 0$, and (ii) firms facing larger capital distortions also face larger labor distortions, so that $\text{cov}_{\log(1+\tau_{ni}^{K}),\log(1+\tau_{ni}^{L})} > 0$.

2.4 Decomposing the Impact of Climate Change on Aggregate TFP

Because all moments in equation (7) are smooth functions of climate conditions $\tilde{\mathbf{T}}_{rt}$, the first-order effect of climate change on TFP can be written in terms of a set of measurable derivatives. We illustrate this in the benchmark case with capital wedges only.

Benchmark case: only capital wedges are present. We focus on the case with capital wedges only, following a large misallocation literature (Asker, Collard-Wexler, and De Loecker 2014; David and Venkateswaran 2019; Sraer and Thesmar 2023). This benchmark is natural because capital is a dynamic input: it is often externally financed, costly to adjust, and hard to reallocate on short notice. As a result, distortions in capital use—such as those arising from credit frictions or investment irreversibility—can be larger and more persistent than distortions in labor use (Gorodnichenko et al. 2018). This focus is therefore motivated ex ante by economic structure and tractability; later sections show empirically that capital is the quantitatively dominant margin, while other channels are comparatively small.

When only the capital wedge is present, the MRPK of each firm i satisfies that: $\text{MRPK}_{nit} = \alpha_{Kn}\frac{\sigma_n-1}{\sigma_n}\frac{P_{it}Y_{it}}{K_{it}} \propto (1+\tau_{nit}^{K})R_{nt}$. Hence, dispersion in capital wedges is directly measurable from the cross-sectional dispersion of MRPK:

$$\text{var}_i\left[\log(1+\tau_{nit}^{K})\right] = \text{var}_i(mrpk_{nit}) = \text{var}_i\left(\log\frac{P_{nit}Y_{nit}}{K_{nit}}\right), \tag{8}$$

where $mrpk_{nit} \equiv \log(\text{MRPK}_{nit})$. The last equality shows that, conditional on a common technology within each region-sector-year, the cross-sectional variance of log MRPK equals the variance of log sales over capital stock. We can now formalize how climate conditions affect aggregate TFP via the following Lemma:

Lemma 2 *With only capital wedges, the first-order effect of climate conditions on region-sector TFP can be*

written as

$$\frac{d\log TFP_{nt}}{d\tilde{\mathbf{T}}_{rt}} = \underbrace{\frac{1}{\sigma_n - 1}\frac{d\log \mathbb{E}_i\left[B_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)\left(A_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)\right)^{\sigma_n-1}\right]}{d\tilde{\mathbf{T}}_{rt}}}_{\textit{Technology channel}} - \underbrace{\frac{\alpha_{Kn}+\alpha_{Kn}^2(\sigma_n-1)}{2}\frac{d\,\mathrm{var}_{mrpk_{ni}}(\tilde{\mathbf{T}}_{rt},\cdot)}{d\tilde{\mathbf{T}}_{rt}}}_{\textit{Misallocation Channel}}, \tag{9}$$

Lemma 2 gives the empirical content of the framework. The first term in (9) captures the direct technology channel, while the second captures the misallocation channel. In the capital-wedge benchmark, measuring climate-induced misallocation reduces to estimating how climate changes within-cell MRPK dispersion. This is the key reduced-form object in Section 4. In practice, we estimate local semi-elasticities around the observed $\tilde{\mathbf{T}}_{rt}$ and $\tilde{\mathbf{X}}_{nt}$. Section 4.4 therefore allows these local effects to vary with long-run climate and development.

3 Data

3.1 Global Firm-level Microdata

We compile firm-level microdata from 30 European countries, China, and India, collectively covering 38.6% of world GDP. European firm data come from Bureau van Dijk's (BvD) Orbis database; Chinese data from the National Bureau of Statistics (NBS) Annual Survey of Industrial Firms; and Indian data from the Annual Survey of Industries (ASI). All three sources are widely used in the literature and provide nationally representative coverage of the formal sector.[6] Table B.1 lists countries, sample periods, and sources.

Key variables and measurement. The central object in our empirical analysis is the cross-firm dispersion of the marginal revenue product of capital, var(log MRPK), within each region–sector–year cell. Constructing this sufficient statistic requires firm-level revenue $P_{it}Y_{it}$ and capital stock K_{it}. We measure revenue as reported operating revenue in Orbis and China NBS, and as total sales in India ASI. Capital is measured as the book value of gross fixed assets, with one exception: India ASI reports net fixed assets more consistently, so we use net values for India.[7] The ASI also records whether a plant has computerized air-conditioning (available 2001–2014), which we exploit as a plant-level adaptation proxy in Section 6. All three datasets additionally report wage bills, employment, and (with varying coverage) material costs, which we use in robustness exercises and in the technology-channel estimation of Section 4.3. We harmonize sectoral classifications into eight USSIC major divisions and define regions as NUTS 3 areas in Europe, prefectures in China,

6. Orbis coverage is regulatory across Europe, extending to small private firms. The China NBS surveys cover all industrial firms with annual sales above CNY 5 million (1998–2007), representing over 90.7% of total output. The India ASI censuses plants with 100+ workers and samples smaller registered plants, ensuring representativeness at the state–industry level.

7. Gross fixed-asset values in ASI contain substantial missing data. In Appendix B we show that our results are not sensitive to this choice.

and districts in India—units roughly comparable to U.S. counties.[8]

Sample construction. We retain all firm-year observations with non-missing revenue and capital stock, and trim the top and bottom 0.1% of MRPK within each country.[9] For region–sector–year level analysis, we further restrict to cells containing at least 30 firm-year observations, so that within-cell variance estimates are not dominated by sampling noise.[10] The resulting panel is an unbalanced sample of 124,567 region–sector–year observations covering 76,826,956 firm-years. For firm-level regressions, we use the full sample after trimming. Appendix B provides details on cleaning, variable construction, and sample selection for each dataset.

3.2 Weather and Forecast Data

Temperature. We obtain daily temperature from ERA5-Land (Sabater 2019), a reanalysis product from ECMWF that provides hourly surface variables at $0.1°\times0.1°$ resolution (approximately 9 km) with global coverage. We compute daily mean temperature as the arithmetic average of the daily maximum and minimum, then construct annual temperature bins at 5°C intervals from below -5°C to above 30°C. Each bin counts the number of days per year a region's average temperature falls within that range.

Temperature bins capture the distributional shifts that define climate change—in particular, the increasing frequency of extreme heat events (Oudin Åström et al. 2013; Christidis, Mitchell, and Stott 2023)—more accurately than changes in annual means alone. Figure 1a documents that hot days (above 25°C) have increased substantially across all three sample regions relative to 1951–1980 baselines, with the shift accelerating in the most recent decade. Figure 1b shows the corresponding rightward shift in the full daily temperature distribution, including projections for 2100.

Climate projections. End-of-century temperature distributions come from the GFDL-ESM4 model under CMIP6's SSP3-4.5 scenario, which represents intermediate emissions with modest mitigation—a pathway widely regarded as among the more plausible (O'Neill et al. 2016).[11]

Weather forecasts. We use ECMWF's long-range (seasonal) forecasts (Copernicus Climate Change Service and Climate Data Store 2018), which provide daily maximum and minimum 2 m temperature forecasts up to seven months ahead at $1°\times1°$ resolution. We extract 30-day-ahead forecasts issued on the first day of each month.[12] Details on how we map these forecasts to annual forecast errors and perceived forecast uncertainty are in Section 7 and Appendix B.

8. Industries are: agriculture, mining, construction, manufacturing, transportation & utilities, wholesale trade, retail trade, FIRE, and services. We adopt USSIC for cross-country comparability, following its availability in Orbis.

9. Trimming extreme log revenue-to-capital ratios is standard in the misallocation literature: Hsieh and Klenow (2009) trim at 1% and Gopinath et al. (2017) apply similar filters, as tail values in administrative data typically reflect reporting errors, dormant firms, or asset revaluations. Our 0.1% threshold is conservative relative to these precedents.

10. Results are robust to alternative thresholds of 5, 10, and 20 firm-year observations per cell. We also drop cells exhibiting discontinuous jumps in firm counts or aggregate sales, which likely reflect changes in reporting requirements rather than economic activity.

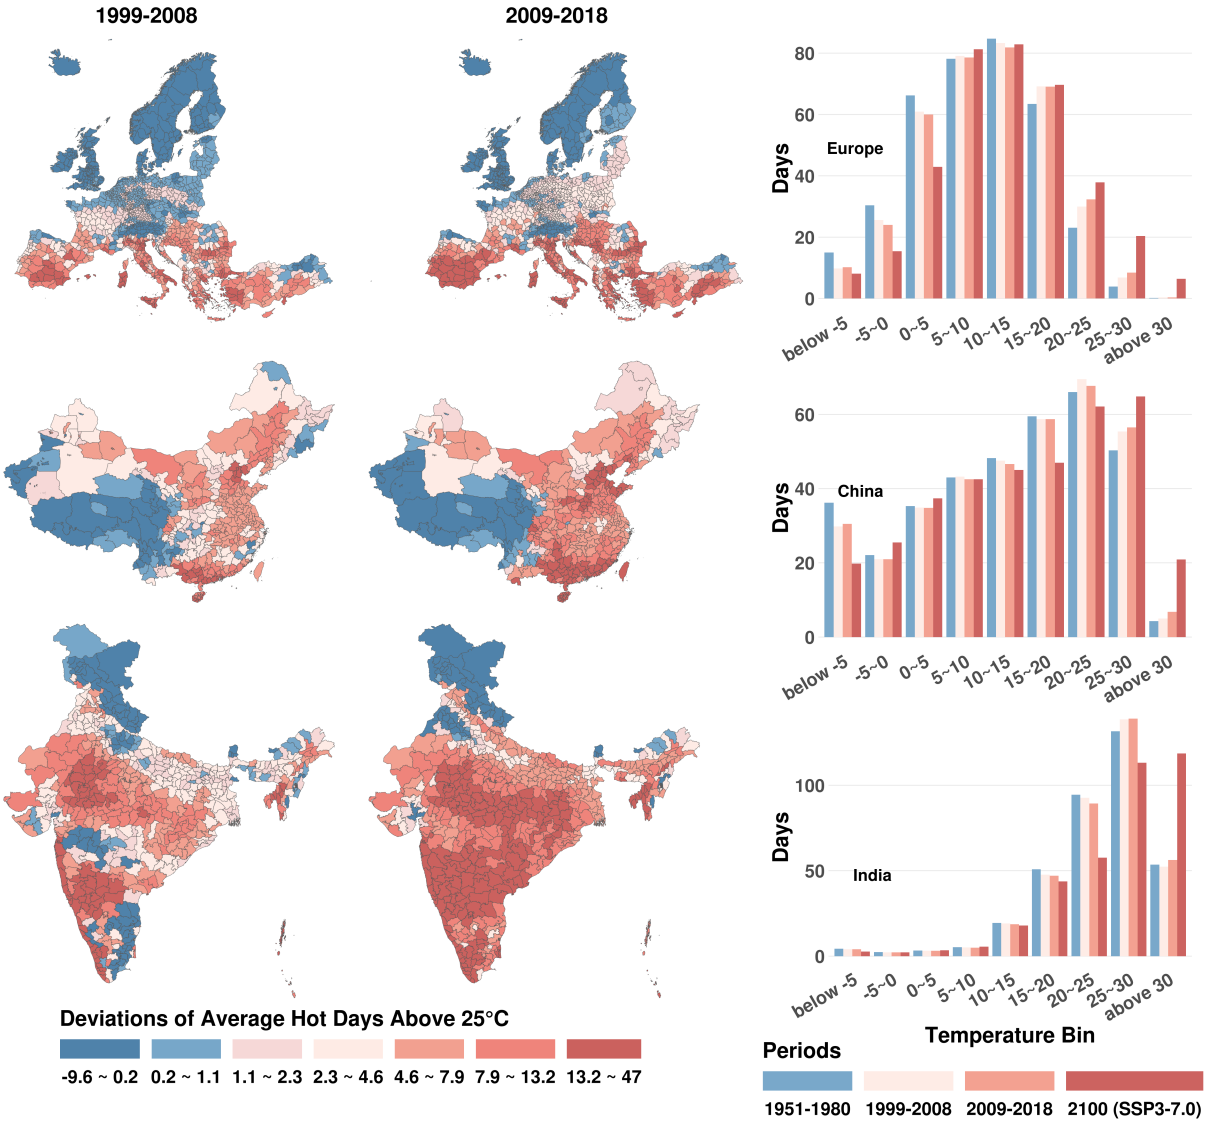
11. GFDL-ESM4, produced by NOAA's Geophysical Fluid Dynamics Laboratory, has a 1-degree nominal horizontal resolution.

12. Pre-2017 values are hindcasts from a 25-member ensemble; from 2017 onward, forecasts are produced monthly with a 51-member ensemble.

Figure 1: Climates of our sample regions

(a) Number of Hot Days Deviation Relative to 1951–1980

(b) Average Daily Temperature Distribution



Notes: Panel (a) plots the change in annual days above 25°C between each sample decade (1999–2008 and 2009–2018) and the 1951–1980 baseline. Panel (b) shows average daily temperature distributions for the baseline, two sample periods, and 2100 under SSP3-7.0.

3.3 Other Data

Regional GDP. We obtain sub-national GDP per capita (PPP-adjusted, 2017 international dollars) from the DOSE dataset (Wenz et al. 2023), matched to our regional definitions via spatial coordinates. Using sub-national rather than country-level income is important because countries like India and China exhibit large within-country income variation across districts and prefectures. We compute long-run regional averages over our sample period and use them in two ways: as a time-invariant moderator interacted with temperature bins in the heterogeneous-effects specification (Section 4.4), and as the key input to the adaptation-cost parameterization in the structural model (Section 7).

Income projections. National income per capita projections under the SSP scenarios are from the OECD EnvGrowth model (Dellink et al. 2017), hosted by the International Institute for Applied Systems Analysis. These projections enter the end-of-century decomposition (Section 4.6), where projected income growth drives the "income effect" component of future misallocation losses.

4 Estimating the Misallocation Channel of Climate Change

This section documents the paper's central empirical fact: temperature shocks increase capital misallocation within region–sector cells. Section 4.1 describes the empirical design. Section 4.2 reports the average effect of temperature on MRPK dispersion, showing a U-shaped relationship that is robust across specifications. Section 4.3 decomposes aggregate TFP into technology and misallocation channels, finding that the misallocation channel is much larger than the technology channel, especially for heat shocks. Section 4.4 documents that misallocation losses are larger in hotter and richer regions, a pattern that sharply distinguishes this channel from conventional climate damages. Section 4.5 shows that incorporating the misallocation channel reconciles micro-based and macro-based climate damage estimates. Section 4.6 projects end-of-century TFP losses under SSP3-4.5. Appendix C reports additional robustness checks.

4.1 Empirical design

We estimate how annual temperature realizations affect capital misallocation within region–sector cells, measured by the cross-firm variance of log MRPK. Following Deschênes and Greenstone (2011) and Carleton et al. (2022), we summarize each region-year temperature realization by the number of days in 5°C bins. Specifically, we define current climate conditions as the vector $\mathbf{T}_{r,t} = \left\{\text{Tbin}_{r,t}^{<-5^\circ\text{C}}, \text{Tbin}_{r,t}^{-5\sim0^\circ\text{C}}, ..., \text{Tbin}_{r,t}^{>30^\circ\text{C}}\right\}$, where each element $\text{Tbin}_{r,t}^{b}$ counts the number of days in year t for which region r's average temperature falls within bin b. This approach uses detailed daily temperature data while allowing for flexible nonlinear climate effects in a linear regression framework.

Our baseline specification is

$$\sigma^2_{mrpk_{(s,r),t}} = \sum_{b \in B/\{5\sim10^\circ\text{C}\}} \lambda^{b}_{\sigma^2_{mrpk}} \times \text{Tbin}^{b}_{r,t} + \boldsymbol{\delta}_{\sigma^2_{mrpk}} \mathbf{X}_{s,r,t} + \alpha_{c(r),t} + \eta_{s,r} + \varepsilon_{r,s,t}, \tag{10}$$

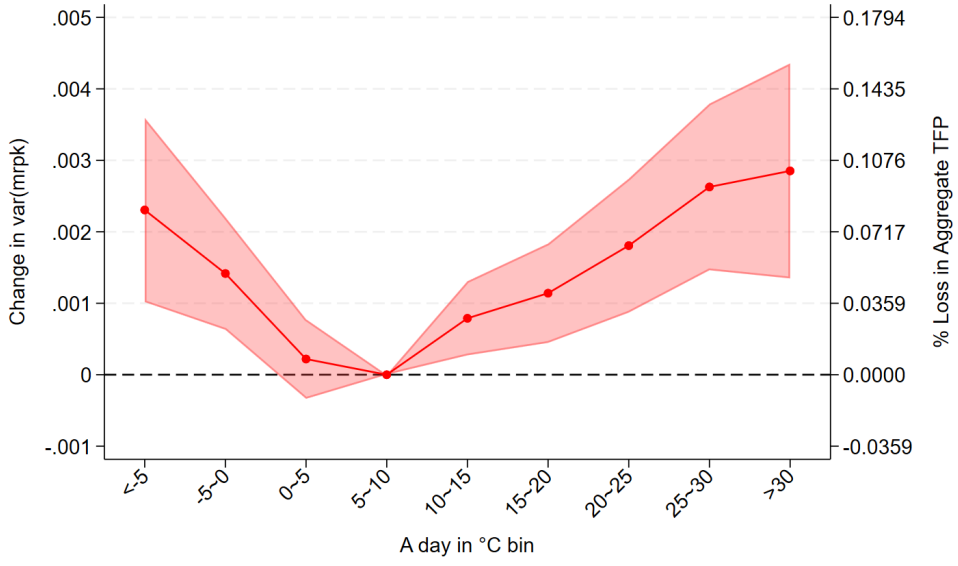
where $\sigma^2_{mrpk_{(s,r),t}}$ denotes the variance of log MRPK across firms within region–sector (s, r) in year t. Region–sector fixed effects $\eta_{s,r}$ absorb time-invariant differences across cells, while country–year fixed effects $\alpha_{c(r),t}$ absorb national shocks common to all regions within a country-year. The coefficients on the temperature bins are therefore identified from year-to-year fluctuations in local temperature realizations within a region–sector, relative to other regions in the same country and year. Such climate variation is plausibly exogenous, and $\boldsymbol{\lambda}_{\sigma^2_{mrpk}} = \mathbb{E}\left[d\,\mathrm{var}_{mrpk_{ni}}(\tilde{\mathbf{T}}_{rt}, \cdot)/d\mathbf{T}_{r,t}\right]$ captures the average effect of climate on misallocation. Because the temperature bins sum to 365 days, coefficients are interpreted relative to the omitted 5°C–10°C bin. Standard errors are clustered at the region level.

In our preferred specification, we omit time-varying region–sector controls. Variables such as the number of firms, average sales, and average MRPK may themselves respond to temperature shocks, so conditioning on them risks absorbing part of the total effect. Table 1 shows that the results are very similar when these controls are added.

4.2 Average effects of temperature on capital misallocation

The baseline estimates of $\hat{\lambda}^b_{\sigma^2_{mrpk}}$ and the implied TFP losses are reported in columns (1) and (5) of Table 1 and shown in Figure 2. We convert $\hat{\lambda}^b_{\sigma^2_{mrpk}}$ into implied TFP losses using the relationship $-\frac{\alpha_{Kn}+\alpha^2_{Kn}(\sigma_n-1)}{2}\hat{\lambda}^b_{\sigma^2_{mrpk}}$ in equation (9), adopting the standard parameter values of capital share $\alpha_{Kn} = 0.35$ and elasticity of demand $\sigma_n = 4$ (Redding and Weinstein 2020) for all region-sector n.

Figure 2: Estimated impact of daily temperature shocks on annual MRPK dispersion and implied TFP loss



Notes: This figure shows the aggregate impact linking annual MRPK dispersion and TFP loss to average daily temperatures. The formula for computing TFP loss from misallocation is $-\frac{\alpha_{Kn}+\alpha^2_{Kn}(\sigma_n-1)}{2}\hat{\lambda}^b_{\sigma^2_{mrpk}}$. The estimation is normalized by setting the range of 5°C - 10°C as the reference category. Therefore, each $\lambda^b_{\sigma^2_{mrpk}}$ represents the estimated effect of an additional day in temperature bin b on annual MRPK dispersion or TFP loss, relative to a day with temperatures between 5°C - 10°C. The figure also includes the 90% confidence interval for these estimates where standard errors are clustered at the region level.

Table 1: Effects of daily mean temperature bins on MRPK dispersion and TFP loss

	(1) σ^2_{mrpk}	(2) σ^2_{mrpk}	(3) σ^2_{mrpk}	(4) σ^2_{mrpk}	(5) Implied TFP Loss
< -5°C	0.0023***	0.0022***	0.0024***	0.0002	0.0827***
	(0.0008)	(0.0008)	(0.0008)	(0.0018)	(0.0281)
$-5 \sim 0$°C	0.0014***	0.0011**	0.0015***	-0.0002	0.0508***
	(0.0005)	(0.0005)	(0.0005)	(0.0011)	(0.0172)
$0 \sim 5$°C	0.0002	-0.0000	0.0003	-0.0014	0.0079
	(0.0003)	(0.0003)	(0.0003)	(0.0008)	(0.0121)
$10 \sim 15$°C	0.0008**	0.0008**	0.0008**	0.0039***	0.0283**
	(0.0003)	(0.0003)	(0.0003)	(0.0008)	(0.0113)
$15 \sim 20$°C	0.0011***	0.0012***	0.0011***	0.0040***	0.0409***
	(0.0004)	(0.0004)	(0.0004)	(0.0010)	(0.0151)
$20 \sim 25$°C	0.0018***	0.0019***	0.0017***	0.0079***	0.0648***
	(0.0006)	(0.0006)	(0.0006)	(0.0016)	(0.0204)
$25 \sim 30$°C	0.0026***	0.0027***	0.0025***	0.0088***	0.0943***
	(0.0007)	(0.0007)	(0.0007)	(0.0017)	(0.0254)
> 30°C	0.0029***	0.0029***	0.0028***	0.0080***	0.1023***
	(0.0009)	(0.0009)	(0.0009)	(0.0017)	(0.0328)
Controls	No	No	Yes	No	No
Region-Sector FE	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	No	Yes	Yes	Yes
Country-Sector-Year FE	No	Yes	No	No	No
1995 VA Weighted	No	No	No	Yes	No
Observations	122,055	121,501	122,055	121,843	122,055
R^2	0.875	0.908	0.878	0.902	0.875

Notes: Standard errors in parentheses.

Our estimates show a clear U-shaped relationship between temperature and capital misallocation. Both extreme heat and extreme cold raise MRPK dispersion, with the largest effects at the tails of the temperature distribution. This pattern translates into an inverted U-shape between TFP and temperature, a well-known result in the climate econometrics literature (Burke, Hsiang, and Miguel 2015).

Relative to a day in the 5°C–10°C bin, an additional day above 30°C raises the MRPK dispersion by 0.29 log points and implies a 0.10% reduction in annual TFP, equivalent to roughly 37% of one day's GDP. An additional day below −5°C raises the MRPK dispersion by about 0.23 log points, corresponding to an annual TFP loss of 0.08%. These estimates are large: larger than the firm-level productivity effect in the worst-hit countries in the global sample of Nath (2025) and orders of magnitude larger than the estimates by Ponticelli, Xu, and Zeume (2023) in the US.

Robustness. The estimates are stable across alternative fixed effects, the inclusion of controls, and value-added weighting (Table 1). Appendix C reports seven additional exercises. First, to examine whether the average response is driven by a single broad industry, we re-estimate the specification sector by sector; Figure C.1 shows that the heat response is strongest in agriculture, manufacturing, transportation, services, and retail. Second, one potential concern is that temperature affects every input return margin, not capital alone. We re-run the specification using MRPL and MRPM dispersion as outcomes; Figures C.2 and C.3 show no comparable U-shape, consistent with the model's timing in which labor and materials adjust after temperature is realized. Third, to examine whether the result depends on the construction of weather exposure, we reconstruct the temperature variable using daily maximum temperature, population-weighted temperature, and a cubic polynomial in daily average temperature; Figure C.4 shows the U-shape is preserved under each alternative. Fourth, inference may be sensitive to the standard-error correction. We report spatial-HAC standard errors at varying distance bandwidths alongside alternative clustering choices; Figure C.5, Table C.2, and Figure C.6 show the heat effects remain statistically significant. Fifth, to examine whether the region-sector cell definition drives the result, we re-estimate at NUTS 2–SIC 1, NUTS 3–SIC 2, and, where available, NUTS 3–SIC 3 cells; Figures C.7, C.8, and C.10 show qualitatively similar estimates at each alternative. Sixth, to examine whether extreme MRPK observations drive the estimates, we winsorize MRPK more aggressively; Figure C.12 shows the U-shape remains. Finally, to examine whether the result relies on Cobb–Douglas production, we measure capital-wedge dispersion in a enrionment with CES production function with factor-augmenting productivity using labor–capital or materials–capital ratios; Figures C.14 and C.15 show the U-shape is preserved and Cobb-Douglas specification gives a lower bond in the TFP cost. Overall, the U-shaped response of MRPK dispersion to temperature is robust to sectoral composition, the choice of input margin, weather construction, inference procedure, region–sector definition, MRPK trimming, and the CES generalization of production.

4.3 Technology vs. Misallocation

We next compare the magnitude of climate-induced misallocation with the technology channel—the response of the CES productivity aggregate to temperature. Our decomposition in equation (9) identifies the two channels separately, but mapping firm-level microdata to an estimate

that is comparable in kind to the misallocation regression is not automatic. With the region-sector-year cell indexed by (n,t), where $n=(r,s)$ follows the notation introduced in Section 2, the comparability rests on two design choices: the first is theory-driven and fixes what object we target; the second is sample-driven and fixes how we weight cells against each other.

The relevant technology object is defined at the cell, not at the firm. Within cell (n,t), we define the CES aggregate:

$$M_{nt}(\tilde{\mathbf{T}}_{rt}) := \mathbb{E}_i\left[\text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}\,\middle|\,n,t\right], \qquad \text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot) := B_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\frac{1}{\sigma_n-1}} A_{ni}(\tilde{\mathbf{T}}_{rt},\cdot),$$

where the expectation is taken over firms within the cell. The cell-level technology channel is $1/(\sigma_n-1)$ times the elasticity of this aggregate:

$$\frac{1}{\sigma_n-1}\frac{d\log M_{nt}(\tilde{\mathbf{T}}_{rt})}{d\tilde{\mathbf{T}}_{rt}}.$$

This is a cell-level aggregate response—not a firm-level elasticity, and not the equal-weight average of firm-level elasticities. The misallocation regression in equation (10) is also defined at the cell, so the technology estimate we need is the one that matches this level of aggregation.

The first design choice—what to target—follows from the theory: we want the elasticity of the aggregate, not the aggregate of elasticities. The elasticity of the aggregate is a productivity-share-weighted average of firm-level elasticities, which is a different object from the equal-weight average of firm elasticities:

$$\underbrace{\frac{d\log \mathbb{E}_i\left[\text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}\right]}{d\tilde{\mathbf{T}}_{rt}}}_{\text{Elasticity of the aggregate}} = \mathbb{E}_i\left[\frac{\text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}}{\mathbb{E}_i\left[\text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}\right]}\cdot\frac{d\log \text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}}{d\tilde{\mathbf{T}}_{rt}}\right] \neq \underbrace{\mathbb{E}_i\left[\frac{d\log \text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}}{d\tilde{\mathbf{T}}_{rt}}\right]}_{\text{Average of firm elasticities}}.$$

The two coincide only under homogeneous firm responses or, degenerately, uniform firm sizes—both of which are assumptions we want to relax. A firm-level OLS regression recovers the right-hand side and is therefore the wrong target whenever a productive firm (thus with higher weight in the aggregate) respond differently from the average firm, which is precisely the margin opened by heterogeneous exposure and adaptation.

Following Tyazhelnikov, Zhou, and Shi (2024), we estimate the left-hand side by Poisson pseudo-maximum likelihood (PPML): M_{nt} is a conditional mean in levels, and PPML is the natural estimator of conditional means in levels, with first-order conditions that match fitted sums to observed sums within each conditioning cell; under the exponential specification in equation (11) below, the bin coefficient β_b then reads off as the elasticity of M_{nt} with respect to Tbin^b_{rt}. This is what makes the technology estimate comparable in kind to the cell-level misallocation regression—both are defined at the cell and take firm heterogeneity as an input rather than assuming it away, whereas a firm-level elasticity is neither.

Our baseline technology specification is built on this PPML target. Define the firm-level outcome:

$$y_{nit} := \widetilde{\text{TFP}}_{nit}^{\sigma_n-1}, \qquad \widetilde{\text{TFP}}_{nit} := \frac{(P_{nit}Y_{nit})^{\sigma_n/(\sigma_n-1)}}{K_{nit}^{\alpha_{Kn}}L_{nit}^{\alpha_{Ln}}},$$

and let o_{nt} denote the cell offset constructed from region-sector-year revenue.[13] We estimate:

$$\mathbb{E}\Big[y_{nit} \,\Big|\, \tilde{\mathbf{T}}_{rt}, o_{nt}, \eta_i, \kappa_{c(r),s,t}\Big] = \exp\left[o_{nt} + \eta_i + \kappa_{c(r),s,t} + \sum_{b \in B/\{5\sim 10^\circ\mathrm{C}\}} \beta_b \,\mathrm{Tbin}^b_{rt}\right]. \tag{11}$$

The PPML first-order conditions deliver within-cell CES weighting of firm responses, so the within-cell aggregation preserves firm heterogeneity rather than averaging it out under a homogeneous-response assumption.

The second design choice—how to weight cells against each other—is driven by the sample. Our panel covers 32 countries, with economies ranging from China and France down to much smaller ones, and with uneven numbers of firm-years across cells. Standard PPML implicitly weights each cell by its baseline outcome mass (formally, Proposition 7 in Appendix D), so the estimate would be dominated by cells that contribute more mass: France and China because their aggregates are larger, and other cells when they simply contribute more observations. We want the average cross-country response, not a mass-weighted one, so we attach analytic weights a_{nit} built from the cell mean:

$$a_{nit} := \frac{1}{N_{nt}\,\bar{y}_{nt}}, \quad \text{where } \bar{y}_{nt} := \frac{1}{N_{nt}} \sum_{j \in (n,t)} y_{njt}.$$

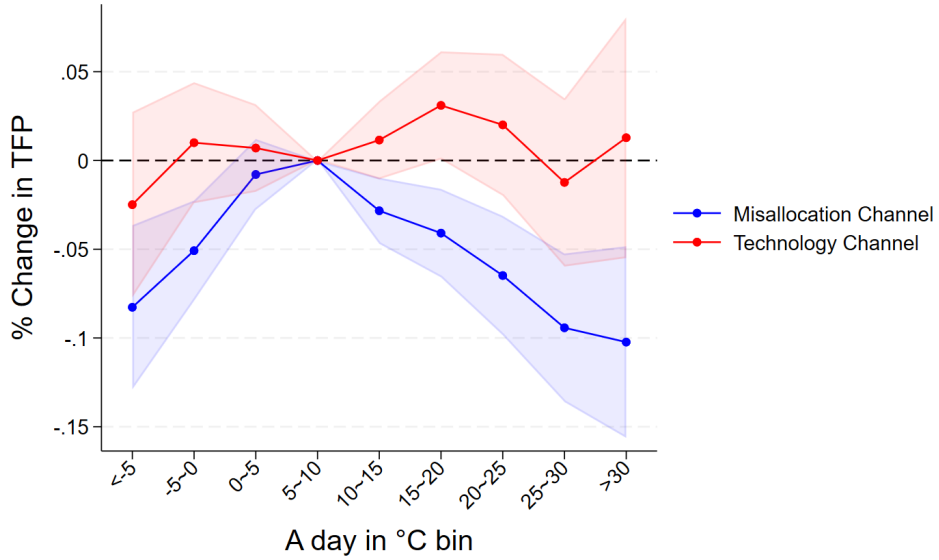
The factor $1/\bar{y}_{nt}$ neutralizes the baseline mass that tilts standard PPML toward larger economies; the factor $1/N_{nt}$ neutralizes the observation count so that cells with more firm-years do not count more. Therefore, the weighted PPML preserves the CES-relevant productivity weighting of firms within each cell, while assigning equal weight across cells in the final estimating objective.

These two corrections together make the technology estimate directly comparable to the misallocation estimate from equation (10). Appendix D proves that both estimators consistently estimate objects at the same level of aggregation. Concretely, $\hat{\beta}_b/(\sigma_n - 1)$ identifies the average cell-level *aggregate technology* effect and $\hat{\lambda}^b_{\sigma^2_{mrpk}}$ identifies the corresponding average cell-level *misallocation* effect, so the two coefficients can be read against each other directly.

The misallocation channel dominates the technology channel most sharply for heat shocks, where the loss is roughly five times as large as the technology channel. Figure 3 plots the two responses side by side. At the upper tail, an additional day above 30°C has null effect on aggregate productivity through the technology channel, against a 0.10% TFP loss through misallocation. At the lower tail, an additional day below -5°C lowers productivity by about 0.02% through technology, against a 0.08% loss through misallocation. Overall, the technology channel is quantitatively small, statistically insignificant, and a minor contributor once the misallocation channel is accounted for.

13. Since $\mathrm{TFP}_{nit} = \widetilde{\mathrm{TFP}}_{nit}/(P_{nt}^{\sigma_n} Y_{nt})$, assuming a common sectoral price index within a country, it suffices to use the offset $\log(P_{nt}Y_{nt})$, with the remaining country-sector-year price normalization absorbed by fixed effects.

Figure 3: Technology vs. Misallocation



Notes: The graphs include 90% confidence intervals, and standard errors are clustered at the regional level. The reference temperature is at 5∼10°C.

These magnitudes also clarify what previous studies do and do not capture. Micro studies that estimate firm-level productivity responses to temperature typically recover the equal-weight average firm elasticity and are silent on the misallocation channel. Our weighted-PPML estimate targets the average cell-level aggregate technology effect, and our misallocation regression recovers the channel those micro studies omit. Macro studies capture the sum of both channels but cannot separate them.

4.4 Heterogeneous effects by long-run climate and development

The average effect masks substantial heterogeneity. We next ask whether the same temperature shock has different effects in places with different long-run climate conditions and levels of economic development. A priori, the answer is ambiguous: warmer or richer regions may be better adapted on average, but they may also exhibit greater cross-firm heterogeneity in exposure and adjustment capacity. In warmer climates, heat-averse firms may have already adapted, so additional heat shocks cause only limited disruption; alternatively, once temperatures exceed adaptation thresholds, limited adaptability (Moscona and Sastry 2023) and the convexity of heat damages (Burke, Hsiang, and Miguel 2015) may sharply widen cross-firm differences in marginal products. Likewise, richer regions tend to feature more specialized production and greater firm-size dispersion (Chen 2022; Poschke 2018), which can amplify MRPK divergence when some firms adjust quickly and others cannot.

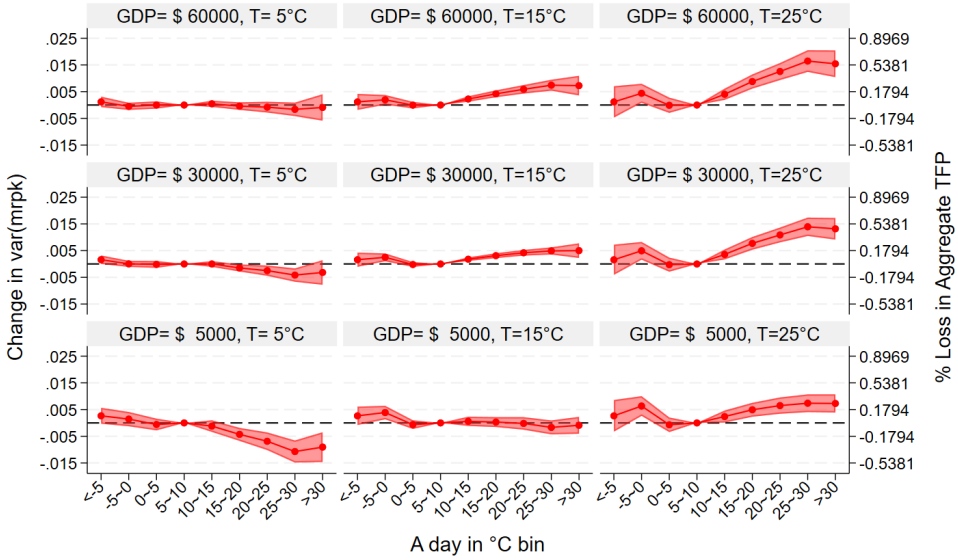
Following Carleton et al. (2022) and Nath (2025), we interact time-invariant measures of long-

run average temperature and GDP per capita with each temperature bin:

$$\begin{aligned}\sigma^2_{mrpk_{s,r,t}} = &\sum_{b\in B_{/(5\sim10^\circ C)}} \lambda^b \times \text{Tbin}^b_{r,t} + \sum_{b\in B_{/(5\sim10^\circ C)}} \lambda^b_{\overline{T}} \times \text{Tbin}^b_{r,t} \times \overline{T}_r \\ &+ \sum_{b\in B_{/(5\sim10^\circ C)}} \lambda^b_{\text{GDP}_{\text{pc}}} \times \text{Tbin}^b_{r,t} \times \ln \overline{\text{GDP}_{\text{pc}}}_r + \boldsymbol{\delta}\tilde{\mathbf{X}}_{s,r,t} + \alpha_{c,t} + \eta_{s,r} + \varepsilon_{s,r,t}.\end{aligned} \tag{12}$$

Here $\overline{T}_r$ and $\ln \overline{\text{GDP}_{pc,r}}$ are region-level averages over 1997–2018.[14] The interaction coefficients show how the effect of an additional day in bin b varies with long-run climate and development.

Figure 4: MRPK Dispersion and TFP Loss Across Climates and Income



Notes: The graphs plot the predicted effect of exposure to daily mean temperature bins on MRPK dispersion and TFP loss at varying levels of income and climates. These predicted effects are derived from the interacted panel regression specified in equation (12). The graphs include 90% confidence intervals, and standard errors are clustered at the regional level. The left y-axis indicates changes in MRPK dispersion, and the right y-axis shows the calculated TFP loss. The reference temperature is at 5∼10°C.

Figure 4 presents the implied marginal effects from equation (12), evaluated at different combinations of long-run climate and income. Two facts stand out. First, extreme heat raises misallocation more in warmer regions: holding income fixed, the effect of a day above 30°C is close to zero or negative in cold climates but positive and increasing in hot climates. Evaluated at baseline incomes near $60,000, an additional day above 30°C implies a TFP loss of 0.26% in Arizona but a TFP gain of 0.16% in Norway. Second, extreme heat raises misallocation more in richer regions: holding long-run climate fixed, the response becomes more U-shaped as GDP per capita rises. Evaluated at long-run temperatures near 10°C, an additional day above 30°C implies a TFP loss of 0.11% in the Netherlands but only 0.03% in Türkiye.

These heterogeneous effects distinguish the misallocation channel sharply from other climate channels, such as labor productivity (Nath 2025) and mortality risk (Carleton et al. 2022), where hotter and richer places often appear less vulnerable because of long-run adaptation. Here the

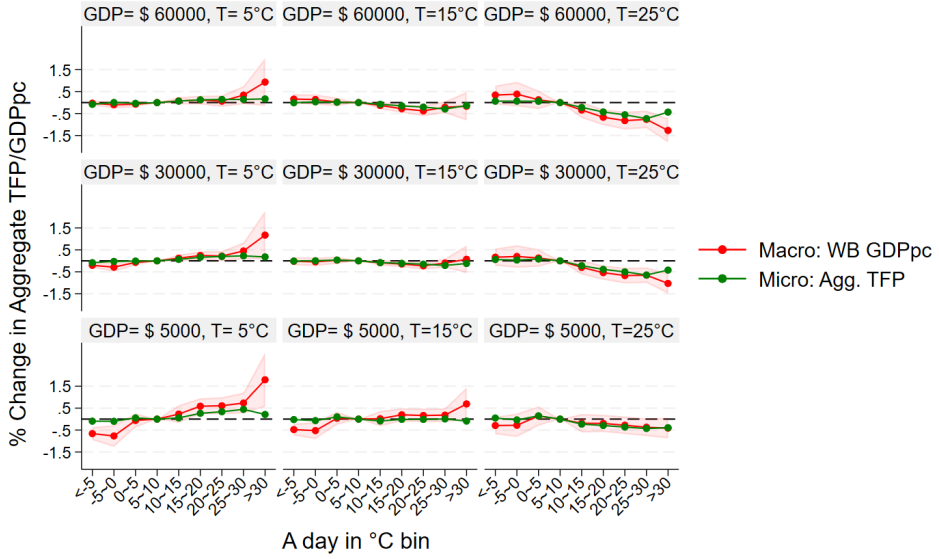
14. The sub-national level (PPP-adjusted) GDP per capita data is from DOSE. It is important to use sub-national level data as countries like India and China admit large income heterogeneity across the districts or prefectures. The GDP per capita data is in 2017 International Dollars.

opposite pattern emerges: hotter and richer regions suffer larger TFP losses through misallocation. As economies develop, production and adaptation become more heterogeneous across firms, and achieving efficient allocations may itself become more difficult under climate stress.

4.5 Micro estimates vs. macro estimates

To show that incorporating the misallocation channel quantitatively bridges the micro–macro disconnect, we compare *micro estimates*—derived from firm-level data and aggregation theory—with *macro estimates* obtained from country-level aggregate data. For the micro estimates, we estimate the technology channel using the PPML specification in equation (11), augmented with the interaction terms from equation (12), and combine the estimated technology effects with our misallocation estimates. If the misallocation channel is shut down, the resulting micro estimates are essentially flat, indicating that cross-country heterogeneity in the total effect is driven primarily by misallocation. For the macro estimates, we use annual GDP per capita data from the World Bank for 1990–2019 and estimate a country-level regression analogous to equation (12), including country and year fixed effects.

Figure 5: Aggregate Impact of Climate Shocks: Micro vs. Macro Estimates



Notes: The graphs include 90% confidence intervals, and standard errors are clustered at the region level. The reference temperature is at 5∼10°C.

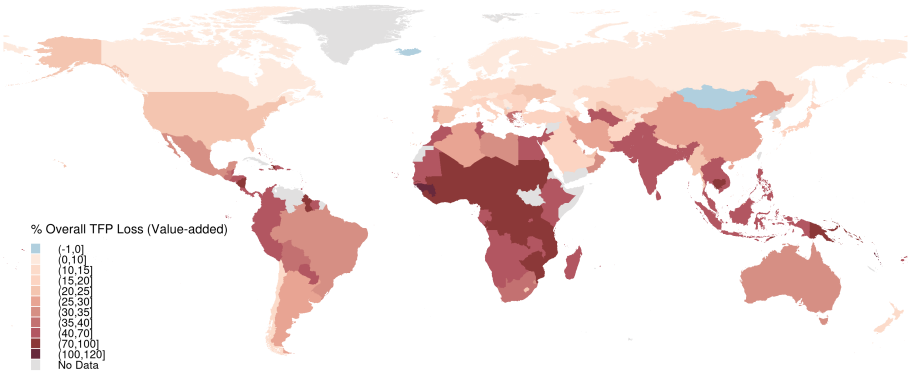
Figure 5 plots the two sets of estimates. The micro-based estimates closely track the macro-based estimates, with two exceptions: they understate damages from extreme heat in regions with high income and high mean temperatures, and they overstate the benefits of heat shocks in colder, poorer regions. Both types of regions are rare in our sample. These results suggest that our micro estimates capture the same cross-country heterogeneity in temperature sensitivity as the macro estimates, making them informative for out-of-sample extrapolation.

4.6 End-of-the-century projection

We use the estimates from Section 4.4 to project how climate-induced misallocation will affect aggregate TFP by the end of the 21st century (2081–2100) for 4,881 regions in 172 countries. We first construct the implied loss for each region by applying the heterogeneous reduced-form estimates to projected changes in temperature-bin exposure, mean temperature, and income relative to the baseline. For the end-of-century map, we then aggregate these regional losses to the country level using projected regional GDP shares. For the global time series, we repeat the same regional calculation period by period and aggregate regions using the average of their current-period and previous-period shares of world GDP; because the projected income data are available every five years, this smoothing is defined over adjacent five-year periods. The formal projection equations are presented in Appendix E.2.[15]

Figure 6 maps the cross-country distribution of projected TFP losses. The losses are systematically larger in already hot countries, where warming shifts more days into the upper temperature bins and where the estimated marginal effect of heat is amplified by higher temperature and income levels. The upper tail of the distribution is concentrated in tropical and subtropical economies, with projected losses of 72.1% in Tanzania, 69.4% in Honduras, 55.0% in Malaysia, and 50.0% in India. Losses remain economically large outside this upper tail, reaching 27.0% in Argentina and 20.8% in the United States, while they are smaller in maritime climates such as the United Kingdom (7.7%).

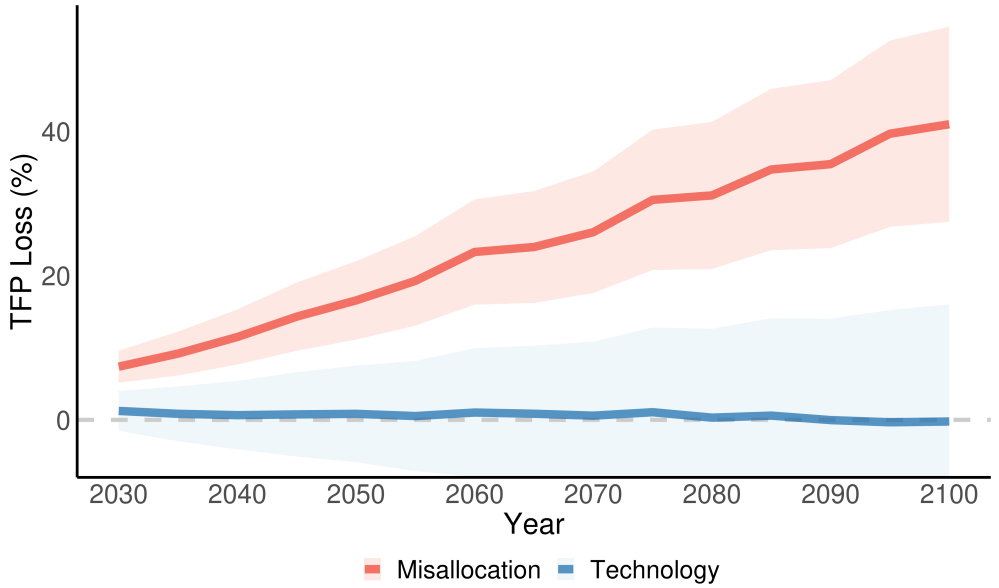
Figure 6: End-of-century Projected TFP Loss Due to the Misallocation Channel



Notes: Figure shows the projected TFP losses from capital misallocation under SSP3-4.5 scenarios. The estimation follows equation (43) in Appendix E.2.

15. For our 32-country sample, regions are defined as in Section 3. For other countries, we use GADM1-level administrative units from the GADM dataset.

Figure 7: Projected Global TFP Loss: Misallocation vs. Technology Channels, 2030–2100



Notes: Figure plots the total TFP loss from the misallocation and technology channels across years, with 90% confidence intervals from parameter uncertainty. All computations are under the SSP3-4.5 scenario and percentage losses are measured relative to the baseline level. The decomposition of misallocation losses into shock, level, and income components is reported in Appendix E.2; see Figure E.17.

After aggregating regional losses with projected GDP weights, the end-of-the-century cross-sectional projection implies a 36.41% decline in aggregate TFP relative to baseline.[16] Appendix E.2 decomposes this total into shock, level, income, and residual effects, while Figure 7 plots the global time-series for both the misallocation and technology channels via a Törnqvist index. The technology channel remains close to zero throughout the century, so the projected increase in global TFP losses is driven primarily by the misallocation channel.

These estimates are substantial and comparable to the overall projected impacts reported by Burke, Hsiang, and Miguel (2015) and, more recently, Bilal and Känzig (2024). Much of the total projected effect arises from the sharp increase in the marginal effect of daily temperature shocks, driven by rising temperatures and incomes. As with any century-scale forecast, our projection is subject to caveats—most notably out-of-sample extrapolation, since end-of-century income and temperature levels lie well beyond historical experience. Overall, these results underscore the quantitative importance of the misallocation channel, which may serve as one of the key forces driving up the social cost of carbon.

5 A Firm Dynamics Model of Climate and MRPK Dispersion

The reduced-form evidence establishes that both long-run climatic conditions and short-run temperature fluctuations raise capital misallocation, with stronger effects in hotter and more developed regions. Why does temperature disperse MRPK, and why does adaptation amplify rather

16. Under this weighting scheme, the resulting TFP losses correspond to a model-implied global TFP loss based on a Cobb-Douglas aggregator of country-level value added. This also serves as a first-order approximation under any constant-returns-to-scale aggregator for global output.

than attenuate this effect? To answer these questions, we build a tractable general-equilibrium model with risk-averse households and heterogeneous firms, featuring time-to-build capital, pre-existing capital wedges, and costly endogenous climate adaptation. A novel feature is that adaptation is fully endogenous: firms choose their effective productivity process by selecting adaptation intensity in response to expected climate conditions and weather risk, building on Kopytov, Taschereau-Dumouchel, and Xu (2024).

The model's demand and production structure mirrors the structure of our accounting framework and is defined at the region-year level. For expositional clarity, we suppress the region-year notation here and introduce regional heterogeneity later.

5.1 Model Structure

5.1.1 Household

In each region, a representative household with King-Plosser-Rebelo preferences consumes C_t, inelastically supplies a mass N of unskilled labor to production, and elastically supplies adaptation resources R_t to firms.[17] The household chooses $\{C_t, R_t\}_{t\geq 0}$ to maximize discounted expected lifetime utility

$$\max_{\{C_t,R_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \frac{(C_t \exp(-\eta R_t))^{1-\gamma}}{1-\gamma},$$

where r is the discount rate and $\gamma \geq 1$ governs relative risk aversion. We interpret R_t as adaptation services such as cooling, insulation, and other resilience inputs purchased by firms. $\eta > 0$ governs the consumption-equivalent unit cost of supplying effective adaptation resources: one more unit of R_t reduces effective log consumption by η. Lower η means a region can provide adaptation more cheaply, so firms face a lower equilibrium adaptation price and choose more adaptation.

Differentiated goods aggregate into the final-good bundle

$$Y_t = \left(\int_0^1 B_{it}^{\frac{1}{\sigma}} Y_{it}^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}},$$

where $\sigma > 1$ is the elasticity of substitution across varieties i and B_{it} is a preference shifter. The same final good is used for consumption and capital accumulation, so the aggregate resource constraint is

$$C_t + K_{t+1} - (1-\delta)K_t = Y_t.$$

We assume markets are complete with respect to aggregate risk, so firms discount future cash flows using the household's stochastic discount factor (SDF)[18]

$$\Lambda_{t+1} = \frac{1}{1+r}\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \exp\left[(\gamma-1)\eta(R_{t+1}-R_t)\right].$$

17. Equivalently, we can think of the household supplying skilled labor $\mathcal{R}_t$ to an adaptation sector with technology $\mathcal{R}_t = \eta R_t$. We suppress $\mathcal{R}_t$ in the main text and work directly with effective adaptation resources R_t.

18. Writing period utility as $u(C_t, R_t) = \frac{(C_t e^{-\eta R_t})^{1-\gamma}}{1-\gamma}$, the marginal utility of consumption is $u_C(C_t, R_t) = C_t^{-\gamma} \exp(-(1-\gamma)\eta R_t)$. Hence $\Lambda_{t+1} = \frac{1}{1+r}\frac{u_C(C_{t+1},R_{t+1})}{u_C(C_t,R_t)}$.

The household's optimality condition is

$$w_{t+1}^R = \eta \frac{\mathbb{E}_t\left[\Lambda_{t+1} C_{t+1}\right]}{\mathbb{E}_t\left[\Lambda_{t+1}\right]},$$

where the equilibrium price of adaptation resources must equal the household's marginal consumption-equivalent cost of supplying one more unit of R_{t+1}. [19]

5.1.2 Firms

Each firm i produces a differentiated variety with Cobb-Douglas technology, $Y_{it} = \tilde{A}_{it} K_{it}^{\tilde{\alpha}_K} N_{it}^{\tilde{\alpha}_N}$, where $\tilde{A}_{it}$ is physical productivity, K_{it} is capital, N_{it} is labor, and $\tilde{\alpha}_K + \tilde{\alpha}_N = 1$. Firms compete monopolistically à la Dixit-Stiglitz. Combining the production function with the demand system implies an equilibrium revenue function of the form

$$P_{it} Y_{it} = Y_t^{\frac{1}{\sigma}} B_{it}^{\frac{1}{\sigma}} \tilde{A}_{it}^{\frac{\sigma-1}{\sigma}} K_{it}^{\alpha_K} N_{it}^{\alpha_N},$$

where $\alpha_F = \left(1 - \frac{1}{\sigma}\right) \tilde{\alpha}_F$ for $F \in \{K, N\}$, and $B_{it}^{\frac{1}{\sigma}} \tilde{A}_{it}^{\frac{\sigma-1}{\sigma}}$ captures the firm's demand-adjusted productivity. Since the market structure is identical to Section 2, an aggregate production function $Y_t = A_t K_t^{\tilde{\alpha}_K} N^{\tilde{\alpha}_N}$ exists for aggregate productivity A_t and capital $K_t = \int_i K_{it}$.

Flexible Inputs and Profits At the beginning of period t, after capital and adaptation decisions are made and all shocks are realized, firms choose labor from a competitive market to maximize profits: $\Pi_{it} = \max_{N_{it}} Y_t^{\frac{1}{\sigma}} B_{it}^{\frac{1}{\sigma}} \tilde{A}_{it}^{\frac{\sigma-1}{\sigma}} K_{it}^{\alpha_K} N_{it}^{\alpha_N} - W_t N_{it}$. Writing the competitive wage W_t as a function of aggregate output, $W_t = \tilde{\alpha}_N \frac{Y_t}{N}$, we can solve for the optimal labor choice and obtain the profit function:

$$\Pi_{it} = \mathcal{Q}_t A_{it} K_{it}^{\alpha}, \tag{13}$$

where $\alpha = \frac{\alpha_K}{1-\alpha_N}$. $\mathcal{Q}_t := G(A_t K_t)^{\theta_A \tilde{\alpha}_K}$ is a common aggregate profitability shifter.[20] In the empirically relevant case, $\theta_A < 0$, reflecting that higher aggregate productivity raises wages enough to reduce firm profits in general equilibrium.

5.1.3 Productivity Process and Endogenous Adaptation

Effective productivity A_{it} combines a firm-specific baseline with a temperature-exposure term that the firm can partially offset through adaptation:

$$\log A_{it} \equiv a_{it} = \left(\beta_i + \xi_{it} + \widetilde{\beta}_{it}\right)\left(T_t - T^*\right) - \frac{\Gamma_\beta}{2} \widetilde{\beta}_{it}^2 + \underbrace{a^0 + \bar{a}_i + z_{it}}_{\text{baseline productivity}}, \qquad \Gamma_\beta > 0,$$

where T^* is the bliss-point temperature emphasized in the climate-econometrics literature.[21]

19. The full budget constraint and derivation are in Appendix F.

20. Here $\theta_A = \frac{1-\sigma\alpha_N}{\sigma(1-\alpha_N)}$ and $G = (1-\alpha_N) N^{\frac{\alpha_N}{1-\alpha_N} + \tilde{\alpha}_N \theta_A}$.

21. T^* typically lies around 10–13°C; Burke, Hsiang, and Miguel (2015) estimate $T^* \approx 12$°C.

Temperature sensitivity and endogenous adaptation. Temperature sensitivity decomposes as $\beta_{it} = \beta_i + \xi_{it} + \widetilde{\beta}_{it}$. The persistent component $\beta_i \sim \mathcal{N}(\bar{\beta}, \sigma_\beta^2)$ captures product-level exposure known to the firm, while the transitory component $\xi_{it} \sim \mathcal{N}(0, \sigma_\xi^2)$ captures unpredictable firm-specific climate disruptions. Because temperature enters through the signed deviation $T_t - T^*$, the same specification captures both heat and cold exposure: $T_t > T^*$ represents heat relative to the bliss point, while $T_t < T^*$ represents cold. The effective loading $\beta_i + \xi_{it} + \widetilde{\beta}_{it}$ determines how this deviation maps into productivity. Thus heat is harmful when the effective loading is negative, whereas cold is harmful when the effective loading is positive.

Firms choose the exposure adjustment $\widetilde{\beta}_{it}$ to reshape their productivity response to temperature deviations. Adaptation affects productivity through $\widetilde{\beta}_{it}(T_t - T^*) - \frac{\Gamma_\beta}{2}\widetilde{\beta}_{it}^2$. The linear term captures the benefit of adjusting exposure to the locally relevant climate risk, while the quadratic term captures diminishing productivity returns from further adjustment. Since this term enters log productivity a_{it}, its profit cost scales with the firm's operating scale. Firms also pay a convex resource cost to implement the exposure adjustment. Choosing $\widetilde{\beta}_{i,t+1}$ also requires

$$R_{i,t+1} = \frac{1}{2} H \widetilde{\beta}_{i,t+1}^2$$

units of adaptation resources, purchased at the equilibrium price w_{t+1}^R. This resource cost is quasi-fixed in firm scale, so larger firms have stronger incentives to adapt: the benefits scale with their profit base, while the direct resource cost does not. Together, the productivity curvature and the convex resource cost deliver an interior, bounded adaptation choice.

Baseline productivity. The baseline consists of a constant level a^0, a permanent component $\bar{a}_i \sim \mathcal{N}(0, \sigma_{\bar{a}}^2)$, and a transitory idiosyncratic component z_{it} that follows an AR(1) process:

$$z_{i,t+1} = \rho_z z_{it} + \varepsilon_{i,t+1}^z, \qquad \varepsilon_{i,t+1}^z \sim \mathcal{N}(0, \sigma_{\varepsilon^z}^2).$$

5.1.4 Temperature Dynamics and Forecast Uncertainty

We assume local temperature follows a stationary AR(p) process:

$$T_t - \bar{T} = \sum_{\tau=1}^{p} \rho_\tau \left(T_{t-\tau} - \bar{T}\right) + \varepsilon_t^T.$$

At the end of period t, agents observe an unbiased one-step-ahead forecast $\mathbb{E}_t[T_{t+1}]$ and perceive forecast uncertainty $\mathbb{VAR}_t(T_{t+1})$, which we model as a smoothed average of past q-period squared forecast errors:

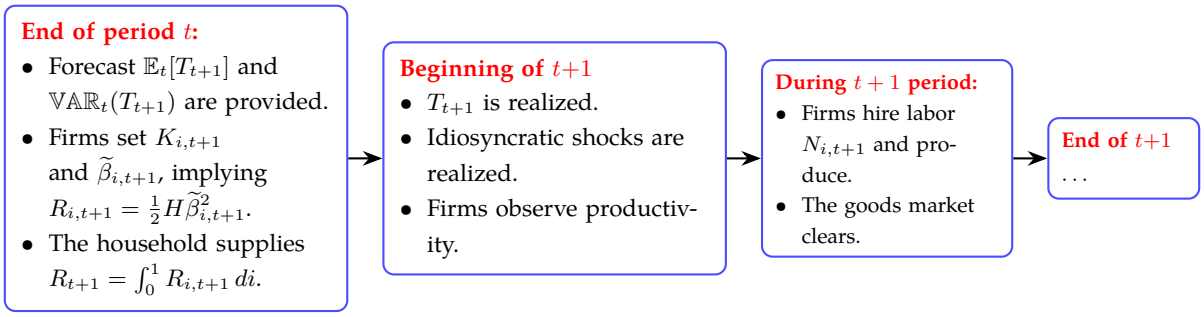
$$\mathbb{VAR}_t(T_{t+1}) = \frac{1}{q} \sum_{\tau=1}^{q} \left(T_{t+1-\tau} - \mathbb{E}_{t-\tau}[T_{t+1-\tau}]\right)^2.$$

Thus, perceived uncertainty rises following periods of large forecast errors.

5.1.5 Timing and Dynamic Investment

At the end of period t, agents observe the forecast $\mathbb{E}_t[T_{t+1}]$ with perceived forecast uncertainty $\mathbb{VAR}_t(T_{t+1})$. Given this information, firms choose next-period capital $K_{i,t+1}$ and next-period adaptation intensity $\widetilde{\beta}_{i,t+1}$. This choice implies firm-level adaptation-resource demand $R_{i,t+1} = \frac{1}{2}H\widetilde{\beta}^2_{i,t+1}$, while the representative household supplies aggregate adaptation resources $R_{t+1} = \int_0^1 R_{i,t+1}\, di$. At the beginning of period $t+1$, temperature T_{t+1} and all idiosyncratic shocks are realized. Firms then observe realized productivity, choose flexible labor input $N_{i,t+1}$ to produce, and the goods market clears. Figure 8 illustrates the full timing of decisions in the model.

Figure 8: Timing of events in the model.



The firm's dynamic problem can be written as the Bellman equation

$$
\begin{aligned}
V(\mathbf{x}_{it}, \mathbf{x}_t, \mathbf{d}_t, \Gamma_{\text{micro},t}) = \max_{K_{i,t+1},\, \widetilde{\beta}_{i,t+1}} \Bigg\{ & \mathcal{Q}_t e^{a_{it}} K_{it}^{\alpha} - \left(\frac{1}{2}H\widetilde{\beta}^2_{i,t+1}\right) w^R_{t+1} - K_{i,t+1} \\
& - \frac{r+\delta}{1+r} K_{i,t+1}\big(\mathcal{T}_{i,t+1} - 1\big) + (1-\delta)K_{it} \\
& + \mathbb{E}_t[\Lambda_{t+1}\, V(\mathbf{x}_{i,t+1}, \mathbf{x}_{t+1}, \mathbf{d}_{t+1}, \Gamma_{\text{micro},t+1})] \Bigg\}.
\end{aligned} \tag{14}
$$

where

$$
\mathbf{x}_t \equiv \big(\mathbb{E}_t[T_{t+1}],\, \mathbb{VAR}_t(T_{t+1})\big), \qquad \mathbf{x}_{it} \equiv \big(K_{it},\, z_{it},\, \bar{a}_i,\, \beta_i,\, \mathcal{T}_{i,t+1}\big),
$$

$$
\mathbf{d}_t \equiv \big(A_t,\, K_t,\, R_t\big).
$$

Here $\mathbf{x}_t$, $\mathbf{x}_{it}$, and $\mathbf{d}_t$ collect the aggregate temperature moments, firm-level states, and aggregate endogenous variables, respectively, while $\Gamma_{\text{micro},t}$ denotes the cross-sectional distribution of firm states. Current productivity a_{it} enters the flow payoff directly, and $K_{i,t+1}$ and $\widetilde{\beta}_{i,t+1}$ are chosen before T_{t+1}, $\xi_{i,t+1}$, and $\varepsilon^z_{i,t+1}$ are realized.

Frictions in the Capital Market Firms face heterogeneous capital market distortions summarized by the (log) wedge $\tau_{i,t+1} \equiv \log \mathcal{T}_{i,t+1}$:

$$
\tau_{i,t+1} = \omega_{\bar{a}}\bar{a}_i + \phi^{\tau}_{i,t+1}, \qquad \omega_{\bar{a}} > 0.
$$

The first component captures *size-dependent* distortions: when $\omega_{\bar{a}} > 0$, firms born with higher productivity $\bar{a}_i$ (and hence larger scale) face systematically larger wedges, in the spirit of Restuccia

and Rogerson (2008). The second component $\phi^{\tau}_{i,t+1}$ captures idiosyncratic distortions and follows an AR(1) process,

$$\phi^{\tau}_{i,t+1} = \rho^{\tau}_{\phi}\,\phi^{\tau}_{i,t} + \varepsilon^{\phi^{\tau}}_{i,t+1}, \qquad \varepsilon^{\phi^{\tau}}_{i,t+1} \sim \mathcal{N}(0, \sigma^2_{\varepsilon^{\phi^{\tau}}}).$$

A firm with wedge $\mathcal{T}_{i,t+1}$ pays an additional $\mathcal{T}_{i,t+1} - 1$ on the risk-free user cost for each unit of installed capital.

Firm's Optimality Conditions The firm's optimal investment policies follow from the Bellman problem in (14). Optimal capital accumulation satisfies the Euler equation

$$\frac{1}{1+r}(r+\delta)\mathcal{T}_{i,t+1} = \alpha K^{\alpha-1}_{i,t+1}\,\mathbb{E}_t[\Lambda_{t+1}\mathcal{Q}_{t+1}e^{a_{i,t+1}}] + (1-\delta)\,\mathbb{E}_t\left[\Lambda_{t+1} - \frac{1}{1+r}\right].$$

This condition equates the wedge-adjusted risk-free user cost of installed capital to the expected SDF-discounted marginal profit product of capital, plus a risk-valuation adjustment of future resale value. Optimal adaptation satisfies

$$\underbrace{Hw^{R}_{t+1}\widetilde{\beta}_{i,t+1}}_{\text{marginal resource cost}} = K^{\alpha}_{i,t+1}\,\mathbb{E}_t\left[\Lambda_{t+1}\mathcal{Q}_{t+1}e^{a_{i,t+1}}\left(\underbrace{T_{t+1}-T^*}_{\text{marginal exposure benefit}} - \underbrace{\Gamma_{\beta}\widetilde{\beta}_{i,t+1}}_{\text{marginal productivity cost}}\right)\right].$$

The optimal adaptation condition equates the marginal resource cost on the left-hand side to the expected SDF-discounted marginal profit benefit on the right-hand side. Because that marginal benefit is proportional to operating profits, larger firms—whose cash flows are more exposed to temperature fluctuations—choose more adaptation. At the same time, the quadratic productivity penalty and convex resource cost keep the adaptation response bounded. Heterogeneity in scale and fundamentals therefore translates directly into heterogeneity in endogenous adaptation choices.

The Homogeneous Deterministic Steady State To build intuition, consider a homogeneous deterministic steady state (HDSS) with no firm heterogeneity and no uncertainty. Steady-state adaptation satisfies

$$\widetilde{\beta} = \frac{\bar{T} - T^*}{\Gamma_{\beta} + (1+r)\eta H\left[\frac{1}{1-\alpha_N} - \frac{\alpha\delta}{(r+\delta)}\right]}.$$

The average net temperature sensitivity of productivity in the HDSS is $\bar{\beta} + \widetilde{\beta}$: $\bar{\beta}$ is the average baseline temperature sensitivity, while $\widetilde{\beta}$ is the endogenous adaptation component. Thus $\bar{\beta}+\widetilde{\beta} < 0$ means that heat remains productivity-reducing on average even after adaptation. Steady-state adaptation is positive in warm regions ($\bar{T} > T^*$) and negative in cold regions ($\bar{T} < T^*$). Its magnitude is bounded by the quadratic curvature in adaptation.[22]

22. The denominator is strictly positive: $\Gamma_{\beta} > 0$, and the bracketed term equals $C/\Pi > 0$ in the HDSS. The full derivation is provided in Appendix F.

5.2 Solving the Model

We solve the model in closed form using a risk-adjusted perturbation around the HDSS, following Coeurdacier, Rey, and Winant (2011) and David and Venkateswaran (2019). The approach proceeds in three steps: (i) we develop a second-order expansion of stochastic variables in conditional moments to approximate the risk-adjusted FOCs; (ii) we approximate firms' policy rules using an affine–bilinear expansion in firm-level and aggregate state vectors; and (iii) we substitute the approximated policy rules and aggregation equations into all equilibrium conditions and solve for the unknown coefficients by matching terms of equal order. This guess-and-verify procedure yields tractable expressions for both micro policies and macro equilibrium objects, including MRPK dispersion and the resulting aggregate TFP losses. We provide the full derivation in Appendix F.

5.3 Investment and Adaptation in the Cross-section

The aggregate policy functions and state dynamics are qualitatively similar to those in a standard heterogeneous-firm real business cycle model. We therefore emphasize the model's novel cross-sectional implications: how heterogeneity in productivity, wedges, and climate exposure shapes firms' capital and adaptation choices, and how these choices generate MRPK dispersion. The analysis focuses on the empirically relevant warm-region case, in which heat remains productivity-reducing on average despite endogenous adaptation.

Definition 1 (A typical warm region.) *A region is a typical warm region if average temperature exceeds the bliss point, $\bar{T} > T^*$, and heat remains productivity-reducing on average after adaptation, $\bar{\beta} + \widetilde{\beta} < 0$.*

Proposition 2 shows that cross-sectional differences in both adaptation and capital investment are driven by expected profitability and the covariance between firm productivity and temperature risk.[23]

Proposition 2 (Investment and Adaptation) *Let $\Delta^i x_{it} \equiv x_{it} - \bar{x}_t$ denote firm i's deviation from the time-t cross-sectional mean. The expected profit is increasing in expected productivity but declining in wedges,*

$$\Delta^i \mathbb{E}_t \hat{\pi}_{it+1} \;=\; \Delta^i \mathbb{E}_t \hat{a}_{it+1} - \alpha\, \Delta^i \hat{\tau}_{it+1}.$$

In equilibrium, optimal policy decisions satisfy:

Adaptation.

$$\Delta^i \widehat{\widetilde{\beta}}_{i,t+1} = \frac{\widetilde{\beta} S_\beta}{\bar{\Delta}(1-\alpha)} \left[\Delta^i \mathbb{E}_t \hat{\pi}_{i,t+1} + \left(\Theta + \frac{1-\alpha}{S_\beta} \right) \Delta^i \mathbb{COV}_t \left(\hat{a}_{i,t+1}, \hat{T}_{t+1} \right) \right], \tag{15}$$

Capital.

$$\Delta^i \hat{k}_{i,t+1} = -\Delta^i \hat{\tau}_{i,t+1} + \frac{1}{1-\alpha} \Delta^i \mathbb{E}_t \hat{\pi}_{i,t+1} + \frac{\Theta}{1-\alpha} \Delta^i \mathbb{COV}_t \left(\hat{a}_{i,t+1}, \hat{T}_{t+1} \right), \tag{16}$$

23. Throughout this subsection, hats denote first-order deviations from the homogeneous deterministic steady state. For any positive variable x, $\hat{x}_{it} := \log x_{it} - \log x$, where x denotes the corresponding HDSS value. For variables already expressed in logs, $\hat{a}_{it} := a_{it} - a$ and $\hat{\tau}_{it} := \tau_{it} - \tau = \tau_{it}$. For temperature, $\hat{T}_t := T_t - \bar{T}$. Endogenous adaptation is written as $\widehat{\widetilde{\beta}}_{it} := \widetilde{\beta}_{it} - \widetilde{\beta}$, and persistent temperature sensitivity as $\beta_i = \bar{\beta} + \hat{\beta}_i$.

where $S_\beta := (\bar{T} - T^*) - \Gamma_\beta \widetilde{\beta}$ *is the local marginal productivity benefit of adaptation around steady state and* $\Theta = \left(\frac{\sigma}{\sigma - 1} - \widetilde{\alpha}_N\right)(\bar{\beta} + \widetilde{\beta}) - \gamma \frac{d\hat{c}_{t+1}}{d\hat{T}_{t+1}}$ *summarizes the effect of temperature on direct payoff and consumption-risk.*

Notice that in a typical warm region, we will have that $\widetilde{\beta} > 0$, $\bar{\beta} + \widetilde{\beta} < 0$ and $S_\beta > 0$. Therefore, firms with higher expected productivity and lower wedges anticipate a larger profit base, and therefore adapt and invest more. The adaptation coefficient is scaled by $\widetilde{\beta} S_\beta / [(\bar{T} - T^*)(1 - \alpha)]$, reflecting that the marginal productivity gain from adaptation is damped by quadratic curvature. Moreover, hot states are bad aggregate states. For large enough risk aversion γ, $\Theta > 0$. In that case, the state-price value of payoffs in hot states dominates the average direct productivity loss from heat. A firm whose productivity is high in hot states becoming less risky and therefore provides a hedge, faces a lower risk-adjusted required return, and chooses more capital.

Combining part (i) and (ii) in Proposition 2 yields the expression for how a firm tilts its choices toward adaptation relative to capital:

Corollary 1 (Within-firm Adaptation-capital Trade-off)

$$\frac{\Delta^i \widehat{\widetilde{\beta}}_{it+1}}{\widetilde{\beta} S_\beta / (\bar{T} - T^*)} - \Delta^i \hat{k}_{it+1} = \underbrace{\Delta^i \widehat{\tau}_{it+1}}_{\text{relative price}} + \underbrace{\frac{1}{S_\beta} \Delta^i \mathbb{COV}_t \left(\hat{a}_{it+1}, \widehat{T}_{t+1}\right)}_{\text{relative climate exposure payoff}}.$$

Corollary 1 isolates the within-firm input tilt between adaptation and capital. The expected-profit component raises both inputs and therefore cancels from the relative expression. What remains are two forces. First, a higher capital wedge raises the relative price of installed capital. This does not necessarily imply that adaptation rises unconditionally, since a higher wedge also lowers the firm's scale. Rather, the corollary says that, conditional on the firm's capital choice and temperature-state covariance, a higher wedge tilts the firm toward adaptation relative to capital. Second, firms whose productivity is more positively state-contingent with temperature tilt toward adaptation, because adaptation directly changes the firm's temperature loading. The strength of this relative state-contingency channel is governed by the local marginal productivity benefit of adaptation, $S_\beta = (\bar{T} - T^*)\Gamma_\beta \widetilde{\beta}$.

Proposition 2 and Corollary 1 characterize equilibrium forces but still involve endogenous objects. We now close this fixed point and express firm choices as functions of firm primitives and temperature forecasts. The key object is the endogenous component of a firm's effective temperature exposure.

Definition 2 (Endogenous effective exposure) $\Psi^{\Gamma,C}_{i,t+1}$ *collects the endogenous component of a firm's effective temperature exposure.*

$$\Psi_{i,t+1} = \Upsilon \left[(\bar{T} - T^*)\hat{\beta}_i + \bar{a}_i + \rho_z z_{it} - \alpha \tau_{i,t+1} \right], \tag{17}$$

where $\Upsilon = \frac{\widetilde{\beta} S_\beta}{(\bar{T} - T^*)(1 - \alpha) - \widetilde{\beta} S_\beta^2}$ *is the endogenous sensitivity multiplier.*

The index increases with the firm's persistent and predictable productivity components ($\bar{a}_i$ and $\rho_z z_{it}$), intrinsic temperature exposure $(\bar{T} - T^*)\hat{\beta}_i$, and decreases with the effective user cost

of capital $\tau_{i,t+1}$. Economically, $\Psi_{i,t+1}$ summarizes the endogenous component of effective exposure that arises from profitability and adaptation incentives after accounting for curvature in the adaptation margin.

We now solve for the firm's policy functions.

Proposition 3 (Policy functions) *The firm's policy functions satisfy:* ***(i) Adaptation.***

$$\Delta^i \widetilde{\widehat{\beta}}_{i,t+1} = \Psi_{i,t+1} + \Upsilon \left(\hat{\beta}_i + \Psi_{i,t+1}\right) \left[\mathbb{E}_t \hat{T}_{t+1} + \left(\Theta + \frac{1-\alpha}{S_\beta}\right) \mathbb{VAR}_t(\hat{T}_{t+1})\right]. \tag{18}$$

(ii) Capital.

$$\Delta^i \hat{k}_{i,t+1} = -\hat{\tau}_{i,t+1} + \frac{(\bar{T} - T^*)}{\widetilde{\beta} S_\beta} \Psi_{i,t+1} + \frac{\hat{\beta}_i + \Psi_{i,t+1}}{1 - \alpha - \widetilde{\beta} S_\beta^2/(\bar{T} - T^*)} \left[\mathbb{E}_t \hat{T}_{t+1} + \left(\Theta + \frac{\widetilde{\beta} S_\beta}{\bar{T} - T^*}\right) \mathbb{VAR}_t(\hat{T}_{t+1})\right]. \tag{19}$$

Proposition 3 gives the closed-form local policy rules after solving the firm-level fixed point. Both policies depend on the firm's *effective exposure* to temperature news, $\hat{\beta}_i + \Psi_{i,t+1}$. The effective exposure combines intrinsic product sensitivity with endogenous sensitivity index, and it governs how strongly a firm responds to expected and unexpected weather shocks. Moreover, the expected scale is mapped to adaptation decisions with the endogenous sensitivity multiplier,

$$\Upsilon = \frac{\widetilde{\beta} S_\beta}{\bar{\Delta}(1-\alpha) - \widetilde{\beta} S_\beta^2} = \frac{\widetilde{\beta} S_\beta}{\bar{\Delta}(1-\alpha)} \sum_{r=0}^{\infty} \left(\underbrace{\frac{\widetilde{\beta} S_\beta}{\bar{\Delta}(1-\alpha)}}_{\text{scale to adaptation}} \times \underbrace{S_\beta}_{\substack{\text{adaptation to} \\ \text{productivity}}} \right)^r.$$

This multiplier summarizes the local fixed-point feedback between scale and adaptation. A higher expected profit base raises the marginal value of adaptation. Additional adaptation changes productivity with local slope $S_\beta = \bar{\Delta} - \Gamma_\beta \widetilde{\beta}$. Higher productivity then raises the value of operating at scale, which further raises the value of adaptation by $\frac{\widetilde{\beta} S_\beta}{\bar{\Delta}(1-\alpha)}$.

With perfectly predictable temperatures, policy rules depend only the expected temperature. Higher expected heat raises both adaptation and capital more for firms with greater effective exposure. In a typical warm region, expected warming therefore disproportionately increases investment and adaptation among firms that are already expected to operate at larger scale (high Ψ_i) or whose products are more temperature exposed (high $\hat{\beta}_i$). Because these firms adjust more strongly, expected heat increases within-region-year dispersion in both adaptation and capital. If risk aversion is high so $\Theta > 0$, when temperatures are harder to predict, the forecast uncertainty further widens firms' differences by inducing larger responses among firms with higher effective exposure.

Interestingly, the presence of endogenous adaptation leads to "self-reinforcing"endogenous productivity and risk premia: by raising effective productivity and improving performance in hot states, it lowers risk premia and required returns, which further strengthens incentives to invest and adapt, amplifying cross-sectional differences.

5.3.1 What Determines MRPK?

Using the profit function in equation (13) and the constant markups implied by the Dixit-Stiglitz competition, log MRPK can be written (up to constants) as

$$\widehat{mrpk}_{i,t+1} = \hat{a}_{i,t+1} - (1-\alpha)\hat{k}_{i,t+1} + \theta_A \hat{y}_{t+1}. \tag{20}$$

Realized MRPK reflects ex post productivity surprises and ex ante capital choices. The aggregate term $\theta_A \hat{y}_{t+1}$ is common within a region-year and therefore drops out of cross-sectional comparisons. Substituting the closed-form policy functions reported in Appendix F into (20) yields the MRPK decomposition.

Proposition 4 (Cross-sectional MRPK differences) *Relative to the time-t cross-sectional mean, realized log MRPK satisfies*

$$\Delta^i \widehat{mrpk}_{i,t+1} = \underbrace{\left(\Delta^i \hat{a}_{i,t+1} - \mathbb{E}_t \Delta^i \hat{a}_{i,t+1}\right)}_{\textit{unexpected productivity}} + \underbrace{\Delta^i \hat{\tau}_{i,t+1}}_{\textit{capital wedge}} - \underbrace{\Theta(\hat{\beta}_i + \Psi_{i,t+1})\mathbb{VAR}_t(\hat{T}_{t+1})}_{\textit{temperature-risk adjustment}}. \tag{21}$$

The unexpected productivity component is

$$\Delta^i \hat{a}_{i,t+1} - \mathbb{E}_t \Delta^i \hat{a}_{i,t+1} = \underbrace{(\hat{\beta}_i + \Psi_{i,t+1})\left(\hat{T}_{t+1} - \mathbb{E}_t \hat{T}_{t+1}\right)}_{\textit{temperature forecast error}} + \underbrace{\xi_{i,t+1}(T_{t+1} - T^*)}_{\textit{random temperature damage}} + \varepsilon^z_{i,t+1}. \tag{22}$$

Proposition 4 points to three distinct sources of cross-sectional heterogeneity in a firm's realized capital return.

First, a positive productivity surprise raises the capital return because capital is predetermined. When realized productivity exceeds its expectation, firms cannot instantaneously adjust installed capital, and therefore the marginal product of capital rises. In a typical warm region, a heat surprise, $\widehat{T}_{t+1} - \mathbb{E}_t(\widehat{T}_{t+1}) > 0$, is common to all firms but generates heterogeneous productivity shocks through firm-specific effective exposure, $\widehat{\beta}_i + \Psi_{i,t+1}$.cFirms that are better adapted, in the sense of having a larger endogenous exposure component $\Psi_{i,t+1}$, or that are inherently less heat-sensitive, with higher $\widehat{\beta}_i$, benefit more from the same heat surprise. For example, an ice-cream parlor ($\widehat{\beta}_i > 0$) experiences a more favorable shock than a ski resort ($\widehat{\beta}_i < 0$). Among ice-cream parlors, the more productive and less distorted firms have a higher endogenous exposure index and therefore gain more from the same heat surprise. In addition, a transitory firm-specific disruption captured by $\xi_{i,t+1}(T_{t+1} - T^*)$ also affects productivity. For example, an unlucky firm hit by a weather-induced plant fire or equipment failure suffers a larger productivity loss than the average firm. The impact of these idiosyncratic disruptions is larger when temperature is farther from the bliss point.

Second, a higher capital wedge lowers investment, which mechanically raises the capital return. This is the standard misallocation force: distorted firms are systematically capital-scarce and therefore have higher marginal products. Finally, when $\Theta > 0$, higher temperature uncertainty lowers realized capital returns for firms with greater effective exposure. A firm whose profitability is more aligned with hot states provides higher payoffs precisely when hot states are high-SDF states. It therefore commands a more favorable risk premium and accumulates more capital ex ante. Holding realized productivity fixed, this larger capital stock reduces realized MRPK ex post.

Thus, the temperature-risk adjustment lowers realized MRPK especially for firms with higher effective exposure.

5.3.2 The Equilibrium MRPK Dispersion

This cross-sectional disparity in realized returns leads directly to the equilibrium misallocation of capital. Proposition 5 presents all the forces driving the dispersion.

Proposition 5 (Equilibrium MRPK dispersion) *Define the aggregate climate-shock component*

$$\mathcal{M}_{t+1} := \hat{T}_{t+1} - \mathbb{E}_t \hat{T}_{t+1} - \Theta \mathbb{VAR}_t(\hat{T}_{t+1})$$

The equilibrium MRPK dispersion is

$$\begin{aligned} \sigma^2_{mrpk,t+1} = \; & \underbrace{\mathbb{VAR}^i(\tau_{i,t+1})}_{\text{baseline wedge dispersion}} + \underbrace{\sigma^2_{\varepsilon^z}}_{\text{non-climate idiosyncratic productivity}} + \underbrace{\sigma^2_\xi (T_{t+1} - T^*)^2}_{\text{random climate damages}} \\ & + \underbrace{\mathbb{VAR}^i(\hat{\beta}_i + \Psi_{i,t+1})\, \mathcal{M}^2_{t+1}}_{\text{exposure-amplified forecast error and risk}} + \underbrace{2\mathbb{COV}^i(\tau_{i,t+1}, \hat{\beta}_i + \Psi_{i,t+1})\, \mathcal{M}_{t+1}}_{\text{wedge–exposure interaction}}. \end{aligned} \tag{23}$$

where

$$\begin{aligned} \mathbb{VAR}^i(\tau_{i,t+1}) &= \omega^2_{\bar{a}} \sigma^2_{\bar{a}} + \sigma^2_{\phi^\tau}, \\ \mathbb{VAR}^i(\hat{\beta}_i + \Psi_{i,t+1}) &= (1 + (\bar{T} - T^*)\Upsilon)^2 \sigma^2_{\hat{\beta}} + \Upsilon^2 \left[(1 - \alpha\omega_{\bar{a}})^2 \sigma^2_{\bar{a}} + \rho^2_z \sigma^2_z + \alpha^2 \sigma^2_{\phi^\tau} \right], \\ \mathbb{COV}^i(\tau_{i,t+1}, \hat{\beta}_i + \Psi_{i,t+1}) &= \Upsilon \left[\omega_{\bar{a}} (1 - \alpha\omega_{\bar{a}}) \sigma^2_{\bar{a}} - \alpha \sigma^2_{\phi^\tau} \right]. \end{aligned}$$

The decomposition in Proposition 5 separates three sources of MRPK dispersion. The first source is baseline misallocation. Dispersion in wedges, $\mathbb{VAR}^i(\tau_{i,t+1})$, and non-climate idiosyncratic productivity shocks, $\sigma^2_{\varepsilon^z}$, would generate MRPK dispersion even in the absence of temperature risk. These terms capture the standard forces emphasized in the misallocation literature: some firms enter the period with systematically different capital costs, and some firms receive non-climate productivity shocks after capital has been chosen.

The second source is idiosyncratic climate damage. The term $\sigma^2_\xi (T_{t+1} - T^*)^2$ shows that unpredictable firm-specific temperature sensitivity generates larger productivity gaps when realized temperature is farther from the bliss point. This channel does not require aggregate forecast errors. Even if aggregate temperature is perfectly anticipated, firms may still experience idiosyncratic climate disruptions, and the effect of those disruptions scales with the realized climate gap. The third source is the aggregate climate-shock component, $\mathcal{M}_{t+1} = \hat{T}_{t+1} - \mathbb{E}_t \hat{T}_{t+1} - \Theta \mathbb{VAR}_t(\hat{T}_{t+1})$. This object combines the realized temperature forecast error with the risk-adjusted temperature-uncertainty term. A common aggregate temperature shock affects firms differently because firms differ in effective exposure, $\hat{\beta}_i + \Psi_{i,t+1}$. The strength of this amplification is governed by $\mathbb{VAR}^i(\hat{\beta}_i + \Psi_{i,t+1})$, while the interaction with pre-existing misallocation is governed by $\mathbb{COV}^i(\tau_{i,t+1}, \hat{\beta}_i + \Psi_{i,t+1})$.

The variance term, $\mathbb{VAR}^i(\hat{\beta}_i + \Psi_{i,t+1})$, captures dispersion in both time-invariant product exposure and the endogenous exposure generated by adaptation incentives. This distinction is impor-

tant. Better average adaptation does not necessarily imply smaller misallocation responses. What matters for MRPK dispersion is not only the average level of adaptation, but also how adaptation changes the cross-sectional distribution of effective exposure. In the model, this distribution is shaped by the endogenous exposure multiplier Υ: firms with different product exposure, productivity, and wedges choose different adaptation intensities, so a common temperature shock can map into larger cross-sectional gaps in realized marginal products.

The covariance term, $\mathbb{COV}^i(\tau_{i,t+1}, \hat{\beta}_i + \Psi_{i,t+1})$, captures whether climate exposure is aligned with pre-existing capital distortions. If firms with high effective exposure also face high capital wedges, then temperature forecast errors and temperature risk operate on a cross-section in which climate exposure and capital misallocation reinforce each other. If exposure and wedges are negatively correlated, the same aggregate climate shock can instead partially offset pre-existing MRPK gaps. Thus, the sign and magnitude of $\mathbb{COV}^i(\tau_{i,t+1}, \hat{\beta}_i + \Psi_{i,t+1})$ determine whether climate shocks amplify or dampen baseline misallocation.

One implication is that cheaper adaptation need not improve allocation. Lower adaptation costs or weaker productivity curvature raise average adaptation $\widetilde{\beta}$, making the representative firm more resilient, but they can also raise the endogenous exposure multiplier Υ and widen dispersion in effective exposure across firms. Long-run warming reinforces this force by increasing the climate gap $\bar{T} - T^*$. As a result, a given temperature shock can generate larger MRPK dispersion even when firms adapt more on average, especially when highly exposed firms also face larger capital wedges. This provides a model-based interpretation of the reduced-form evidence that climate-induced misallocation is stronger in hotter and more developed regions.

Moreover, climate-induced misallocation is jointly shaped by aggregate climate risk and within-region firm heterogeneity. Aggregate temperature realizations and forecast quality determine the size of $\mathcal{M}_{t+1}$, but the mapping from $\mathcal{M}_{t+1}$ into MRPK dispersion depends on the regional distribution of product exposure, productivity, adaptation incentives, and capital wedges. Regions with similar climates can therefore display very different misallocation responses if their firms differ in how exposure is distributed or how exposure aligns with distortions. Conversely, regions with similar firm distributions can experience different misallocation outcomes under different temperature realizations and forecast uncertainty.

6 Reduced-Form Evidence for the Model Mechanisms

The model delivers two testable mechanisms linking temperature to MRPK dispersion: a *level channel* through which climate extremes raise damage volatility, and a *forecast-and-risk channel* through which temperature surprises and uncertainty are amplified by heterogeneous effective exposure. This section provides reduced-form evidence for both, first at the firm level, and then at the region-sector level.

6.1 Firm-Level Evidence

6.1.1 Adaptation Rises with Capital and Distortions

The linearized policy rules imply a within-firm capital–adaptation trade-off:

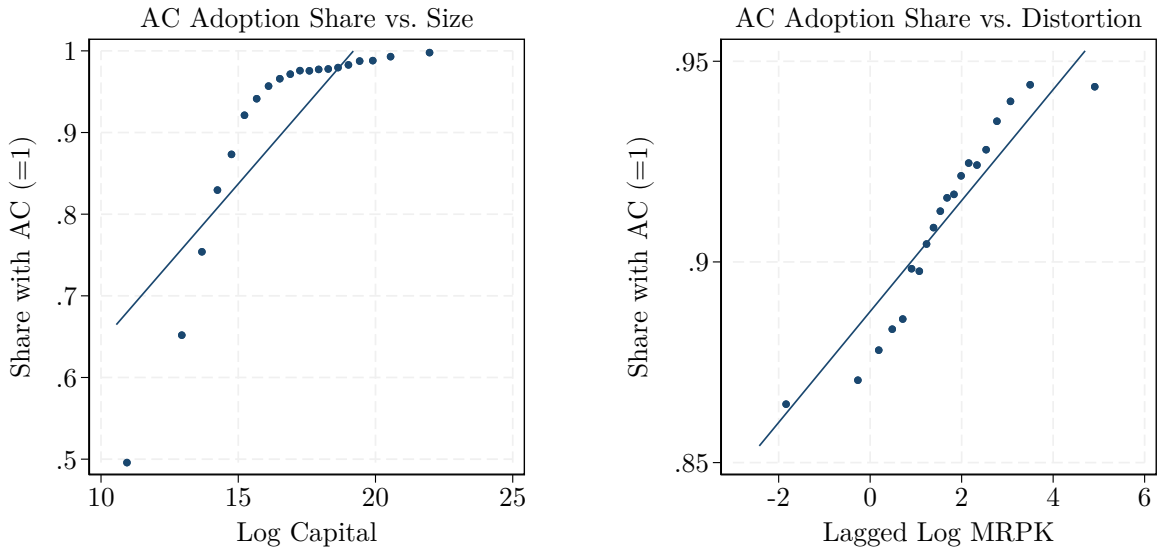
$$\frac{\Delta^i \widehat{\widetilde{\beta}}_{i,t+1}}{\widetilde{\beta} S_\beta/(\bar{T}-T^*)} = \Delta^i \hat{k}_{i,t+1} + \Delta^i \hat{\tau}_{i,t+1} + \frac{1}{S_\beta}\Delta^i \mathbb{COV}_t\left(\hat{a}_{i,t+1}, \hat{T}_{t+1}\right), \tag{24}$$

where Δ^i denotes deviations from the region-year mean and $S_\beta = (\bar{T}-T^*) - \Gamma_\beta \widetilde{\beta}$. In typical warm regions ($\widetilde{\beta} > 0$ and $S_\beta > 0$), larger firms choose higher adaptation, and firms facing higher capital wedges tilt toward adaptation relative to capital.[24] We test these predictions by estimating

$$\mathbf{1}\{AC_{it}=1\} = \lambda_1 k_{it} + \lambda_2 mrpk_{i,t-1} + \eta_{r,s,t} + u_{it}, \tag{25}$$

where $\eta_{r,s,t}$ are region–sector–year fixed effects and lagged MRPK proxies for capital wedges.[25] Consistent with (24), AC adoption is increasing in firm capital and in lagged MRPK conditional on scale (Figure 9).

Figure 9: AC adoption, capital and wedges



Notes: Figure 9 is constructed using the `binsreg` while controlling for firm and region-sector-year fixed effects.

6.1.2 Heterogeneous MRPK Responses to Heat Shocks

The model predicts that a common temperature surprise generates heterogeneous MRPK responses proportional to firms' effective exposure, $\hat{\beta}_i + \Psi_{i,t+1}$:

$$\Delta^i \widehat{mrpk}_{i,t+1} \;\propto\; \Delta^i(\hat{\beta}_i + \Psi_{i,t+1})\left(\hat{T}_{t+1} - \mathbb{E}_t[\hat{T}_{t+1}]\right).$$

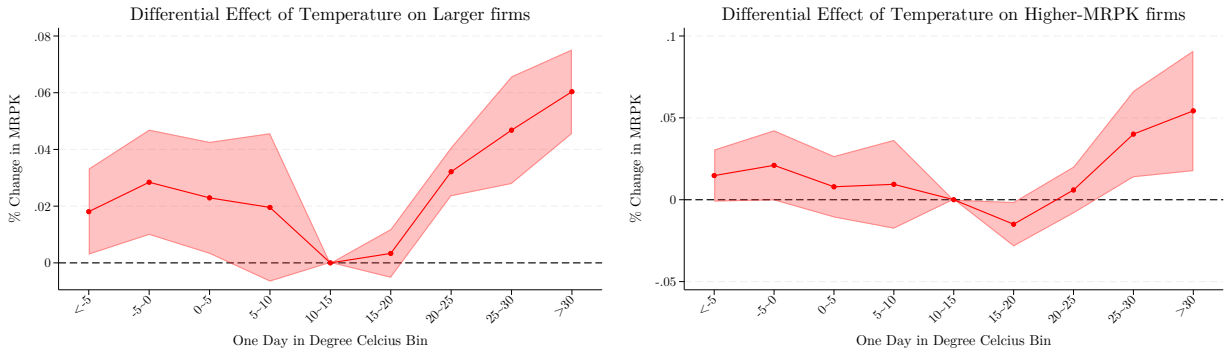
24. Direct measures of firm-level adaptation are rare. We proxy adaptation with computerized air-conditioning (AC) adoption from the Indian ASI, which records this variable for 2001–2014.

25. In the accounting framework, wedge-induced capital scarcity maps into systematically higher MRPK, making lagged MRPK a natural proxy for τ_{it}.

Conditional on intrinsic sensitivity, the endogenous exposure component $\Psi_{i,t+1}$ is shaped by firm scale and capital wedges. We therefore test this prediction by interacting temperature bins with lagged capital and lagged MRPK in a firm-level panel regression with firm and country–sector–year fixed effects.[26]

Figure 10 confirms the prediction: conditional on distortions, extreme heat shifts MRPK more for firms with more capital; conditional on capital, extreme heat shifts MRPK more for more distorted firms.

Figure 10: Differential MRPK response to heat shocks interacted with capital and distortions



(a) Heterogeneous Effect with Capital

(b) Heterogeneous Effect with Distortion

6.2 Region–Sector Evidence

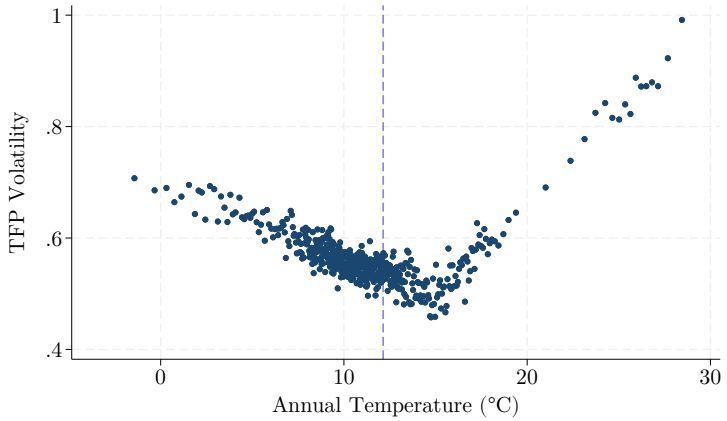
6.2.1 Mechanism 1: Climate Extremes Raise Damage Volatility

The level channel predicts that TFP volatility is minimized near T^* and rises as temperature departs from the bliss point. We measure TFP volatility as the within-cell variance of first-differenced firm productivity, $\mathbb{VAR}_{i\in(r,s,t)}(\hat{a}_{it}-\hat{a}_{i,t-1})$, following Asker, Collard-Wexler, and De Loecker (2014).[27] Figure 11 plots this measure against annual average temperature after residualizing on region–sector and country–sector–year fixed effects. The U-shape is clear: volatility is minimized around 12–15°C and rises as temperature moves away from that range.

26. Specifically, we estimate $mrpk_{i,t} = \sum_b \lambda_b\, Tbin^b_{r,t} + \sum_b \lambda_{b,k}\,(Tbin^b_{r,t} \times \widetilde{k}_{i,t-1}) + \sum_b \lambda_{b,\tau}\,(Tbin^b_{r,t} \times \widetilde{mrpk}_{i,t-1}) + \boldsymbol{\delta X}_{it} + \delta_i + \alpha_{s,c(r),t} + \varepsilon_{i,t}$, where $Tbin^b_{r,t}$ counts days in temperature bin b, $\widetilde{k}_{i,t-1}$ and $\widetilde{mrpk}_{i,t-1}$ are demeaned, and the omitted bin is 5–10°C. Using lagged variables aligns with the model's timing: capital and adaptation are chosen before temperature is realized.

27. We use first differences rather than innovations from the productivity law of motion to avoid sensitivity to specification of the AR process and agents' information set.

Figure 11: TFP volatility and average temperature.

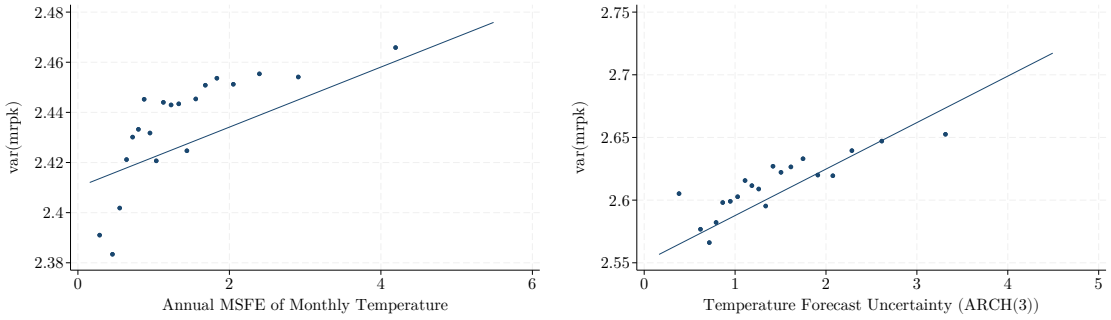


Notes: Figure 11 is constructed using the `binsreg` while controlling for region-sector and country-sector-year fixed effects.

6.2.2 Mechanism 2: Forecast Errors and Uncertainty Amplify MRPK Dispersion

Figure 12 shows that, after controlling for panel fixed effects, both larger forecast errors (annual MSFE of monthly temperature) and higher perceived forecast uncertainty (computed under ARCH(3)) are associated with greater MRPK dispersion. These patterns confirm the forecast-and-risk channel: when mid-range forecasts are less precise, the aggregate-weather component $\hat{T}_{t+1} - \mathbb{E}_t\hat{T}_{t+1} - \Theta\mathbb{VAR}_t(\hat{T}_{t+1})$ translates more strongly into cross-firm dispersion in capital returns.

Figure 12: Forecast Accuracy, Uncertainty and Misallocation



(a) Contemporaneous Forecast Accuracy (b) ARCH(3) Forecast uncertainty

Notes: Figure 12 is constructed using the `binsreg` while controlling for region-sector and country-sector-year fixed effects.

7 Model Estimation and Quantitative Results

7.1 Estimation

This section describes how we bring the model in Section 5 to the data and quantify the contribution of temperature-driven misallocation to aggregate climate damages. We estimate the model parameters jointly using firm panels spanning 1,916 regions in the EU, China, and India. The

closed-form solution delivers explicit policy rules and an accounting decomposition of realized MRPK dispersion into (i) wedge-driven dispersion, (ii) idiosyncratic climate damages that scale with distance from the bliss point, and (iii) forecast-error and forecast-uncertainty components that are amplified by heterogeneity in firms' effective exposure. This structure yields a transparent mapping from primitives to observable moments, which we exploit in a Generalized Method of Moments (GMM) estimation of regional adaptation costs and climate-damage parameters.

The need for Joint Estimation Two features of the model make joint estimation essential. First, the temperature bliss point T^* is identified from how temperature affects volatility—TFP volatility and MRPK dispersion—but volatility is itself shaped by endogenous adaptation. Pinning down T^* therefore requires disciplining adaptation costs, because heterogeneity in adaptation contributes mechanically to measured volatility. Second, the firm-level fundamentals that drive the relevant cross-sectional moments are not directly observed. In the model, realized productivity and MRPK reflect not only permanent and transitory productivity components $(\bar{a}_i, z_{it})$ and wedges (τ_{it}) , but also heterogeneous temperature exposure through intrinsic product sensitivity $(\hat{\beta}_i)$ and endogenous adaptation (Ψ_{it}). As a result, identification cannot proceed by simply inverting firm outcomes to recover primitives. Instead, bliss point T^*, product sensitivity $\hat{\beta}_i$ and the regional effective adaptation cost $\eta_r H_r$ governing the strength of Ψ_{it} must be jointly identified together with all micro-fundamentals in a model-consistent way under regional general equilibrium.

7.1.1 Parameterization and Measurement

We split parameters into (i) standard preference and technology parameters that we calibrate externally and impose to be common across regions, and (ii) climate- and region-specific objects that we discipline using moments from the data.

We take one model period to be one year and set the discount factor to $1/(1+r) = 0.96$. We set the CRRA coefficient to $\gamma = 5$, which lies within the median range commonly used in the macro literature (e.g., Mehra and Prescott (1985)). On the production side, we adopt the same demand elasticity as in the accounting exercise and set $\sigma = 4$. We set the Cobb–Douglas factor shares to $\hat{\alpha}_K = 0.35$ and $\hat{\alpha}_N = 0.65$. Under monopolistic competition, these imply revenue exponents $\alpha_F = (1 - 1/\sigma)\hat{\alpha}_F$ for $F \in \{K, N\}$, and hence the profit curvature $\alpha \equiv \frac{\alpha_K}{1-\alpha_N} = 0.51$. These parameters are held fixed across all regions in our baseline. With these parameters, we can map observables to model objects. Consistent with the model's profit function, we measure firm productivity as $a_{it} \equiv py_{it} - \alpha k_{it}$, where py_{it} and k_{it} are log sales and log capital, respectively. Similarly, up to constants common within a region-year (e.g., markups and aggregate price terms), log MRPK can be measured as $mrpk_{it} \equiv py_{it} - k_{it}$, which matches our reduced-form construction and is the sufficient statistic that enters the MRPK-dispersion moments.

As is standard in the literature, we assume the temperature bliss point T^* is common across regions, since it is identified from differential responses of TFP volatility to temperature variation across climates. We also impose that the volatility of random climate damages, σ_ξ^2, is common across regions, because all heterogeneity in product exposure and adaptation is captured by β_i and Ψ_i.[28] For product-level sensitivity, β_i, we assume it is common within each 3-digit USSIC

28. We also find limited heterogeneity in the TFP-volatility response to temperature across levels of economic devel-

sector.[29] For all other firm-level fundamentals (e.g., $\bar{a}_i$, z_{it}, and wedge components), we allow them to be firm-specific, with their cross-sectional moments varying by region.

Adaptation costs. The key regional object governing endogenous adaptation is the effective cost of adaptation resources, $\eta_r H_r$, which enters the model only through the product $\eta_r H_r$ and therefore cannot be separately identified. To allow adaptation costs to vary systematically with development, we parameterize the inverse cost as a flexible function of long-run income: $\frac{1}{\eta_r H_r} \equiv f\left(\overline{\ln \text{GDP}_r}\right)$, where $f(\cdot)$ is a fourth-order polynomial in $\overline{\ln \text{GDP}_r}$. Under this parameterization, steady-state adaptation intensity in region r can be written as $\widetilde{\beta}_r = \frac{1-\alpha_N}{1+r} f\left(\overline{\ln \text{GDP}_r}\right) (\bar{T}_r - T^*)$. Moreover, the endogenous-sensitivity multiplier that appears in the policy rules and in the MRPK dispersion, $\widetilde{\beta}_r / \big(1 - \alpha - \widetilde{\beta}_r(\bar{T}_r - T^*)\big)$, admits a convenient local approximation around $\bar{T}_r = T^*$ and $f = f^*$:

$$\frac{\widetilde{\beta}_r}{1 - \alpha - \widetilde{\beta}_r(\bar{T}_r - T^*)} = \frac{1 - \alpha_N}{(1+r)(1-\alpha)} \Big[f^*(\bar{T}_r - T^*) + (\bar{T}_r - T^*)\Big(f(\overline{\ln \text{GDP}_r}) - f^*\Big)\Big] + O\big(\Delta f^3, \Delta \bar{T}_r^3\big), \tag{26}$$

so that, under the cubic specification for $f(\cdot)$, the multiplier becomes a simple polynomial function of long-run temperature and long-run GDP per capita.

Temperature and Forecasts. Lastly, we estimate the temperature process. Following Nath, Ramey, and Klenow (2023), we fit a climate-dependent autoregressive process for annual regional temperature with two-way fixed effects:

$$T_{rt} = \sum_{\tau=1}^{p} \left(\rho_\tau^0 + \rho_\tau^{\bar{T}} \bar{T}_r^{1951\text{–}2018}\right) T_{r,t-\tau} + \chi_r + \chi_t + \varepsilon_{rt}^T, \tag{27}$$

where the persistence coefficients vary flexibly with a region's long-run mean temperature $\bar{T}_r^{1951\text{–}2018}$. We set $p = 6$ as our baseline to capture medium-run temperature persistence while remaining parsimonious in annual data. The fitted values from (27) define the one-step-ahead statistical forecast $\widehat{\mathbb{E}}_{t-1}[T_{rt}]$, and we use the estimated innovations $\widehat{\varepsilon}_{rt}^T \equiv T_{rt} - \widehat{\mathbb{E}}_{t-1}[T_{rt}]$ as our empirical counterpart to the model's annual forecast error $T_{rt} - \mathbb{E}_{t-1}[T_{rt}]$.

To discipline squared forecast errors and hence forecast uncertainty, we use mid-range forecast data from ECMWF as in our reduced-form exercise. Let MSFE_{rt} denote the region-year mean squared error of monthly ECMWF forecasts, aggregated to the annual level. Because the ECMWF product is monthly (rather than year-ahead), we rescale the implied monthly forecast errors to be consistent with the unconditional variance of annual forecast innovations in our sample (1998–2018). Concretely, we map MSFE_{rt} into an annual forecast-error variance by preserving ECMWF's relative improvement over a monthly AR(6) benchmark:

$$\big(T_{rt} - \mathbb{E}_{t-1} T_{rt}\big)^2 = \frac{\text{MSFE}_{rt}}{\frac{1}{12}\sum_{m=1}^{12} \mathbb{VAR}^{1998\text{–}2018}\big(\widehat{\varepsilon}_{rmt}^T\big)} \, \mathbb{VAR}^{1998\text{–}2018}\big(\widehat{\varepsilon}_{rt}^T\big), \tag{28}$$

opment.

29. To implement this in India and China, we harmonize the NIC-04 and GB-2002 classifications to USSIC at the 3-digit level.

where $\hat{\varepsilon}^T_{rmt}$ denotes the month-m AR(6) forecast error and $\hat{\varepsilon}^T_{rt}$ denotes the annual AR(6) forecast innovation. This scaling preserves time variation in forecast precision from ECMWF while matching the annual scale of temperature innovations in the estimation sample. Finally, we model perceived temperature uncertainty as a smoothed ARCH(3) process:

$$\mathbb{VAR}_{t-1}(T_{rt}) = \frac{1}{3}\sum_{\tau=1}^{3}\left(T_{r,t-\tau} - \mathbb{E}_{t-\tau-1}T_{r,t-\tau}\right)^2, \tag{29}$$

so that agents perceive temperature to be more volatile following periods with larger recent forecast errors.

7.1.2 Identification and Estimation

The model is identified from four sets of moment conditions, detailed in Appendix G. First, standard firm fixed-effect and Yule–Walker conditions recover the permanent and transitory components of productivity and wedges for each region, given temperature exposure parameters. Second, intrinsic product sensitivity β_i (common within 3-digit sectors) is identified from the orthogonality between residualized MRPK and temperature surprises, after netting out endogenous adaptation and region–year aggregates. Third, regional adaptation costs $(\eta_r H_r)^{-1}$, parameterized as a fourth-order polynomial in long-run income, are identified from how cross-firm variation in MRPK responses to temperature surprises covaries with model-implied drivers of endogenous adaptation. Fourth, the common bliss point T^* and random-damage dispersion σ^2_ξ are jointly identified from how residual TFP-growth volatility and MRPK dispersion scale with $(T_{rt} - T^*)^2$, after purging the components explained by heterogeneous exposure and forecast moments.

We implement the estimator via a fixed-point algorithm that iterates between inverting the model to recover firm-level latent objects (given current parameter guesses) and updating parameters using the moment conditions. The algorithm converges in fewer than 20 iterations across a wide range of initial conditions. Taken together, these moments jointly identify the common parameters $\{T^*, \sigma^2_\xi\}$ and the region-specific objects $\{\bar{\beta}_r,\ (\eta_r H_r)^{-1},\ \sigma^2_{\hat{\beta},r},\ \sigma^2_{\bar{a},r},\ \sigma^2_{\varepsilon^z,r},\ \sigma^2_{\varepsilon^{\phi^\tau},r},\ \rho_{\phi^\tau,r},\ \rho_{z,r},\ \omega_{a,r}\}_r$.

7.1.3 Estimation Results

We briefly summarize the estimation results. Table 2 reports the common parameters and the coefficients governing the regional adaptation-cost shifter $f(\cdot)$. For region-specific moments and implied parameters, we report the cross-region distribution (p5, p50, p95) in Table 3.

The model-consistent temperature bliss point remains close to its reduced-form counterpart, at roughly 11°C. We also identify substantial dispersion in idiosyncratic (random) climate damages, with $\sigma^2_\xi \simeq 0.0046$, implying that even conditional on the same (predictable) temperature realization, firms can experience materially different productivity impacts.

Regions' average intrinsic exposure $\bar{\beta}_r$ is tightly centered around zero and slightly negative: the 5th percentile is -0.0074, the median is essentially zero, and the 95th percentile is 0.006. This pattern suggests that, absent of endogenous adaptation, most regions exhibit a weakly negative response to temperature surprises, while a non-trivial right tail of regions responds positively.

Table 2: Estimated and Calibrated Parameters

Parameter	Description	Value
Panel A. Common estimated parameters		
T^*	Temperature bliss point (°C)	11.03
σ_ξ^2	Dispersion of random climate damages	0.0046
Panel B. Adaptation-cost shifter: $(\eta_r H_r)^{-1} = \sum_{j=0}^{4} c_j\, g_r^j,\ g_r \equiv \overline{\ln \mathrm{GDP}_r}$		
c_0	Constant	6.9462
c_1	Linear	−3.1359
c_2	Quadratic	0.5255
c_3	Cubic	−0.0387
c_4	Quartic	0.0011
Panel C. Calibrated parameters		
$\tilde{\alpha}_K$	Capital share	0.35
r	Interest rate	0.042
γ	CRRA risk aversion	5
σ	Demand elasticity of substitution	4

Table 3: Distribution of Region-Specific Estimates

Parameter	Description	p5	p50	p95
Climate exposure and adaptation				
$\bar{\beta}_r$	Avg. intrinsic product sensitivity	−0.0074	≈ 0	0.0060
$(\eta_r H_r)^{-1}$	Inverse effective adaptation cost	0.0006	0.0013	0.0063
$\sigma^2_{\hat{\beta},r}$	Within-region dispersion of $\hat{\beta}_i$	0.0001	0.0003	0.0007
Productivity and wedge dynamics				
$\sigma^2_{\bar{a},r}$	Permanent productivity dispersion	0.6516	1.5601	3.3667
$\sigma^2_{\varepsilon^z,r}$	Transitory productivity innovation var.	0.0925	0.1946	0.3019
$\sigma^2_{\varepsilon^{\phi^\tau},r}$	Wedge innovation variance	0.0327	0.1386	0.4110
$\rho_{z,r}$	Transitory productivity persistence	0.6401	0.7723	0.9086
$\rho_{\phi^\tau,r}$	Wedge persistence	0.6310	0.8103	0.9388
$\omega_{a,r}$	Size-dependent wedge coefficient	0.4713	0.9671	1.1840

Consistent with this interpretation, a sizeable share of regions in our sample have average temperatures below the bliss point.

At the product level, intrinsic sensitivities vary meaningfully across 3-digit categories, spanning roughly -0.11 to 0.09. Within a region, however, dispersion in intrinsic exposure is comparatively modest: $\sigma^2_{\hat{\beta},r}$ ranges from 10^{-4} to 7×10^{-4} (p5–p95). This stands in sharp contrast to the much larger dispersion in random climate damages captured by σ^2_ξ, highlighting that a substantial share of climate-induced heterogeneity operates through unpredictable firm-level disruptions rather than through persistent cross-product differences in exposure.

All regions exhibit positive adaptation responses. The inverse effective adaptation cost $(\eta_r H_r)^{-1}$ has a median of 0.0013 and spans 0.0006 to 0.0063 (p5–p95), implying nearly an order-of-magnitude variation in adaptation capacity across regions. Under the estimated income-based shifter $f(\overline{\ln \text{GDP}_r})$, adaptation is cheaper in higher-income regions on average, which translates into stronger predicted endogenous adaptation—especially in climates farther from the bliss point, where the marginal value of adaptation is higher.

Turning to region-specific parameters governing firm dynamics, we find large dispersion in permanent productivity and substantial idiosyncratic volatility. Permanent productivity dispersion $\sigma^2_{\bar{a},r}$ is sizable ($\text{p5} = 0.65$, $\text{p50} = 1.56$, $\text{p95} = 3.37$), and both transitory productivity and wedge components are highly persistent, with $\rho_{z,r}$ and $\rho_{\phi^\tau,r}$ typically between 0.6 and 0.95. Innovation variances in transitory productivity and idiosyncratic wedges are also economically large, indicating that both long-run heterogeneity and short-run shocks contribute meaningfully to within-region dispersion.

Finally, correlated distortions are strong. The size-dependent wedge coefficient $\omega_{a,r}$ is close to one at the median, implying that a 1% increase in permanent productivity (and hence scale) is associated with roughly a 1% higher effective user cost of capital. This magnitude is consistent with size-dependent distortion mechanisms recently estimated in Ayerst, Nguyen, and Restuccia (2024).

7.2 Quantitative Results

With the model estimated for regions spanning the EU, China, and India, we now quantify the aggregate importance of the misallocation channel. We use the estimated region-level moments to (i) measure the global productivity cost of climate-induced misallocation in an average year, (ii) track how this cost evolves from 1981–2019 as global warming unfolds and (iii) revisit the classic geography-development relationship by comparing the model-implied misallocation losses to cross-country productivity gaps in the macro data.

7.2.1 The Global Cost of Climate-Induced Misallocation

We first measure the total cost of climate-driven misallocation in our framework. We define the *global cost of climate-induced misallocation* as the aggregate TFP (or output) loss when comparing the observed economy to a counterfactual in which all within-region dispersion in climate sensitivities is eliminated. Operationally, for each region-year we shut down the three within-region dispersion objects that generate climate-driven MRPK dispersion: $\sigma^2_\xi = 0$, $\mathbb{VAR}^i_{rt}(\hat{\beta}_i + \Psi_{it}) = 0$

and $\mathbb{COV}^i_{rt}\left(\tau_{it+1}, \widehat{\beta}_i + \Psi_{i,t+1}\right) = 0$. All other region-level fundamentals are held fixed. Intuitively, this counterfactual removes the within-region heterogeneity in (effective) temperature exposure and idiosyncratic climate damages, as well as its interaction with wedges, thereby isolating the allocative-efficiency losses that operate through dispersion in MRPK.

Global implementation. To implement this exercise at the global level, we pair our micro estimates with gridded climate data (ERA5) and seasonal forecast data (ECMWF) for all countries worldwide. We partition the globe into roughly 4,894 regions. For China, India, and EU countries, we adopt the same regional definitions used in our empirical analysis; for all other countries, we use GADM level-1 partitions.

For regions in our estimation sample, we use the estimated firm-level moments directly. For regions outside the sample, we extrapolate these moments using long-run GDP per capita, leveraging the fact that firm-level dispersion moments covary systematically with development.[30] We also impute regional adaptation costs using the estimated mapping $\eta_r H_r = f(\bar{T}_r, \overline{\ln \text{GDP}}_r)$.

Aggregation. We aggregate the world economy using a Cobb–Douglas output index across countries and regions (see Appendix G for details). This implies

$$\log \text{TFP}_t^{\text{Global}} = \sum_r \omega_{rt} \log \text{TFP}_{rt}, \qquad \omega_{rt} \equiv \omega_{ct}\omega^c_{rt},$$

where ω_{ct} is country c's global output share and ω^c_{rt} is region r's share of country c's output.

Sufficient-statistics decomposition. Under this aggregation, the global misallocation cost implied by the model can be written as a weighted average of three sufficient statistics: (i) squared deviations from the optimal (bliss) temperature $(T_{rt} - T^*)^2$; (ii) forecast errors and forecast uncertainty, amplified by within-region dispersion in effective exposure; and (iii) the interaction between wedges and exposure. Define the relevant temperature innovation (forecast error net of the risk adjustment) as:

30. For parameters governing climate sensitivity, we extrapolate using both long-run GDP per capita and long-run climate.

$$\underbrace{\overline{\Delta^{T,\text{Mis}} \log \text{TFP}_t^{\text{Global}}}}_{\approx -12.5\%} = -\frac{\tilde{\alpha}_K + \tilde{\alpha}_K^2(\sigma-1)}{2}\Bigg[\underbrace{\sigma_\xi^2\Big(\overline{\sum_r \omega_{rt}(T_{rt}-T^*)^2}\Big)}_{\substack{\text{Random Climate Damage}\\ \approx -10.38\%}}$$
$$+\underbrace{\Big(\overline{\sum_r \omega_{rt}\mathbb{VAR}^i_{rt}(\hat{\beta}_i+\Psi_{it})\cdot\Big(T_{rt}-\mathbb{E}_{t-1}[T_{rt}]-(\bar{\beta}_r+\widetilde{\beta}_r)\Theta\mathbb{VAR}_{t-1}[T_{r,t}]\Big)^2}\Big)}_{\substack{\text{Exposure-amplified Forecast Error \& Risk}\\ \approx -2.11\%}}$$
$$+\underbrace{\Big(\overline{2\sum_r \omega_{rt}\mathbb{COV}^i_{rt}(\tau_{it},\hat{\beta}_i+\Psi_{it})\cdot\Big(T_{rt}-\mathbb{E}_{t-1}[T_{rt}]-(\bar{\beta}_r+\widetilde{\beta}_r)\Theta\mathbb{VAR}_{t-1}[T_{r,t}]\Big)}\Big)}_{\substack{\text{Wedge–exposure Interaction}\\ \approx -0.01\%}}\Bigg]. \tag{30}$$

Here $\overline{(\cdot)}$ denotes the time average over 1981–2019. Averaging these output-weighted sufficient statistics over 1981–2019, we find that climate-induced misallocation lowers global TFP by about 12.5% in an average year. Of this total, 10.38% arises from the *random climate damage* channel (persistent displacement from the bliss point, which magnifies heterogeneous and unpredictable firm-level disruptions), while 2.11% comes from forecast errors and climate uncertainty amplified by dispersion in effective exposure. The wedge–exposure interaction, while heterogeneous in the cross-section of regions, is quantitatively negligible in the global aggregate (about 0.01%).

The quantitative message is clear: most climate-induced misallocation reflects persistent deviations of regional climates from the optimal temperature—which magnify heterogeneous, unpredictable firm-level disruptions—rather than short-run forecast errors or forecast uncertainty. In other words, long-run climate displacement dominates weather predictability as a driver of global allocative inefficiency in our estimates.

The Changing Cost with Global Warming. We next quantify how the global misallocation cost evolves over time, feeding the observed paths of regional temperatures and ECMWF forecasts into (30) from $t_0 = 1981$ to $t_1 = 2019$. Figure 13a plots the implied evolution of the three components. Over this period, climate-induced misallocation becomes more costly, lowering global TFP by an additional

$$\Delta^{T,\text{Mis}}_{t_1,t_0} \log \text{TFP}^{\text{Global}} \approx -4.81\%.$$

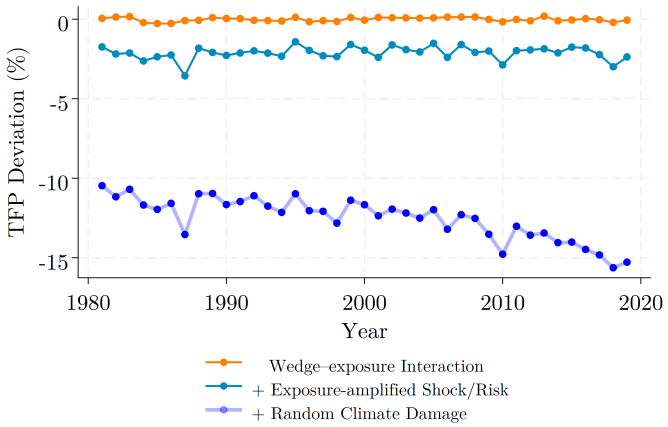
The dominant driver is again the level/random-damage channel: warming pushes many already-warm, high-output regions further away from T^*, more than it brings cooler regions closer, mechanically raising disruption-driven volatility and worsening allocative efficiency. The exposure-amplified forecast/risk component and the wedge–exposure interaction contribute little to the aggregate time trend.

7.2.2 Climate and Productivity Differences Across Countries: Macro Data vs. Model Estimates

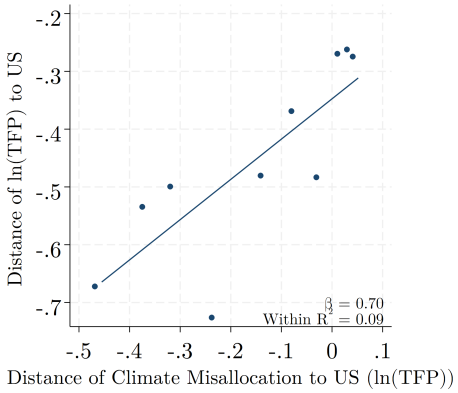
We next use our framework to revisit a central question in development economics: what explains large cross-country productivity gaps? Geography and climate have long been empha-

Figure 13: The Cost of Climate-Induced Misallocation and Cross-country Development

(a) The Changing Cost of Climate-Induced Misallocation



(b) Cross-country TFP Dispersion: Model vs. Data



sized in comparative development (Montesquieu 1748; Sachs and Warner 1997; Gallup, Sachs, and Mellinger 1999; Nordhaus 2006; Dell, Jones, and Olken 2012), but the mechanisms linking climate to aggregate productivity remain debated. Here we focus on one mechanism—the misallocation channel—and ask how much of the observed cross-country TFP dispersion can be explained with the model-implied climate-misallocation losses.

We match our model-generated country-year misallocation costs to country-level TFP from the Penn World Table (PWT) 10.01 (Feenstra, Inklaar, and Timmer 2015). For each country-year, we compare the measured TFP gap relative to the United States, $\ln \text{TFP}_{c,t} - \ln \text{TFP}_{\text{US},t}$, to the model-predicted climate-misallocation gap relative to the United States, $\Delta^{T,\text{Mis}} \log \text{TFP}_{c,t} - \Delta^{T,\text{Mis}} \log \text{TFP}_{\text{US},t}$. The U.S. is a natural benchmark as a productivity frontier, and it is also close to the estimated temperature optimum ($T^* \approx 11.03°\text{C}$), implying comparatively low climate-induced misallocation.

Figure 13b reports a binned scatter plot of macro TFP gaps against model-implied misallocation gaps, with both variables demeaned by year fixed effects to remove common global trends. The fitted relationship has a slope of about 0.7: when a country's predicted TFP distance to the U.S. due to climate misallocation rises by 1%, its measured TFP distance rises by about 0.7% on average. The corresponding R^2 of 0.094 implies that roughly 9% of observed cross-country TFP differences can be accounted for by climate through the misallocation channel. Notably, the model object is constructed solely from micro estimates combined with granular temperature and forecast data, yet it aligns well with the macro evidence.

8 Conclusion

This paper provides the first causal estimates of the misallocation effect from temperature shocks using firm-level microdata from 32 countries. On average, an additional hot day with $> 30°\text{C}$ of temperature increases MRPK dispersion by 0.29 log points and contributes to a 0.10% decline in annual aggregate TFP. Intriguingly, the detrimental impact of extreme heat on capital misallocation is more severe in regions with hotter climates and higher incomes, highlighting a significant

productivity cost of climate change coupled with a limited capacity for market-based adaptation. A novel firm dynamics model of endogenous adaptation rationalizes these pattern and reveals the quantitative important of random climate damages on TFP volatility and undermining the allocative efficiency investment. Climate-induced misallocation explains 9% of cross-country aggregate productivity differences.

References

Addoum, Jawad M, David T Ng, and Ariel Ortiz-Bobea. 2020. "Temperature shocks and establishment sales." *The Review of Financial Studies* 33 (3): 1331–1366.

Allcott, Hunt, Allan Collard-Wexler, and Stephen D O'Connell. 2016. "How do electricity shortages affect industry? Evidence from India." *American Economic Review.*

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Dynamic Inputs and Resource (Mis)Allocation." *Journal of Political Economy* 122 (5): 1013–1063.

Ayerst, Stephen, Duc M Nguyen, and Diego Restuccia. 2024. *The micro and macro productivity of nations.* Technical report. National Bureau of Economic Research.

Bakkensen, Laura A, and Lint Barrage. 2021. "Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics." *The Review of Financial Studies* 35 (8): 3666–3709.

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica* 87 (4): 1155–1203.

Basu, Susanto, and John G. Fernald. 2002. "Aggregate productivity and aggregate technology." *European Economic Review* 46 (6): 963–991.

Bau, Natalie, and Adrien Matray. 2023. "Misallocation and Capital Market Integration: Evidence From India." *Econometrica* 91 (1): 67–106.

Bilal, Adrien, and Diego R Känzig. 2024. *The Macroeconomic Impact of Climate Change: Global vs. Local Temperature.* Technical report. National Bureau of Economic Research.

Bils, Mark, Peter J. Klenow, and Cian Ruane. 2021. "Misallocation or Mismeasurement?" *Journal of Monetary Economics* 124:S39–S56.

Brandt, Loren, Johannes Van Biesebroeck, and Yifan Zhang. 2014. "Challenges of working with the Chinese NBS firm-level data." *China Economic Review.*

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

Caggese, Andrea, Andrea Chiavari, Sampreet Goraya, and Carolina Villegas-Sanchez. 2023. "Climate Change, Firms, and the Aggregate Productivity." Working Paper.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, et al. 2022. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *The Quarterly Journal of Economics* 137 (4): 2037–2105.

Casey, Gregory, Stephie Fried, and Matthew Gibson. 2022. "Understanding Climate Damages: Consumption versus Investment." Federal Reserve Bank of San Francisco Working Paper 2022-21.

Chen, Zhang. 2022. "Economic Growth and the Rise of Large Firms."

Christidis, Nikolaos, Dann Mitchell, and Peter A. Stott. 2023. "Rapidly increasing likelihood of exceeding 50°C in parts of the Mediterranean and the Middle East due to human influence." *npj Climate and Atmospheric Science* 6, no. 1 (May 26, 2023): 45.

Coeurdacier, Nicolas, Helene Rey, and Pablo Winant. 2011. "The risky steady state." *American Economic Review* 101 (3): 398–401.

Copernicus Climate Change Service and Climate Data Store. 2018. *Seasonal forecast daily and sub-daily data on single levels.*

Cruz, José-Luis, and Esteban Rossi-Hansberg. 2023. *The Economic Geography of Global Warming.*

David, Joel M., Hugo A. Hopenhayn, and Venky Venkateswaran. 2016. "INFORMATION, MISALLOCATION, AND AGGREGATE PRODUCTIVITY." *The Quarterly Journal of Economics* 131 (2): 943–1006. Accessed November 27, 2023.

David, Joel M., and Venky Venkateswaran. 2019. "The Sources of Capital Misallocation." *American Economic Review* 109 (7): 2531–67.

David, Joel M., and David Zeke. 2021. *Risk-Taking, Capital Allocation and Optimal Monetary Policy.* Federal Reserve Bank of Chicago.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics* 4 (3): 66–95.

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

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3 (4): 152–185.

Dundas, Steven J., and Roger H. von Haefen. 2020. "The Effects of Weather on Recreational Fishing Demand and Adaptation: Implications for a Changing Climate." *Journal of the Association of Environmental and Resource Economists* 7 (2): 209–242.

Fan, Jingting. 2024. "Talent, Geography, and Offshore R&D." *The Review of Economic Studies* (April): rdae044. ISSN: 0034-6527.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer. 2015. "The next generation of the Penn World Table." *American economic review* 105 (10): 3150–3182.

Gallup, John Luke, Jeffrey D. Sachs, and Andrew D. Mellinger. 1999. "Geography and Economic Development." *International Regional Science Review* 22 (2): 179–232.

Ginglinger, Edith, and Quentin Moreau. 2023. "Climate risk and capital structure." *Management Science* 69 (12): 7492–7516.

Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. 2017. "Capital Allocation and Productivity in South Europe." *Quarterly Journal of Economics* 132 (4).

Gorodnichenko, Yuriy, Debora Revoltella, Jan Svejnar, and Christoph T Weiss. 2018. *Resource misallocation in European firms: The role of constraints, firm characteristics and managerial decisions.* Technical report. National Bureau of Economic Research.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4): 1403–1448.

Kalemli- Özcan, Şebnem, Bent E. Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş. 2024. "How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration." *American Economic Journal: Macroeconomics* 16, no. 2 (April): 353–74.

Kopytov, Alexandr, Mathieu Taschereau-Dumouchel, and Zebang Xu. 2024. "The Origin of Risk." *Available at SSRN.*

Krusell, Per, and Anthony A. Smith Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867–896.

Lemoine, Derek. 2018. *Estimating the Consequences of Climate Change from Variation in Weather.* Working Paper, Working Paper Series 25008. National Bureau of Economic Research.

Mehra, Rajnish, and Edward C Prescott. 1985. "The equity premium: A puzzle." *Journal of monetary Economics* 15 (2): 145–161.

Montesquieu, Baron de, Charles de Secondat. 1748. *The Spirit of Laws.* Printed for T. Evans.

Moscona, Jacob, and Karthik A Sastry. 2023. "Does directed innovation mitigate climate damage? evidence from us agriculture." *The Quarterly Journal of Economics* 138 (2): 637–701.

Nath, Ishan. 2025. "Climate change, the food problem, and the challenge of adaptation through sectoral reallocation." *Journal of Political Economy.*

Nath, Ishan, Valerie Ramey, and Peter Klenow. 2023. "How Much Will Global Warming Cool Global Growth?"

Nordhaus, William D. 1992. "An optimal transition path for controlling greenhouse gases." *Science* 258 (5086): 1315–1319.

———. 2006. "Geography and Macroeconomics: New Data and New Findings." *Proceedings of the National Academy of Sciences* 103 (10): 3510–3517.

O'Neill, B. C., C. Tebaldi, D. P. van Vuuren, V. Eyring, P. Friedlingstein, G. Hurtt, R. Knutti, et al. 2016. "The Scenario Model Intercomparison Project (ScenarioMIP) for CMIP6." *Geosci. Model Dev.*

Oberfield, Ezra, and Devesh Raval. 2021. "Micro data and macro technology." *Econometrica.*

Oudin Åström, Daniel, Bertil Forsberg, Kristie L. Ebi, and Joacim Rocklöv. 2013. "Attributing mortality from extreme temperatures to climate change in Stockholm, Sweden." *Nature Climate Change* 3, no. 12 (December 1, 2013): 1050–1054.

Piao, Shilong, Philippe Ciais, Yao Huang, Zehao Shen, Shushi Peng, Junsheng Li, Liping Zhou, Hongyan Liu, Yuecun Ma, Yihui Ding, et al. 2010. "The impacts of climate change on water resources and agriculture in China." *Nature.*

Ponticelli, Jacopo, Qiping Xu, and Stefan Zeume. 2023. *Temperature and Local Industry Concentration.* Technical report w31533. National Bureau of Economic Research.

Poschke, Markus. 2018. "The Firm Size Distribution across Countries and Skill-Biased Change in Entrepreneurial Technology." *American Economic Journal: Macroeconomics* 10 (3): 1–41.

Redding, Stephen J, and David E Weinstein. 2020. "Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences." *The Quarterly Journal of Economics* 135 (1): 503–560.

Restuccia, Diego, and Richard Rogerson. 2008. "Policy distortions and aggregate productivity with heterogeneous establishments." *Review of Economic dynamics* 11 (4): 707–720.

Rudik, I., G. Lyn, W. Tan, and A. Ortiz-Bobea. 2021. *The Economic Effects of Climate Change in Dynamic Spatial Equilibrium.*

Sabater, Joaquin Muñoz. 2019. *ERA5-Land monthly averaged data from 1950 to present.* Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Sachs, Jeffrey D., and Andrew M. Warner. 1997. "Sources of Slow Growth in African Economies." *Journal of African Economies* 6 (3): 335–376.

Somanathan, E., Rohini Somanathan, Anant Sudarshan, and Meenu Tewari. 2021. "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing" [in en]. *Journal of Political Economy* 129 (6). Accessed November 27, 2023.

Sraer, David, and David Thesmar. 2023. "How to Use Natural Experiments to Estimate Misallocation." *American Economic Review* 113 (4): 906–38.

Tarsia, Romano. 2023. "Heterogeneous effects of weather shocks on firm economic performance." *Available at SSRN 4672552.*

Tyazhelnikov, Vladimir, Xinbei Zhou, and Xuetao Shi. 2024. *PPML, Gravity, and Heterogeneous Trade Elasticities.* Technical report.

Wang, Tingting, and Fubao Sun. 2022. "Global gridded GDP data set consistent with the shared socioeconomic pathways." *Scientific Data.*

Wenz, Leonie, Robert Devon Carr, Noah Kögel, Maximilian Kotz, and Matthias Kalkuhl. 2023. "DOSE –Global data set of reported sub-national economic output." *Scientific Data* 10 (1): 425.

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

Zivin, Joshua Graff, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics* 32 (1): 1–26.

Online Appendix

A Additional Derivations for Section 2

This appendix provides the derivation of the expressions and propositions featured in Section 2.

A.1 Equilibrium Allocations in Distorted Equilibrium

In this appendix, we solve the firm's profit maximization problem in equation (3) and derive equation (4).

Subject to the inverse demand and wedges, each firm i in region-sector n engages in monopolistic competition and optimally chooses its quantity of inputs and price to maximize profits

$$\max_{P_{nit},K_{nit},L_{nit}} \left(1-\tau^Y_{nit}\right) P_{nit} \underbrace{A_{nit}K^{\alpha_{Kn}}_{nit}L^{\alpha_{Ln}}_{nit}}_{Y_{nit}} - \left(1+\tau^K_{nit}\right) R_{nt}K_{nit} - \left(1+\tau^L_{nit}\right) W_{nt}L_{nit}$$
$$\text{subject to : } Y_{nit} = B_{nit}Y_{nt}\left[\frac{P_{nit}}{P_{nt}}\right]^{-\sigma_n}.$$

After substituting the demand curve $P_{nit}Y_{nit} = P_{nt}Y^{\frac{1}{\sigma_n}}_{nt}B^{\frac{1}{\sigma_n}}_{nit}Y^{\frac{\sigma_n-1}{\sigma_n}}_{nit}$ into the objective function, we derive the first-order condition with respect to any factor input F_{nit} used (where F_{nit} can be either K_{nit} or L_{nit})

$$MRPF_{nit} = \frac{\partial P_{nit}Y_{nit}}{\partial F_{nit}} = \alpha_{F_n}\frac{\sigma_n-1}{\sigma_n}\frac{P_{nit}Y_{nit}}{F_{nit}} = \frac{1+\tau^F_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)}{1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)}P^F_{nt},$$

where P^F_{nt} denotes the factor price, specifically W_{nt} for labor and R_{nt} for capital. $\frac{\sigma_n-1}{\sigma_n}$ is the optimal CES markup. This is equation (4) in the main text.

equation (4) can be rearranged to yield the following:

$$K_{nit} = \alpha_{K_n}(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\frac{P_{nit}Y_{nit}}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))P^K_{nt}}$$
$$L_{nit} = \alpha_{L_n}(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\frac{P_{nit}Y_{nit}}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))P^L_{nt}}$$

Plugging the labor and capital first order conditions into the expression for production yields a relationship between output Y_{nit} and sale $P_{nit}Y_{nit}$

$$\begin{aligned} Y_{nit} &= A_{nit}K^{\alpha_{Kn}}_{nit}L^{\alpha_{Ln}}_{nit} \\ &= A_{nit}\alpha^{\alpha_{Kn}}_{K_n}\alpha^{\alpha_{Ln}}_{L_n}(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}P_{nit}Y_{nit}\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))P^K_{nt}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))P^L_{nt}}\right]^{\alpha_{Ln}} \end{aligned} \tag{31}$$

Using demand curve once again solves for sales $P_{nit}Y_{nit}$

$$
\begin{aligned}
P_{nit}Y_{nit} &= P_{nt}Y_{nt}^{\frac{1}{\sigma_n}}B_{nit}^{\frac{1}{\sigma_n}}Y_{nit}^{\frac{\sigma_n-1}{\sigma_n}} \\
&= P_{nt}Y_{nt}^{\frac{1}{\sigma_n}}B_{nit}^{\frac{1}{\sigma_n}}\left\{A_{nit}\alpha_{K_n}^{\alpha_{K_n}}\alpha_{L_n}^{\alpha_{L_n}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}P_{nit}Y_{nit}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}}\right\}^{\frac{\sigma_n-1}{\sigma_n}} \\
(P_{nit}Y_{nit})^{\frac{1}{\sigma_n}} &= P_{nt}Y_{nt}^{\frac{1}{\sigma_n}}B_{nit}^{\frac{1}{\sigma_n}}\left\{A_{nit}\alpha_{K_n}^{\alpha_{K_n}}\alpha_{L_n}^{\alpha_{L_n}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}}\right\}^{\frac{\sigma_n-1}{\sigma_n}} \\
P_{nit}Y_{nit} &= P_{nt}^{\sigma_n}Y_{nt}B_{nit}\left\{A_{nit}\alpha_{K_n}^{\alpha_{K_n}}\alpha_{L_n}^{\alpha_{L_n}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}
\end{aligned}
\tag{32}
$$

We aggregate sales across all firms and obtain

$$
\begin{aligned}
\int_i P_{nit}Y_{nit} &= \int_i P_{nt}^{\sigma_n}Y_{nt}B_{nit}\left\{A_{nit}\alpha_{K_n}^{\alpha_{K_n}}\alpha_{L_n}^{\alpha_{L_n}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1} \\
&= P_{nt}^{\sigma_n}Y_{nt}\left(\alpha_{K_n}^{\alpha_{K_n}}\alpha_{L_n}^{\alpha_{L_n}}\frac{\sigma_n-1}{\sigma_n}\left(P_{nt}^{K}\right)^{-\alpha_{K_n}}\left(P_{nt}^{L}\right)^{-\alpha_{L_n}}\right)^{\sigma_n-1}\int_i B_{nit}\left\{A_{nit}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}
\end{aligned}
\tag{33}
$$

Then the sales share of firm i in region-sector n, θ_{nit} can be expressed in closed form of the fundamentals and parameters as

$$
\begin{aligned}
\theta_{nit} &= \frac{P_{nit}Y_{nit}}{\int_i P_{nit}Y_{nit}} \\
&= \frac{B_{nit}\left\{A_{nit}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}}{\int_i B_{nit}\left\{A_{nit}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}},
\end{aligned}
$$

A.2 Proof of Proposition 1

Proposition 1. Equilibrium (Mis)Allocation. The (log) ratio of firm i's distorted and efficient equilibrium allocation of factor, $\frac{F_{nit}}{F_{nit}^*}$,

$$
\log(\frac{F_{nit}}{F_{nit}^*}) = -\underbrace{\log(\frac{1+\tau_{nit}^{F}}{1+\tau_{nt}^{F}})}_{\substack{\text{relative wedge}\\ \text{effect}}} + \underbrace{\log(\frac{\theta_{nit}}{\theta_{nit}^*})}_{\substack{\text{size}\\ \text{effect}}}
$$

is decreasing in the ratio of firm i's own factor wedge comparing to the aggregate factor wedge $1+\tau_{nt}^{F} \equiv \left(\int_0^{J_n}\frac{1}{(1+\tau_{nit}^{F})}\theta_{nit}di\right)^{-1}$, and increasing in the ratio of the firm's sales share θ_{nit} comparing to its efficient counterfactual θ_{nit}^*. Moreover, the efficient allocation of inputs, $\frac{F_{nit}^*}{F_{nt}} = \theta_{nit}^* = \frac{B_{nit}A_{nit}^{\sigma_n-1}}{\int_0^{J_n}B_{nit}A_{nit}^{\sigma_n-1}di}$, is entirely determined by firm i's relative preference shifter and physical productivity within the region-sector and aligns with sales share θ_{nit}^*.

Proof. We first re-write the expression for output Y_{nit} in equation (31) by plugging in for P_{nit}

$$\begin{aligned}
Y_{nit} &= A_{nit}K_{nit}^{\alpha_{Kn}}L_{nit}^{\alpha_{Ln}} \\
&= A_{nit}\alpha_{Kn}^{\alpha_{Kn}}\alpha_{Ln}^{\alpha_{Ln}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}\underbrace{P_{nit}}_{\left(\frac{Y_{nit}}{B_{nit}Y_{nt}}\right)^{-\frac{1}{\sigma_n}}P_{nt}}Y_{nit}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}} \\
Y_{nit}^{\frac{1}{\sigma_n}} &= A_{nit}\alpha_{Kn}^{\alpha_{Kn}}\alpha_{Ln}^{\alpha_{Ln}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}P_{nt}\left(B_{nit}Y_{nt}\right)^{\frac{1}{\sigma_n}}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}} \\
Y_{nit} &= B_{nit}Y_{nt}\left\{A_{nit}\alpha_{Kn}^{\alpha_{Kn}}\alpha_{Ln}^{\alpha_{Ln}}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\frac{\sigma_n-1}{\sigma_n}P_{nt}\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{K}}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))P_{nt}^{L}}\right]^{\alpha_{Ln}}\right\}^{\sigma_n}
\end{aligned} \tag{34}$$

Next we solve for factor allocation K_{nit} and L_{nit} and express them to be a function of the wedges, aggregate variables, prices and parameters. Plugging equation (34) back to the first-order conditions of capital and labor, and rearranging terms in the FOCs to connect to sales share θ_{nit} yields

$$\begin{aligned}
F_{nit} = Z_{nt}\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}\int_i B_{nit}\left\{A_{nit}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}di \\
F\in(K,L),
\end{aligned} \tag{35}$$

where $Z_{nt}=\alpha_{K_n}\left[\frac{\sigma_n-1}{\sigma_n}P_{nt}\right]^{\sigma_n}Y_{nt}\frac{1}{P_{nt}^{K}}\left[\left(\frac{\alpha_{Kn}}{P_{nt}^{K}}\right)^{\alpha_{Kn}}\left(\frac{\alpha_{Ln}}{P_{nt}^{L}}\right)^{\alpha_{Ln}}\right]^{\sigma_n-1}$.

Aggregate across firms:

$$\begin{aligned}
F_{nt} &= \int_i F_{nit}di \\
&= Z_{nt}\int_i B_{nit}\left\{A_{nit}(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau_{ni}^{K}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau_{ni}^{L}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}di \\
&\cdot\int_i\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}di
\end{aligned}$$

The ratio between firm and region allocation $\frac{F_{nit}}{F_{nt}}$ is

$$\frac{F_{nit}}{F_{nt}}=\frac{\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}}{\int_i\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}di}$$

We can evaluate the expression at the efficient equilibrium where all distortions are eliminated and obtain

$$\frac{F_{nit}^{*}}{F_{nt}^{*}}=\frac{\theta_{nit}^{*}}{\underbrace{\theta_{nt}^{*}}_{=1}}=\theta_{nit}^{*}$$

Finally, combining this with equation (A.2) yields

$$\frac{F_{nit}/F_{nt}}{F_{nit}^{*}/F_{nt}^{*}}=\frac{F_{nit}}{F_{nit}^{*}}=\frac{\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}}{\int_i\frac{(1-\tau_{ni}^{Y}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau_{ni}^{F}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}di}\frac{\theta_{nit}}{\theta_{nit}^{*}},$$

where the aggregate factor is exogenously given and $K_{nt} = K^*_{nt}$. Taking log yields a more intuitive expression

$$\begin{aligned}\log \frac{F_{nit}}{F^*_{nit}} &= \log \frac{\frac{(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau^F_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}}{\int_i \frac{(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau^F_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\theta_{nit}di} + \log \frac{\theta_{nit}}{\theta^*_{nit}} \\ &= \log \frac{\frac{(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}{(1+\tau^F_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}}{(1+\tau^F_{nt})} + \log \frac{\theta_{nit}}{\theta^*_{nit}}\end{aligned}.$$

When factor-F distortions $(1+\tau^F_{nit})$ are the only sources of distortions present, we arrive at the exact expression as in (5) in Lemma 1.

■

A.3 Proof of Proposition 2

Proposition 2. Aggregation and TFP Decomposition. Each region–sector n admits an aggregate production function of the form

$$Y_{nt} = \text{TFP}_{nt} K_{nt}^{\alpha_{Kn}} L_{nt}^{\alpha_{Ln}},$$

where, under a second-order approximation, region–sectoral aggregate total factor productivity $\text{TFP}_{nt} := \text{TFP}_n\left(\tilde{\mathbf{T}}_{rt}, \tilde{\mathbf{X}}_{nt}\right)$ can be decomposed as follows:

$$\begin{aligned}\log \text{TFP}_n(\tilde{\mathbf{T}}_{rt},\cdot) = &\underbrace{\frac{1}{\sigma_n-1}\log\left[\mathbb{E}_i\left[B_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)\left(A_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)\right)^{\sigma_n-1}\right]\right]}_{\text{Technology}(\log \text{TFP}^E_n)} + \frac{1}{\sigma_n-1}\log J_n \\ &\underbrace{-\frac{\sigma_n}{2}\,\text{var}_{\log\left(1-\tau^Y_{ni}\right)}(\tilde{\mathbf{T}}_{rt},\cdot)}_{\text{Output Wedge Dispersion}} \\ &\underbrace{-\sum_{F\in\{K,L\}}\frac{\alpha_{Fn}+\alpha^2_{Fn}(\sigma_n-1)}{2}\,\text{var}_{\log\left(1+\tau^F_{ni}\right)}(\tilde{\mathbf{T}}_{rt},\cdot)}_{\text{Factor Wedge Dispersion}} \\ &\underbrace{+\sigma_n\sum_{F\in\{K,L\}}\alpha_{Fn}\,\text{cov}_{\log\left(1-\tau^Y_{ni}\right),\log\left(1+\tau^F_{ni}\right)}\left(\tilde{\mathbf{T}}_{rt},\cdot\right)}_{\text{Output-Factor Mixed Distortion}} \\ &\underbrace{-(\sigma_n-1)\alpha_{Kn}\alpha_{Ln}\,\text{cov}_{\log\left(1+\tau^K_{ni}\right),\log\left(1+\tau^L_{ni}\right)}\left(\tilde{\mathbf{T}}_{rt},\cdot\right)}_{\text{Factor Mixed Distortion}}\end{aligned} \tag{36}$$

Proof. Note that $P_{nt} = 1$. Rearranging equation (35) yields

$$\left(\frac{\alpha_{L_n}}{W_{nt}}\right)^{\alpha_{L_n}(\sigma_n-1)} = \frac{1}{\left(\frac{\sigma_n-1}{\sigma_n}\right)^{\sigma_n-1}\left(\frac{\alpha_{K_n}}{R_{nt}}\right)^{\alpha_{K_n}(\sigma_n-1)} I_{nt}},$$

where $I_{nt} = \int_i B_{nit}\left\{A_{nit}(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$. Plugging equa-

tion (A.3) to aggregate capital yields

$$K_{nt} = Y_{nt}\frac{\alpha_{K_n}}{R_{nt}}\left[\frac{\sigma_n-1}{\sigma_n}\right]^{\sigma_n}\left[\left(\frac{\alpha_{K_n}}{R_{nt}}\right)^{\alpha_{Kn}}\left(\frac{\alpha_{L_n}}{W_{nt}}\right)^{\alpha_{Ln}}\right]^{\sigma_n-1}J_{nt}$$
$$\frac{K_{nt}}{Y_{nt}} = \left(\frac{\alpha_{K_n}}{R_{nt}}\right)\frac{\sigma_n-1}{\sigma_n}\frac{J_{nt}}{I_{nt}},$$

where $J_{nt} = \int_i \left[(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\right]^{\sigma_n} B_{nit}\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\left\{A_{nit}\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$.
Taking logs yield

$$\log\frac{Y_{nt}}{K_{nt}} = \log I_{nt} - \log J_{nt} - \log\left(\frac{\sigma_n-1}{\sigma_n}\right)\frac{\alpha_{K_n}}{R_{nt}}.$$

Using equation (A.3) once again to obtain aggregate labor

$$L_{nt} = Y_{nt}\left(\frac{\sigma_n-1}{\sigma_n}\right)^{\frac{-\alpha_{K_n}}{1-\alpha_{K_n}}}\left(\frac{\alpha_{K_n}}{R_{nt}}\right)^{-\frac{\alpha_{K_n}}{1-\alpha_{K_n}}}Q_{nt}I_{nt}^{-1-\frac{1}{\alpha_{L_n}(\sigma_n-1)}},$$

where $Q_{nt} = \int_i \left[(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\right]^{\sigma_n} B_{nit}\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\left\{A_{nit}\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$.
Taking logs yield

$$\log\frac{Y_{nt}}{L_{nt}} = \left(1+\frac{1}{\alpha_{L_n}(\sigma_n-1)}\right)\log I_{nt} - \log Q_{nt} + \frac{\alpha_{K_n}}{1-\alpha_{K_n}}\log\frac{\sigma_n-1}{\sigma_n}\frac{\alpha_{K_n}}{R_{nt}}.$$

This allows to find aggregate TFP in the distorted economy:

$$\begin{aligned}\text{TFP}_{nt} &:= \frac{Y_{nt}}{K_{nt}^{\alpha_{Kn}}L_{nt}^{\alpha_{Ln}}}\\ \log TFP_{nt} &= \alpha_{K_n}\log\frac{Y_{nt}}{K_{nt}} + \alpha_{L_n}\log\frac{Y_{nt}}{L_{nt}}\\ &= \frac{\sigma_n}{\sigma_n-1}\log I_{nt} - \alpha_{K_n}\log J_{nt} - \alpha_{L_n}\log Q_{nt},\end{aligned}$$

where

$$\log I_{nt} = \log\int_i B_{nit}\left\{A_{nit}(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$$
$$\log J_{nt} = \log\int_i \left[(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\right]^{\sigma_n} B_{nit}\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\left\{A_{nit}\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$$
$$\log Q_{nt} = \log\int_i \left[(1-\tau^Y_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))\right]^{\sigma_n} B_{nit}\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\left\{A_{nit}\left[\frac{1}{(1+\tau^K_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Kn}}\left[\frac{1}{(1+\tau^L_{ni}(\tilde{\mathbf{T}}_{rt},\cdot))}\right]^{\alpha_{Ln}}\right\}^{\sigma_n-1}$$

We use the following second-order approximation. Let x_i be any firm-level term and normalize the measure so that $\int di = 1$ (so $\int \cdot\, di = \mathbb{E}[\cdot]$). Define

$$\mu \equiv \mathbb{E}[x_i], \qquad \tilde{x}_i \equiv x_i - \mu, \qquad \mathbb{E}[\tilde{x}_i] = 0.$$

Then

$$\ln\left(\int e^{x_i}\,di\right) = \ln\left(\int e^{\mu+\tilde{x}_i}\,di\right) = \mu + \ln\left(\int e^{\tilde{x}_i}\,di\right).$$

A second-order expansion of $e^{\tilde{x}_i}$ around zero implies

$$\int e^{\tilde{x}_i}\, di \approx \int \left(1 + \tilde{x}_i + \frac{1}{2}\tilde{x}_i^2\right) di = 1 + \frac{1}{2}\mathbb{E}[\tilde{x}_i^2] = 1 + \frac{1}{2}\mathbb{VAR}(x_i).$$

Next, apply the second-order expansion of $\ln(1+u)$:

$$\ln(1+u) = u - \frac{u^2}{2} + O(u^3), \qquad u = \frac{1}{2}\mathbb{VAR}(x_i),$$

to obtain

$$\ln\left(\int e^{\tilde{x}_i}\, di\right) = \ln\left(1 + \frac{1}{2}\mathbb{VAR}(x_i)\right) = \frac{1}{2}\mathbb{VAR}(x_i) - \frac{1}{8}\mathbb{VAR}(x_i)^2 + O\big(\mathbb{VAR}(x_i)^3\big)\,.$$

Combining terms yields

$$\ln\left(\int e^{x_i}\, di\right) = \mathbb{E}[x_i] + \frac{1}{2}\mathbb{VAR}(x_i) - \frac{1}{8}\mathbb{VAR}(x_i)^2 + O\big(\mathbb{VAR}(x_i)^3\big)\,.$$

Writing $x_i = \mu + \varepsilon z_i$ with $\mathbb{E}[z_i] = 0$ and $\mathbb{VAR}(z_i) = 1$ implies $\mathbb{VAR}(x_i) = \varepsilon^2$, so $\mathbb{VAR}(x_i)^2 = O(\varepsilon^4)$. Hence, under a second-order approximation,

$$\ln\left(\int e^{x_i}\, di\right) = \mathbb{E}[x_i] + \frac{1}{2}\mathbb{VAR}(x_i) + O(\varepsilon^4),$$

For a multivariate case when x_i is a sum of logs, $x_i = \log A_i + \log B_i + \log C_i$, this implies

$$\begin{aligned}\ln\left(\int e^{\log A_i + \log B_i + \log C_i}\, di\right) \approx\ & \mathbb{E}[\log A_i] + \mathbb{E}[\log B_i] + \mathbb{E}[\log C_i] \\ & + \frac{1}{2}\mathbb{VAR}(\log A_i) + \frac{1}{2}\mathbb{VAR}(\log B_i) + \frac{1}{2}\mathbb{VAR}(\log C_i) \\ & + \mathbb{COV}(\log A_i, \log B_i) + \mathbb{COV}(\log A_i, \log C_i) + \mathbb{COV}(\log B_i, \log C_i).\end{aligned}$$

We apply this second-order approximation to all the $\log I_{nt}, \log J_{nt}, \log Q_{nt}$ to obtain

$$\begin{aligned}\log \mathrm{TFP}_n(\tilde{\mathbf{T}}_{rt}, \cdot) =\ & \underbrace{\frac{1}{\sigma_n - 1}\log \mathbb{E}_i\left[B_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot)\left(A_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot)\right)^{\sigma_n - 1}\right]}_{\text{Technology}(\log \mathrm{TFP}_n^E)} - \underbrace{\frac{\sigma_n}{2}\,\mathrm{var}_{\log\left(1-\tau_{ni}^Y\right)}(\tilde{\mathbf{T}}_{rt}, \cdot)}_{\text{Output Wedge Dispersion}} \\ & \underbrace{- \sum_{F \in \{K,L\}} \frac{\alpha_{Fn} + \alpha_{Fn}^2(\sigma_n - 1)}{2}\,\mathrm{var}_{\log\left(1+\tau_{ni}^F\right)}(\tilde{\mathbf{T}}_{rt}, \cdot)}_{\text{Factor Wedge Dispersion}} \\ & \underbrace{+ \sigma_n \sum_{F \in \{K,L\}} \alpha_{Fn}\,\mathrm{cov}_{\log\left(1-\tau_{ni}^Y\right), \log\left(1+\tau_{ni}^F\right)}\left(\tilde{\mathbf{T}}_{rt}, \cdot\right)}_{\text{Output-Factor Mixed Distortion}} \\ & \underbrace{- (\sigma_n - 1)\alpha_{Kn}\alpha_{Ln}\,\mathrm{cov}_{\log\left(1+\tau_{ni}^K\right), \log\left(1+\tau_{ni}^L\right)}\left(\tilde{\mathbf{T}}_{rt}, \cdot\right)}_{\text{Factor Mixed Distortion}}\end{aligned} \tag{37}$$

which is equation (7) in the main text.

■

B Data Sources and Variable Construction

This appendix details the data sources, cleaning procedures, and variable construction that supplement Section 3.

B.1 Orbis Data (Europe)

Our European firm data come from the Orbis Historical Disk Product maintained by Bureau van Dijk (now Moody's). Orbis aggregates administrative filings from over 170 information providers across countries, including national business registries and regulatory bodies (Kalemli- Özcan et al. 2024). Under EU Directive 2013/34/EU, limited liability companies in all 30 countries in our sample are required to file financial statements, which gives Orbis broad coverage of the formal sector—approximately 99% of covered firms are private entities. Since our analysis centers on MRPK dispersion, we retain every firm-year observation with non-missing revenue and fixed assets, which maximizes coverage of the revenue-to-capital ratio.[31] The pooled sample spans 1998–2018; country-specific coverage periods are listed in Table B.1.

Sample cleaning. Following Kalemli- Özcan et al. (2024), we link multiple vintages of Orbis products over time and match each firm's descriptive information with its financial records via the unique BvD firm identifier (BVDID). We then apply the following cleaning steps:

1. We restrict to firms whose country of residence (latest address) matches the ISO-2 code embedded in their BVDID. Firms with mismatched codes are excluded.
2. For firms lacking address information but possessing postcodes, we manually map postcodes to NUTS 3 regions within each country.
3. We harmonize fiscal years with calendar years based on the closing date: if the closing date is on or after July 1, we assign the current calendar year; otherwise, the previous year.
4. Following Fan (2024), we use unconsolidated accounts to avoid double-counting that can arise with consolidated reporting. When a firm reports multiple sales figures from different sources (e.g., local registries, annual reports), we prioritize the unconsolidated entry.
5. We keep only observations with non-missing and positive sales (`operating_revenue_turnover`) and fixed assets (`fixed_assets`).
6. We compute firm-level log MRPK as

$$\log \text{MRPK}_{it} = \log\Big(\alpha_{Kn} \frac{\sigma_n - 1}{\sigma_n} \frac{P_{it}Y_{it}}{K_{it}}\Big),$$

   where $P_{it}Y_{it}$ is measured with sales and K_{it} with fixed assets.

31. This choice follows the same principle as Kalemli- Özcan et al. (2024), Gopinath et al. (2017), and Nath (2025), who each restrict their samples to the variables required by their respective analyses.

7. We winsorize MRPK, fixed assets, and sales at the top and bottom 0.1% of the distribution within each country to prevent outliers from distorting variance estimates.

B.2 China NBS Data

The NBS Annual Survey covers all industrial firms with annual sales above CNY 5 million (1998–2007). These "above-scale" firms account for over 90.7% of total industrial output when benchmarked against the 2004 NBS economic census (Brandt, Van Biesebroeck, and Zhang 2014). Manufacturing constitutes more than 90% of observations; the remaining firms are in mining and utilities. We restrict the sample to 1998–2007 because of well-documented reporting inconsistencies after 2008 (Brandt, Van Biesebroeck, and Zhang 2014; Nath 2025).

Sample cleaning. We follow Zhang et al. (2018). Sales are measured using "□□□□□□" (product sales revenue) and capital using "□□□□□□" (total fixed assets). Each firm is classified by its four-digit Chinese Industry Classification (CIC) code, which we harmonize to the USSIC division level. Firm locations are mapped to prefecture-level divisions. The remaining cleaning (variable construction and winsorization) follows items 5–7 of the Orbis procedure in Appendix B.1.

B.3 India ASI Data

The ASI is a census of large plants (100+ workers) combined with a probability sample of smaller registered plants under the Indian Factories Act (Allcott, Collard-Wexler, and O'Connell 2016).[32] Almost all sampled plants are in manufacturing. The sample period is 1998–2018.

Sample cleaning. We match plants to Indian districts following Somanathan et al. (2021) and harmonize industry codes first to NIC-04, then to USSIC divisions. Sales are measured using the gross sale value of all products; capital is measured as the average of opening and closing net book value of total fixed assets, following Bils, Klenow, and Ruane (2021).[33] The ASI also records whether a plant has computerized air-conditioning (2001–2014), which we exploit as an adaptation proxy in Section 6. The remaining cleaning follows items 5–7 of Appendix B.1.

B.4 Temperature Variable Construction

Starting from hourly ERA5-Land data at 0.1° resolution, we compute the daily maximum and minimum 2 m air temperature for each grid cell and take their geometric mean as the daily average.[34] We aggregate to the regional level using area-weighted averages over all grid cells intersecting each NUTS 3 region, prefecture, or district. Annual temperature bins are then constructed by counting the number of days each region's daily average temperature falls within each 5°C interval: < -5, $[-5, 0)$, $[0, 5)$, ..., $[25, 30)$, ≥ 30°C. Because the bins sum to 365 (or 366) days per year, we omit the $[5, 10)$°C bin as the reference category in all regressions to avoid perfect collinearity.

32. The census threshold was 200 workers during 1997–2003. The sampling fraction for smaller plants was one-third until 2004 and one-fifth thereafter.

33. We use net rather than gross fixed assets because gross values in ASI contain substantial missing data. In Section 3 we note that results are not sensitive to this choice.

34. The daily maximum (minimum) is the highest (lowest) value among the 24 hourly readings within a grid cell.

B.5 Forecast Error and Uncertainty Construction

We construct three forecast-related objects from the ECMWF seasonal forecast data: monthly forecast errors, annual mean squared forecast errors (MSFE), and perceived forecast uncertainty. This subsection details each step.

Monthly forecast errors. For each region r and month m, we compute the raw forecast error as the difference between realized monthly mean temperature and the ensemble-mean ECMWF seasonal forecast issued on the first day of that month. We then remove a region-by-month fixed effect $\widehat{\text{Bias}}_{m,r}$, to correct for systematic biases in the forecast product that may reflect differences in forecast methodology across regions and seasons.

Annual MSFE. We aggregate monthly squared forecast errors to the annual level by averaging across all twelve months within each region-year:

$$\text{MSFE}_{\text{annual},r,t} = \frac{1}{12}\sum_{m=1}^{12}\left(T_{m,r,t} - \mathbb{E}_{m-1}[T_{m,r,t}] - \widehat{\text{Bias}}_{m,r}\right)^2.$$

We also construct seasonal variants (summer: April–September; winter: October–March) for robustness. Because the ECMWF product provides month-ahead forecasts rather than year-ahead forecasts, the monthly MSFE and the annual forecast-error variance from the AR(6) model in Section 7 operate at different horizons. To map the ECMWF-based MSFE into an annual-scale squared forecast error consistent with the model, we rescale by preserving the ECMWF forecast's relative improvement over a monthly AR(6) benchmark:

$$\left(T_{rt} - \mathbb{E}_{t-1}T_{rt}\right)^2 = \frac{\text{MSFE}_{rt}}{\frac{1}{12}\sum_{m=1}^{12}\mathbb{VAR}^{1998\text{–}2018}\left(\widehat{\varepsilon}^{T}_{rmt}\right)}\,\mathbb{VAR}^{1998\text{–}2018}\left(\widehat{\varepsilon}^{T}_{rt}\right),$$

where $\widehat{\varepsilon}^{T}_{rmt}$ is the month-m AR(6) forecast error and $\widehat{\varepsilon}^{T}_{rt}$ is the annual AR(6) innovation. This scaling preserves time variation in forecast precision from ECMWF while matching the annual scale of temperature innovations in the estimation sample.

Perceived forecast uncertainty. We model perceived temperature uncertainty as a smoothed ARCH(3) process:

$$\mathbb{VAR}_{t-1}(T_{rt}) = \frac{1}{3}\sum_{\tau=1}^{3}\left(T_{r,t-\tau} - \mathbb{E}_{t-\tau-1}T_{r,t-\tau}\right)^2,$$

so that agents perceive temperature to be more volatile following periods with larger recent forecast errors. This object enters the model's risk-adjustment channel through the households' stochastic discount factor.

B.6 Regional GDP Matching

The DOSE dataset (Wenz et al. 2023) provides sub-national GDP in PPP-adjusted 2017 international dollars. We match DOSE units to our regional definitions (NUTS 3, prefectures, districts)

using spatial joins on administrative boundary shapefiles from GADM. Long-run averages of regional GDP per capita and mean temperature (computed over 1997–2018) serve as time-invariant moderators in the heterogeneity analysis of Section 4.4. For end-of-century projections, we use grid-level SSP3 GDP shares from Wang and Sun (2022) to construct projected regional output weights.

B.7 Descriptive Statistics

Table B.1 lists the countries, sample periods, number of regions, and total firm-year observations in the final sample used for empirical analysis.

Table B.1: Descriptive Table of Dataset by Country

Dataset	Country	Coverage	Number of Regions	Number of Firm-Year Obs
NBS China Industrial	China	1998-2007	325	2069100
India ASI	India	1998-2017	267	473646
BvD Orbis	Austria	2004-2018	34	185029
	Belgium	1998-2018	44	781322
	Bulgaria	1998-2018	28	1439871
	Switzerland	2003-2018	3	1701
	Cyprus	2005-2018	1	17356
	Czech Republic	1998-2018	14	1425901
	Germany	1998-2018	346	765981
	Denmark	1999-2018	11	392015
	Estonia	1998-2018	5	791340
	Greece	1998-2018	44	342636
	Spain	1998-2018	52	12210663
	Finland	1998-2018	19	1909080
	France	1998-2018	96	15151185
	Croatia	1998-2018	21	1202515
	Hungary	2004-2018	20	3095326
	Ireland	2000-2018	8	115344
	Italy	1998-2018	107	12083926
	Lithuania	1998-2018	10	129442
	Luxembourg	1998-2018	1	74287
	Latvia	2010-2018	6	533640
	Malta	2000-2018	1	21249
	Netherlands	1998-2018	36	201182
	Norway	1998-2018	12	2462277
	Poland	1998-2018	73	1207428
	Portugal	1998-2018	24	3882515
	Romania	1998-2018	42	4636047
	Sweden	1998-2018	21	3934403
	Slovenia	1998-2018	12	744495
	Slovakia	1998-2018	8	1008353
	United Kingdom	1998-2018	178	3537701

B.8 Sample Trimming

1. 0.83% of raw sample firms are dropped for lacking nuts_id, (mention how we went through this, and if this is mentioned previously then mention this early on, mention where are these firms

from geographically so people are less scared

2. observations lacking all of the following variables are dropped (Y,K,WL,L, vadir)
3. trimming happened in within country, 0.1-99 percent.
4. winsorizing happened in within province-sector-year

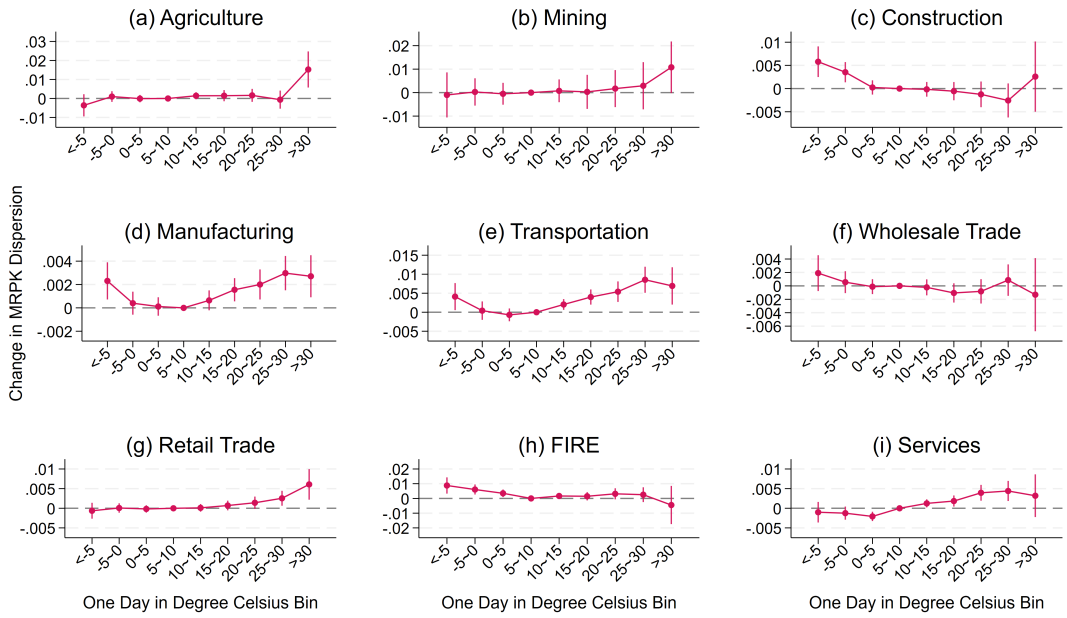
C Robustness for the Empirical Results in Section 4

This appendix examines the sensitivity of the main reduced-form results (Figure 2) to the choice of outcome variable, region–sector definition, weather measure, standard-error clustering, winsorization threshold, and production-function specification.

C.1 Sectoral Results

Figure C.1 reports the estimated effect of daily temperature shocks on regional MRPK dispersion separately for each SIC division. Extreme heat (above 25–30°C) significantly increases MRPK dispersion in agriculture, manufacturing, transportation, services, and retail trade, whereas the effects are modest in mining, construction, wholesale trade, and FIRE. This pattern is consistent with the model's prediction that sectors with greater heterogeneity in product-level climate sensitivity—and with production processes more directly exposed to ambient temperature—exhibit larger MRPK dispersion responses to heat.

Figure C.1: Estimated impact of daily temperature shocks on annual MRPK dispersion by SIC Divisions



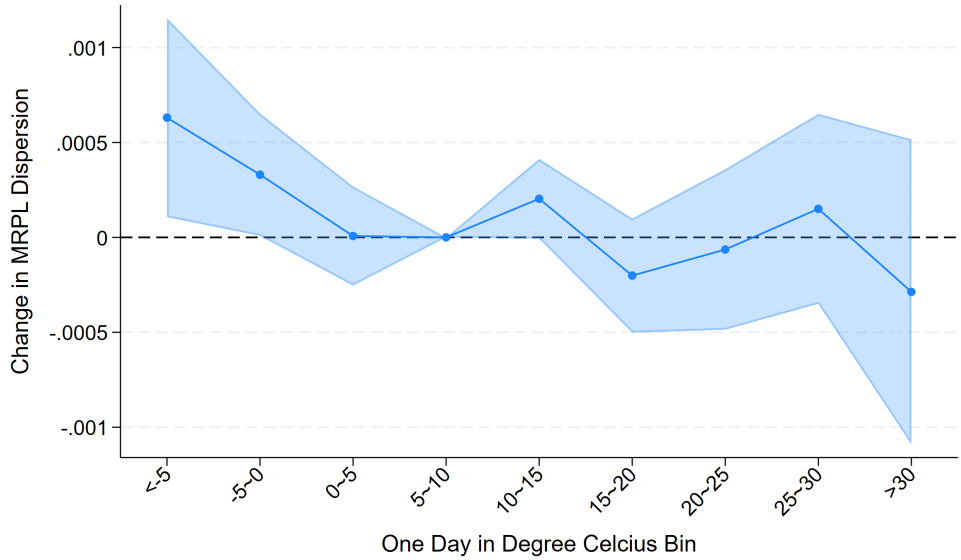
C.2 Temperature's Impact on MRPL and MRPM Dispersion

A natural question is whether temperature shocks also disperse the marginal revenue products of labor and materials. If so, the misallocation channel would extend beyond capital. Figure C.2 and

Figure C.3 show that, in contrast to MRPK, neither MRPL nor MRPM dispersion exhibits a statistically significant U-shaped response to temperature. This is consistent with the model's timing: capital and adaptation are committed before temperature is realized, so temperature surprises create wedges between the ex ante capital choice and ex post productivity. Labor and materials, by contrast, can be adjusted after temperature is observed, allowing firms to partially re-optimize and attenuating dispersion in their marginal products.

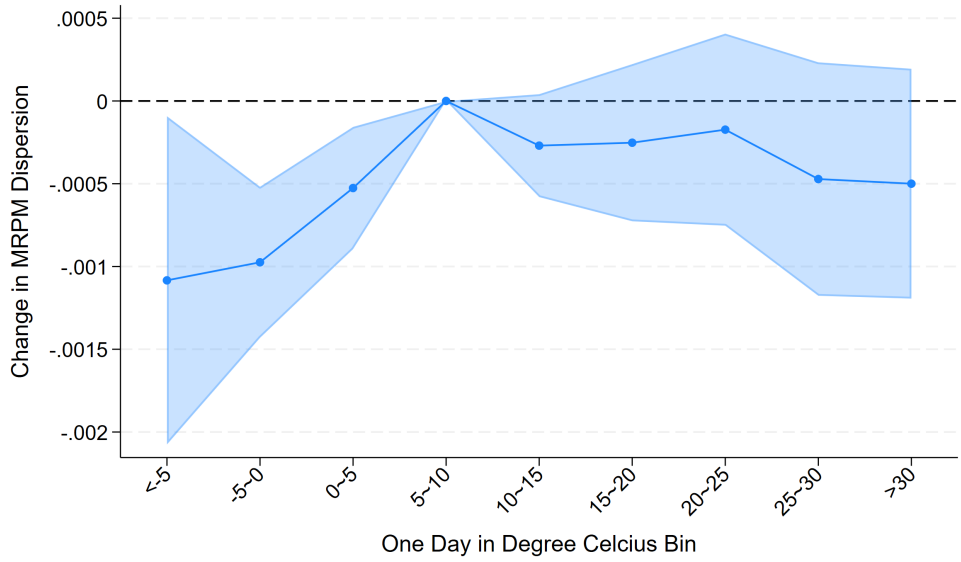
C.2.1 MRPL Dispersion

Figure C.2: Estimated impact of daily temperature shocks on annual MRPL dispersion



C.2.2 MRPM Dispersion

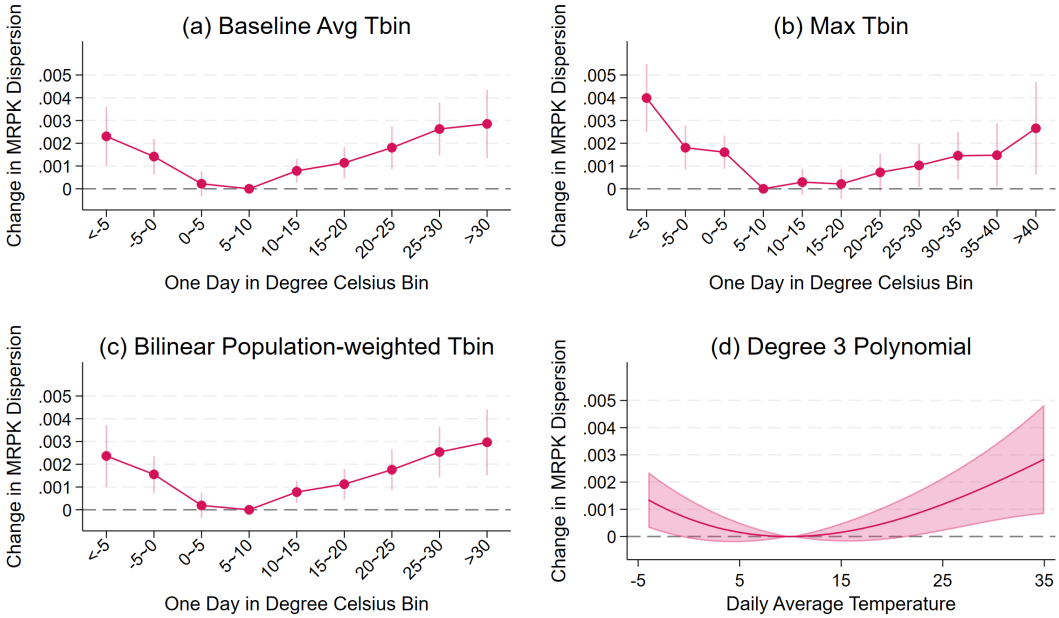
Figure C.3: Estimated impact of daily temperature shocks on annual MRPM dispersion



C.3 Alternative Weather Variables and Weighting

Figure C.4 demonstrates that the main results are robust to alternative constructions of the temperature variable. Panel (a) reproduces the baseline using daily mean temperature bins. Panel (b) replaces mean with daily maximum temperature and extends the bin structure to cover the resulting warmer distribution ($> 30, > 35, > 40$°C). Panel (c) uses population-weighted regional temperature. Panel (d) replaces temperature bins with a cubic polynomial in daily average temperature. The U-shaped relationship between temperature and MRPK dispersion is present across all four specifications.

Figure C.4: Robustness with Alternative Weather Variables



Notes: Panel (a): baseline (daily mean temperature bins). Panel (b): daily maximum temperature bins. Panel (c): population-weighted regional temperature bins. Panel (d): cubic polynomial in daily average temperature.

C.4 Alternative Standard-Error Clustering

Figure C.5 reports alternative standard-error specifications: clustering by region, two-way clustering by region and year, and two-way clustering by region and country–year. Table C.2 reports the corresponding regression coefficients. Figure C.6 shows robustness to spatial HAC standard errors. Our results remain statistically significant even when allowing for spatial correlation up to a cutoff distance of 1,000 km.

Figure C.5: Robustness with Alternative Standard-Error Clustering

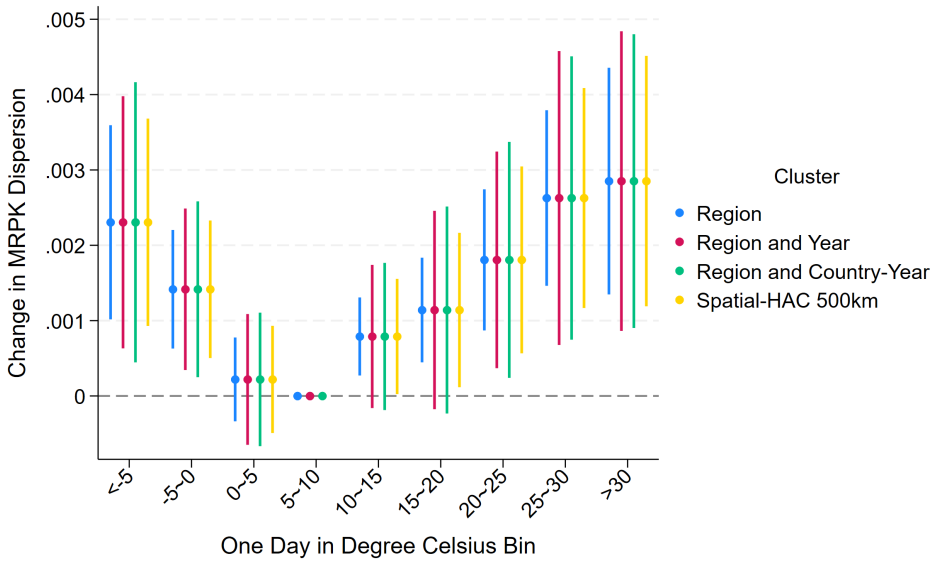


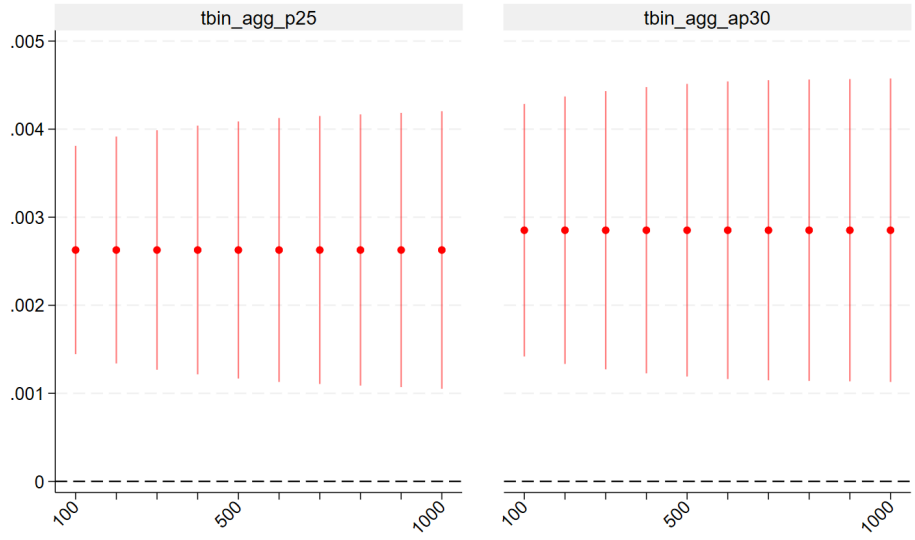
Table C.2: Regression Coefficients: Alternative Standard-Error Clustering

	(1)	(2)	(3)	(4)
tbin_agg_bm5	0.0019**	0.0019*	0.0019*	0.0019**
	(0.0008)	(0.0010)	(0.0011)	(0.0008)
tbin_agg_m5	0.0009*	0.0009	0.0009	0.0009*
	(0.0005)	(0.0007)	(0.0007)	(0.0005)
tbin_agg_p0	-0.0000	-0.0000	-0.0000	-0.0000
	(0.0003)	(0.0005)	(0.0005)	(0.0004)
tbin_agg_p10	0.0006**	0.0006	0.0006	0.0006
	(0.0003)	(0.0005)	(0.0006)	(0.0004)
tbin_agg_p15	0.0010**	0.0010	0.0010	0.0010*
	(0.0004)	(0.0008)	(0.0008)	(0.0006)
tbin_agg_p20	0.0016***	0.0016*	0.0016*	0.0016**
	(0.0006)	(0.0008)	(0.0009)	(0.0007)
tbin_agg_p25	0.0024***	0.0024**	0.0024**	0.0024***
	(0.0007)	(0.0011)	(0.0011)	(0.0009)
tbin_agg_ap30	0.0028***	0.0028**	0.0028**	0.0028***
	(0.0009)	(0.0012)	(0.0012)	(0.0010)
Observations	129,105	129,105	129,105	129,105
R^2	0.871	0.871	0.871	0.000

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure C.6: Robustness with Alternative Distance Cutoffs under Spatial HAC



C.5 Alternative Region–Sector Definitions

The baseline defines region–sector cells as NUTS 3 × SIC 1-digit. To verify that results are not driven by the level of spatial or sectoral aggregation, we re-estimate the main specification under three alternative definitions: a coarser regional unit (NUTS 2) and two finer sectoral units (SIC 2- and SIC 3-digit).

NUTS 2–SIC 1. Figure C.7 redefines the unit as NUTS 2 × SIC 1-digit and re-estimates the baseline under alternative fixed-effect structures. Results are quantitatively similar to the baseline. Table C.3 reports the underlying coefficients.

Figure C.7: Results Using NUTS 2–SIC 1 as the Region–Sector Unit, Under Alternative Fixed Effects

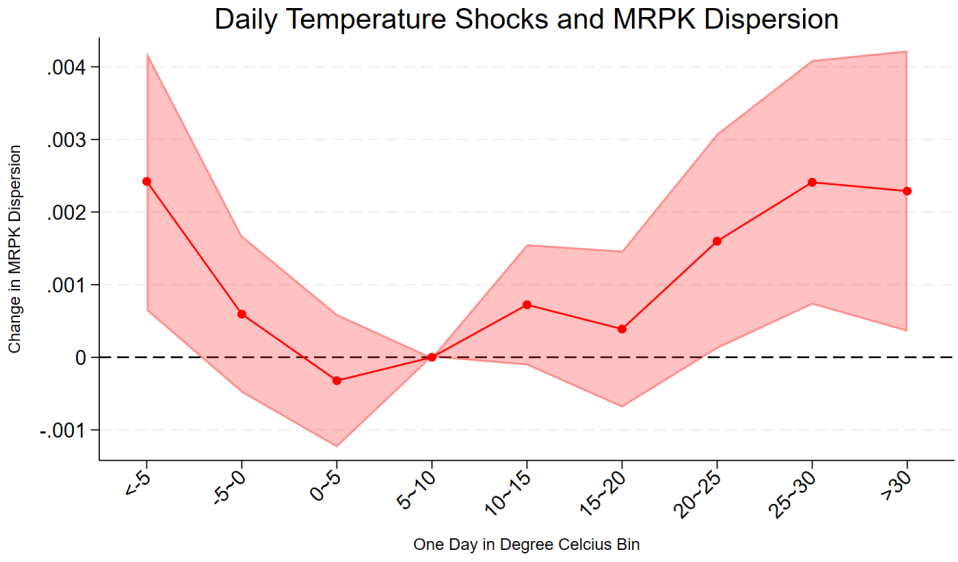


Table C.3: Regression Coefficients: NUTS 2–SIC 1

	(1) σ^2_mrpk	(2) σ^2_mrpk	(3) Implied TFP Loss
< -5°C	0.0024**	0.0028**	0.0869**
	(0.0011)	(0.0011)	(0.0388)
$-5 \sim 0$°C	0.0006	0.0006	0.0213
	(0.0007)	(0.0006)	(0.0236)
$0 \sim 5$°C	-0.0003	-0.0004	-0.0115
	(0.0006)	(0.0006)	(0.0200)
$10 \sim 15$°C	0.0007	0.0008	0.0259
	(0.0005)	(0.0005)	(0.0182)
$15 \sim 20$°C	0.0004	0.0002	0.0140
	(0.0007)	(0.0007)	(0.0235)
$20 \sim 25$°C	0.0016*	0.0014	0.0573*
	(0.0009)	(0.0009)	(0.0323)
$25 \sim 30$°C	0.0024**	0.0021**	0.0864**
	(0.0010)	(0.0010)	(0.0367)
> 30°C	0.0023*	0.0020*	0.0821*
	(0.0012)	(0.0012)	(0.0421)
Region-Sector FE	Yes	Yes	Yes
Country-Year FE	Yes	No	Yes
Country-Sector-Year FE	No	Yes	No
Observations	47,731	46,942	47,731
R^2	0.901	0.936	0.901

Notes: Standard errors in parentheses.

NUTS 3–SIC 2. Figure C.8 redefines the unit as NUTS 3 $\times$ SIC 2-digit, recomputes Var(MRPK), and re-estimates the baseline under alternative fixed-effect structures. The results are quantitatively similar to the baseline. Table C.4 reports the underlying coefficients, and Figure C.9 shows the corresponding heterogeneous-effect estimates across hot and cold regions.

Figure C.8: Results Using NUTS 3–SIC 2 as the Region–Sector Unit, Under Alternative Fixed Effects

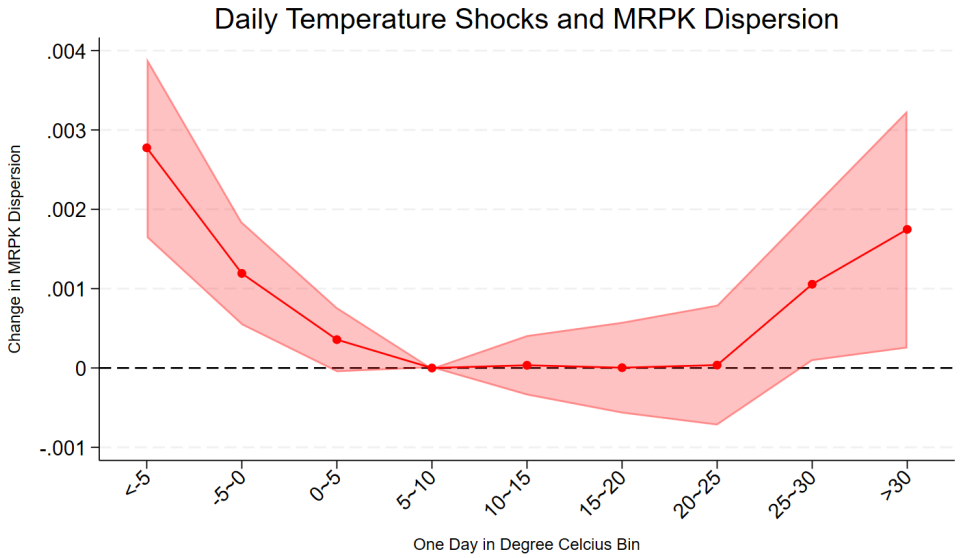
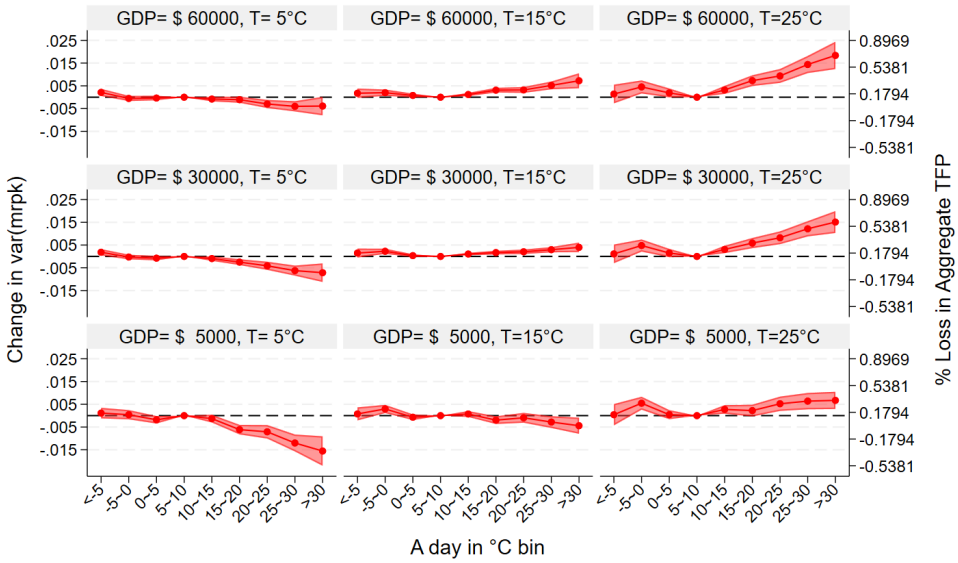


Table C.4: Regression Coefficients: NUTS 3–SIC 2

	(1) σ^2_mrpk	(2) σ^2_mrpk	(3) Implied TFP Loss
< -5°C	0.0028***	0.0029***	0.0996***
	(0.0007)	(0.0007)	(0.0247)
$-5 \sim 0$°C	0.0012***	0.0012***	0.0428***
	(0.0004)	(0.0004)	(0.0142)
$0 \sim 5$°C	0.0004	0.0003	0.0128
	(0.0002)	(0.0002)	(0.0089)
$10 \sim 15$°C	0.0000	0.0001	0.0012
	(0.0002)	(0.0002)	(0.0083)
$15 \sim 20$°C	0.0000	0.0001	0.0001
	(0.0004)	(0.0004)	(0.0126)
$20 \sim 25$°C	0.0000	0.0002	0.0013
	(0.0005)	(0.0005)	(0.0166)
$25 \sim 30$°C	0.0011*	0.0012**	0.0379*
	(0.0006)	(0.0006)	(0.0211)
> 30°C	0.0017*	0.0019**	0.0627*
	(0.0009)	(0.0010)	(0.0327)
Region-Sector FE	Yes	Yes	Yes
Country-Year FE	Yes	No	Yes
Country-Sector-Year FE	No	Yes	No
Observations	317,542	312,377	317,542
R^2	0.840	0.882	0.840

Notes: Standard errors in parentheses.

Figure C.9: Heterogeneous Temperature Effects on MRPK Dispersion Across Regions: NUTS 3–SIC 2



NUTS 3–SIC 3. For European countries, we further disaggregate to NUTS 3 $\times$ SIC 3-digit. Figure C.10 shows that the main patterns are preserved. Table C.5 reports the underlying coefficients, and Figure C.11 shows the corresponding heterogeneous-effect estimates across hot and cold regions.

Figure C.10: Results Using NUTS 3–SIC 3 as the Region–Sector Unit, Under Alternative Fixed Effects

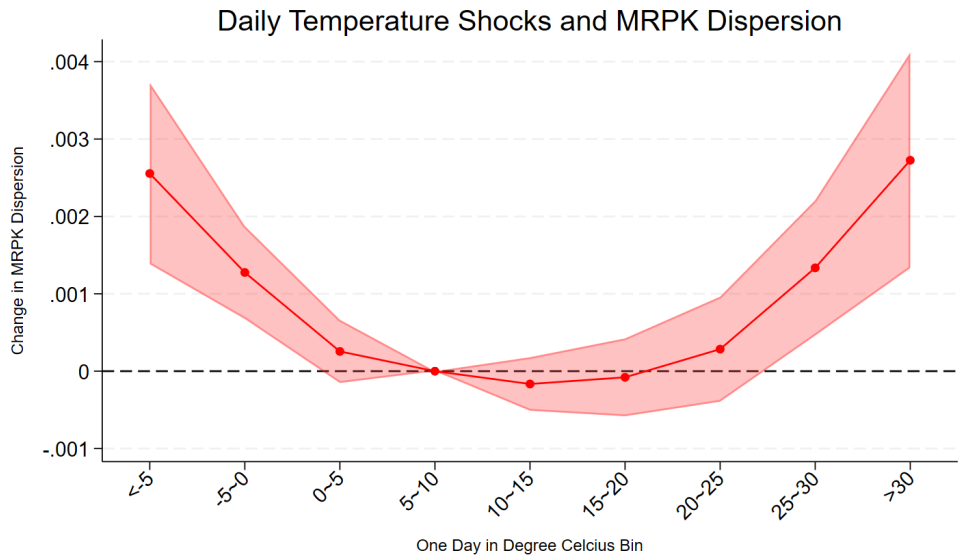
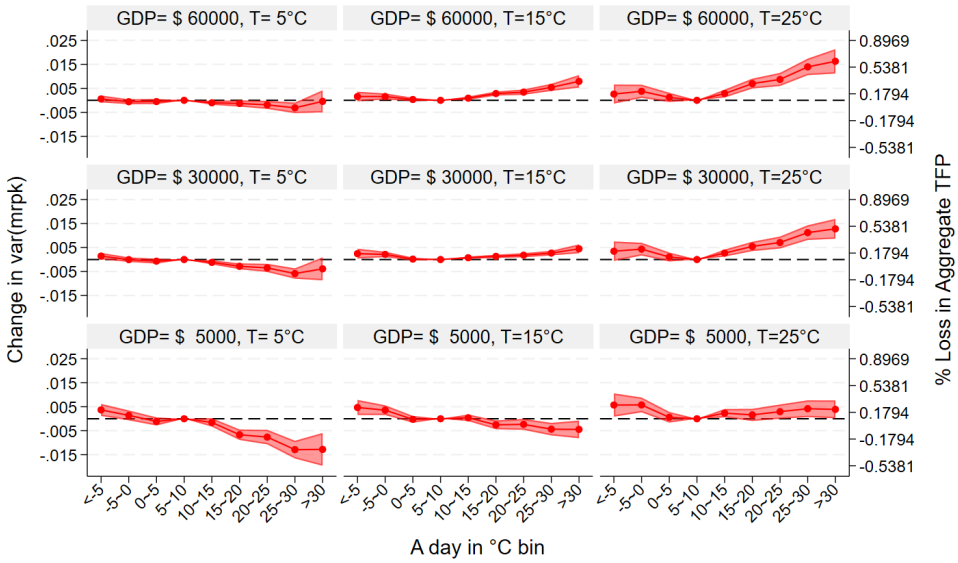


Table C.5: Regression Coefficients: NUTS 3–SIC 3

	(1) σ^2_mrpk	(2) σ^2_mrpk	(3) Implied TFP Loss
< -5°C	0.0026***	0.0023***	0.0917***
	(0.0007)	(0.0007)	(0.0255)
$-5 \sim 0$°C	0.0013***	0.0013***	0.0458***
	(0.0004)	(0.0004)	(0.0131)
$0 \sim 5$°C	0.0003	0.0003	0.0092
	(0.0002)	(0.0002)	(0.0089)
$10 \sim 15$°C	-0.0002	-0.0001	-0.0059
	(0.0002)	(0.0002)	(0.0075)
$15 \sim 20$°C	-0.0001	0.0001	-0.0028
	(0.0003)	(0.0003)	(0.0110)
$20 \sim 25$°C	0.0003	0.0004	0.0102
	(0.0004)	(0.0004)	(0.0148)
$25 \sim 30$°C	0.0013**	0.0013**	0.0479**
	(0.0005)	(0.0006)	(0.0191)
> 30°C	0.0027***	0.0031***	0.0978***
	(0.0008)	(0.0010)	(0.0304)
Region-Sector FE	Yes	Yes	Yes
Country-Year FE	Yes	No	Yes
Country-Sector-Year FE	No	Yes	No
Observations	403,748	388,912	403,748
R^2	0.837	0.884	0.837

Notes: Standard errors in parentheses.

Figure C.11: Heterogeneous Temperature Effects on MRPK Dispersion Across Regions: NUTS 3–SIC 3



C.6 Alternative Winsorization

To verify that results are not driven by extreme outliers in the revenue-to-capital distribution, we winsorize MRPK at the 3% level (rather than the baseline 0.1%) before computing region–sector–year variances. Figure C.12 shows that the results are robust to this substantially more aggressive trimming. Table C.6 reports the regression coefficients, and Figure C.13 confirms that the heterogeneous-effect pattern across hot and cold regions parallels the baseline.

Figure C.12: Estimated Effects of Daily Temperature Shocks on Annual MRPK Dispersion: 3% Winsorized MRPK

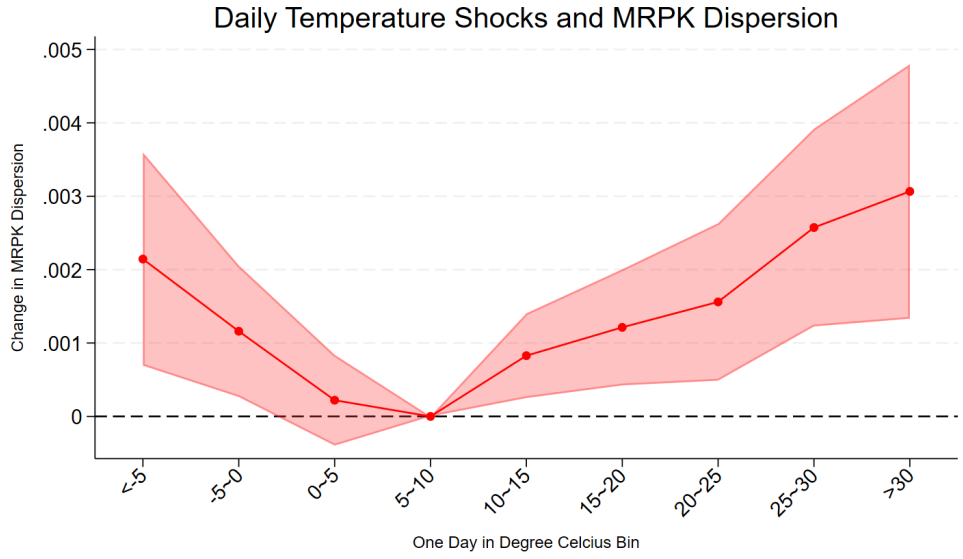
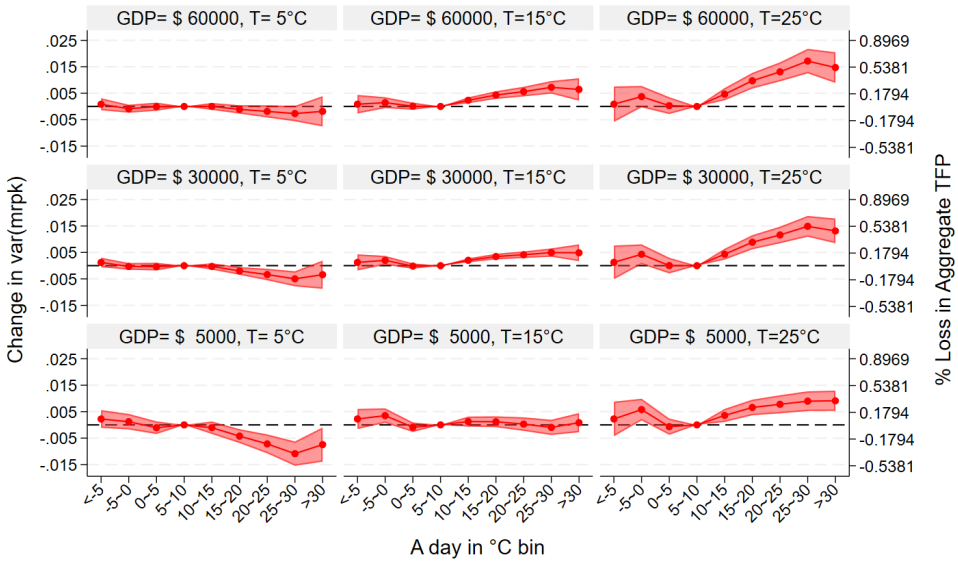


Table C.6: Regression Coefficients: 3% Winsorized MRPK

	(1) σ^2_mrpk	(2) σ^2_mrpk	(3) Implied TFP Loss
< -5°C	0.0021**	0.0021**	0.0769**
	(0.0009)	(0.0009)	(0.0316)
$-5 \sim 0$°C	0.0012**	0.0010*	0.0416**
	(0.0005)	(0.0005)	(0.0195)
$0 \sim 5$°C	0.0002	0.0000	0.0079
	(0.0004)	(0.0004)	(0.0135)
$10 \sim 15$°C	0.0008**	0.0008**	0.0297**
	(0.0004)	(0.0004)	(0.0126)
$15 \sim 20$°C	0.0012**	0.0012**	0.0435**
	(0.0005)	(0.0005)	(0.0172)
$20 \sim 25$°C	0.0016**	0.0016**	0.0560**
	(0.0007)	(0.0006)	(0.0234)
$25 \sim 30$°C	0.0026***	0.0026***	0.0923***
	(0.0008)	(0.0008)	(0.0294)
> 30°C	0.0031***	0.0031***	0.1100***
	(0.0011)	(0.0010)	(0.0378)
Region-Sector FE	Yes	Yes	Yes
Country-Year FE	Yes	No	Yes
Country-Sector-Year FE	No	Yes	No
Observations	122,055	121,501	122,055
R^2	0.871	0.901	0.871

Notes: Standard errors in parentheses.

Figure C.13: Heterogeneous Temperature Effects on MRPK Dispersion Across Regions: 3% Winsorized MRPK



C.7 Alternative Production-Function Specifications

The baseline analysis measures MRPK using the revenue-to-capital ratio under Cobb–Douglas production ($\sigma_{\text{CES}} = 1$). This subsection verifies that results hold under CES production, which relaxes the unit-elasticity restriction and allows for different mappings from factor ratios to capital wedge dispersion.

C.7.1 CES with Hicks-Neutral Productivity

Consider a two-input CES production function,

$$Y_{it} = A_{it}\left[\gamma K_{it}^{\frac{\sigma-1}{\sigma}} + (1-\gamma)B_{it}L_{it}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

where A_{it} is Hicks-neutral productivity with firm-specific temperature sensitivity, $A_{it} = \tilde{A}_{it}e^{a_{it}T_{rt}}$, and $B_{it} = \tilde{B}_{it}e^{b_{nt}T_{rt}}$ is labor-augmenting productivity with a common regional-sectoral temperature component. When $\sigma \to 1$, this reduces to Cobb–Douglas.

Under monopolistic competition with firm-specific capital wedges τ_{it}^k, the ratio of the capital and labor first-order conditions yields

$$\frac{1}{\sigma}\log\frac{L_{it}}{K_{it}} = \log(1+\tau_{it}^k) + \frac{\sigma-1}{\sigma}\log B_{it} + \log r - \log w - \log\frac{\gamma}{1-\gamma}.$$

Within a region–sector–year, factor prices are common, so the within-cell variance reduces to

$$\frac{1}{\sigma^2}\text{Var}\left(\log\frac{L_{it}}{K_{it}}\right) = \text{Var}\left(\log(1+\tau_{it}^k)\right) + \left(\frac{\sigma-1}{\sigma}\right)^2\text{Var}\left(\log\tilde{B}_{it}\right) + 2\frac{\sigma-1}{\sigma^2}\text{Cov}\left(\log(1+\tau_{it}^k), \log\tilde{B}_{it}\right).$$

Under the assumption that the cross-sectional distributions of $\tilde{B}_{it}$ and its covariance with wedges

do not respond to temperature,[35] the temperature response of capital wedge dispersion is identified as

$$\frac{d\,\mathrm{Var}\big(\log(1+\tau_{it}^{k})\big)}{dT} = \frac{1}{\sigma^2}\,\frac{d\,\mathrm{Var}\Big(\log\frac{L_{it}}{K_{it}}\Big)}{dT}.$$

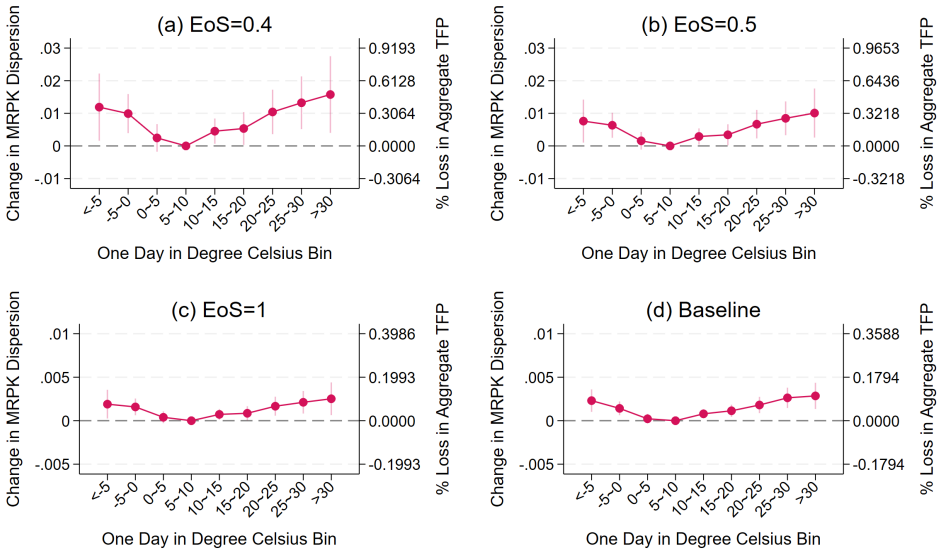
TFP losses. Following Baqaee and Farhi (2019), under CES production over capital and labor, the TFP loss implied by changes in MRPK dispersion is

$$\frac{(\sigma_n-1)\alpha_K^2+\alpha_K^2+\sigma\cdot\alpha_K(1-\alpha_K)}{2}\,\Delta\sigma^2_{mrpk},$$

where σ_n is the demand elasticity and σ is the capital–labor elasticity of substitution.

Results. Figure C.14 reports the estimation under $\sigma \in \{0.4,\ 0.5,\ 1.0\}$, where $\sigma = 0.4$ corresponds to the IV estimates of Oberfield and Raval (2021) and $\sigma = 0.5$ to the average across specifications in the micro–macro literature. The U-shaped relationship between temperature and MRPK dispersion is preserved across all values of σ.

Figure C.14: Estimated Impact of Daily Temperature Shocks on Annual MRPK Dispersion: CES with Hicks-Neutral Productivity



Notes: Panels (a)–(c) measure MRPK dispersion via $\mathrm{Var}(\log(L_{it}/K_{it}))/\sigma^2$ under $\sigma = 0.4$, 0.5, and 1.0 (Cobb–Douglas), respectively. Panel (d) reproduces the baseline using $\mathrm{Var}(\log(PY_{it}/K_{it}))$.

C.7.2 CES with Labor-Augmenting Productivity and Materials

To allow for labor-augmenting productivity while still identifying capital wedge dispersion, we introduce a third factor—materials M_{it}—into the CES aggregator:

$$Y_{it} = A_{it}\left[K_{it}^{\frac{\sigma-1}{\sigma}} + (B_{it}L_{it})^{\frac{\sigma-1}{\sigma}} + M_{it}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}.$$

35. This requires that the firm-specific component of labor-augmenting productivity is orthogonal to temperature, and that temperature does not systematically shift the covariance between wedges and labor productivity within a cell.

The capital and materials first-order conditions do not involve B_{it}, so their ratio directly identifies capital wedge dispersion:

$$\log(1+\tau_{it}^k) = \frac{1}{\sigma}\log\frac{M_{it}}{K_{it}} + \log\frac{p^m}{r},$$

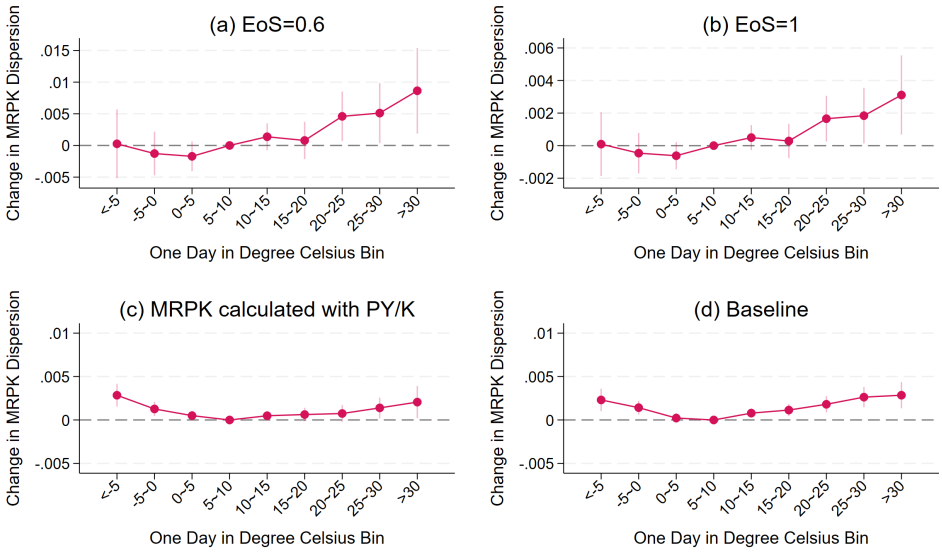
and therefore

$$\mathrm{Var}\Big(\log(1+\tau_{it}^k)\Big) = \frac{1}{\sigma^2}\,\mathrm{Var}\left(\log\frac{M_{it}}{K_{it}}\right),$$

since material prices and rental rates are common within a region–sector–year. This approach eliminates labor-augmenting productivity from the identification without additional assumptions.

Results. Figure C.15 reports the results. The U-shaped pattern is again robust, confirming that the main findings do not depend on the Cobb–Douglas restriction or on the treatment of labor-augmenting productivity.

Figure C.15: Estimated Impact of Daily Temperature Shocks on Annual MRPK Dispersion: CES with Materials



Notes: Panels (a)–(c) measure MRPK dispersion via $\mathrm{Var}(\log(M_{it}/K_{it}))/\sigma^2$ under $\sigma = 0.4$, 0.5, and 1.0, respectively. Panel (d) reproduces the baseline using $\mathrm{Var}(\log(PY_{it}/K_{it}))$.

D Technology Estimation Details for Section 4

This appendix derives the cross-cell object identified by the weighted-PPML technology estimator and shows that it is directly comparable to the cross-cell object identified by the misallocation OLS regression. We first establish the misallocation benchmark: Proposition 6 characterizes the fixed-effects OLS coefficient as a projection-weighted average of cell-specific slopes. We then define the technology object as a cell-level CES aggregate elasticity, explain why firm-level OLS targets an equal-weight average of firm elasticities rather than the aggregate, and show why PPML is the natural estimator for a conditional mean in levels. Proposition 7 shows that regular PPML solves within-cell aggregation but still mass-weights across cells. Proposition 8 shows that inverse-mass

weighting recovers the comparable cross-cell average. The final subsection presents the feasible raw plug-in estimator, Proposition 9 that establishes its consistency, the four-step implementation used in the code, and the algebraic equivalence between the implemented Stata specification and the original weighted form.

D.1 Misallocation OLS and technology WPPML: comparable average effects

This subsection establishes the cross-cell benchmark on the misallocation side. The baseline misallocation regression in Equation (10) is estimated by fixed-effects OLS at the region-sector-year level.[36] Once we allow temperature responses to differ across cells, the relevant question is not whether OLS recovers a single slope, but which average of the heterogeneous cell-level slopes it recovers.

Let $\tilde{\mathbf{T}}_{rt}$ denote the vector of non-omitted temperature bins after partialling out the included fixed effects and controls, and let $\widetilde{\sigma^2_{mrpk,(n,t)}}$ denote the corresponding residualized misallocation outcome in cell (n,t). Then

$$\widetilde{\sigma^2_{mrpk,(n,t)}} = \tilde{\mathbf{T}}'_{rt}\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)} + u_{nt}, \qquad \boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)} := (\tau^b_{\sigma^2_{mrpk},(n,t)})_b,$$

where $\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}$ is the cell-specific slope vector and $\mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}u_{nt}] = 0$.

Proposition 6 characterizes which average of these heterogeneous slopes the OLS coefficient recovers.

Proposition 6 (What Equation (10) averages) *Let*

$$Q := \mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}].$$

If Q is finite and nonsingular, the pseudo-true fixed-effects OLS coefficient vector is

$$\boldsymbol{\lambda}^*_{\sigma^2_{mrpk}} = Q^{-1}\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right].$$

Thus, in general, the misallocation coefficient is a projection-weighted average of cell-specific slopes. If, in addition,

$$\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\,\middle|\,\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right] = Q, \tag{38}$$

then

$$\boldsymbol{\lambda}^*_{\sigma^2_{mrpk}} = \mathbb{E}_{n,t}[\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}].$$

Proof. The population normal equations imply

$$Q\boldsymbol{\lambda}^*_{\sigma^2_{mrpk}} = \mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}\widetilde{\sigma^2_{mrpk,(n,t)}}].$$

36. Throughout this appendix we follow the notation of Section 2 and index region-sector-year cells by (n,t), where $n = (r,s)$ is a region-sector and t is a year. Firm observations in cell (n,t) are indexed by nit.

Substituting $\widetilde{\sigma^2_{mrpk,(n,t)}} = \tilde{\mathbf{T}}'_{rt}\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)} + u_{nt}$ and using $\mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}u_{nt}] = 0$ gives

$$Q\boldsymbol{\lambda}^*_{\sigma^2_{mrpk}} = \mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right],$$

which proves the first result. Under Equation (38),

$$\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right] = \mathbb{E}_{n,t}\left[\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\,\middle|\,\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right]\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}\right] = Q\,\mathbb{E}_{n,t}[\boldsymbol{\tau}_{\sigma^2_{mrpk},(n,t)}],$$

and premultiplying by Q^{-1} yields the claim. ■

Proposition 6 is the interpretation used in the main text. The baseline misallocation regression should therefore be read as a projection-weighted average of cell-specific misallocation responses. Under the mean-independence condition in Equation (38), that projection average coincides with the arithmetic average effect across region-sector-year cells for each coefficient $\lambda^b_{\sigma^2_{mrpk}}$.

Equation (38) is a sufficient condition for the simple-average interpretation, not the only possible one. In plain words, it says that after partialling out the included fixed effects and controls, cells with different temperature sensitivities are not systematically associated with different residual temperature variation. If that condition fails, OLS still has a clean interpretation, but it is the projection-weighted average in Proposition 6 rather than the literal arithmetic mean.

Thus, the misallocation coefficient identifies a well-defined cross-cell average; the remaining steps of the appendix ask which technology-side estimator targets the directly comparable object.

D.2 Average technology effects with weighted PPML

This subsection defines the technology object that we want to estimate. The key distinction is between an equal-weight average firm elasticity and a cell-level *aggregate* technology elasticity. The latter is the object relevant for aggregate TFP in the accounting framework of Section 2, and it is the object that the weighted-PPML estimator is built to recover.

Following the notation introduced in Section 4.3, the cell-level CES technology mass is

$$M_{nt}(\tilde{\mathbf{T}}_{rt}) := \mathbb{E}_i\left[\mathrm{TFP}_{ni}(\tilde{\mathbf{T}}_{rt},\cdot)^{\sigma_n-1}\,\middle|\,n,t\right],$$

where the expectation is taken over firms within cell (n,t). For each temperature bin b, the corresponding cell-level aggregate technology effect is the bth component of

$$\frac{1}{\sigma_n-1}\frac{d\log M_{nt}(\tilde{\mathbf{T}}_{rt})}{d\tilde{\mathbf{T}}_{rt}}.$$

This is the technology object relevant for aggregate TFP.

The rest of this subsection follows the five steps of the technology-side argument. We first show why firm-level OLS does not estimate this object. We then show why PPML is aligned with the CES aggregate and where the transformed outcome and offset come from. We then state the oracle target of regular PPML, explain why the across-cell weighting needs to be corrected, and state the oracle target of weighted PPML.

D.2.1 Why firm-level OLS is not the right technology estimator

Firm-level OLS gives an equal-weight average of firm elasticities, which is not the technology object in the accounting framework. To make that distinction explicit, define the transformed firm-level technology component

$$y^E_{nit} := \text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot)^{\sigma_n - 1}, \qquad \boldsymbol{\psi}_{nit} := \frac{\partial \log y^E_{nit}}{\partial \tilde{\mathbf{T}}_{rt}} = (\psi^b_{nit})_b.$$

Then the cell-level aggregate technology elasticity can be written as

$$\frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \tilde{\mathbf{T}}_{rt}} = \sum_{i \in (n,t)} s^E_{nit}\, \boldsymbol{\psi}_{nit}, \qquad s^E_{nit} := \frac{y^E_{nit}}{\sum_{j \in (n,t)} y^E_{njt}}.$$

Firms that matter more for the CES technology mass therefore receive larger weight in the derivative of the aggregate.

Now consider the natural but incorrect alternative of estimating a firm-level fixed-effects OLS regression of $\log y^E_{nit}$ on the temperature bins. Let

$$\widetilde{\log y^E_{nit}} = \tilde{\mathbf{T}}'_{rt} \boldsymbol{\phi}_{nit} + v_{nit}, \qquad \boldsymbol{\phi}_{nit} := (\phi^b_{nit})_b,$$

where the tilde denotes residualization with respect to the included fixed effects and controls and $\mathbb{E}_{n,i,t}[\tilde{\mathbf{T}}_{rt} v_{nit}] = 0$. If

$$Q_f := \mathbb{E}_{n,i,t}[\tilde{\mathbf{T}}_{rt} \tilde{\mathbf{T}}'_{rt}]$$

is finite and nonsingular, the pseudo-true firm-level OLS coefficient is

$$\boldsymbol{\beta}^{firm}_{OLS} = Q_f^{-1} \mathbb{E}_{n,i,t}\left[\tilde{\mathbf{T}}_{rt} \tilde{\mathbf{T}}'_{rt} \boldsymbol{\phi}_{nit}\right],$$

and under the corresponding mean-independence condition,

$$\boldsymbol{\beta}^{firm}_{OLS} = \mathbb{E}_{n,i,t}[\boldsymbol{\phi}_{nit}].$$

Firm-level OLS therefore recovers an average firm elasticity across firm observations: a large and a small firm each contribute one observation. The within-cell gap between this equal-weight average and the aggregate technology object is

$$\frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \tilde{\mathbf{T}}_{rt}} - \frac{1}{N_{nt}} \sum_{i \in (n,t)} \boldsymbol{\psi}_{nit} = \sum_{i \in (n,t)} \left(s^E_{nit} - \frac{1}{N_{nt}} \right) \boldsymbol{\psi}_{nit}.$$

Unless firm economic importance is unrelated to temperature sensitivity, the equal-weight average firm elasticity and the CES-share-weighted aggregate elasticity are different objects.

Thus, the problem with firm-level OLS is not only statistical but conceptual: it estimates the wrong aggregation of heterogeneous firm responses.

D.2.2 Why PPML targets the CES technology object

PPML is aligned with the CES technology aggregate because the aggregate enters naturally as a conditional mean in levels, and PPML is the natural estimator of conditional means in levels. Before formalizing the estimator, we derive where the firm-level PPML outcome and the cell offset come from in the accounting framework.

Consider the efficient benchmark in which all wedges are shut down, so that $\tau_{ni}^{Y} = \tau_{ni}^{K} = \tau_{ni}^{L} = 0$. Then firm production satisfies

$$Y_{nit} = A_{nit} K_{nit}^{\alpha_{Kn}} L_{nit}^{\alpha_{Ln}},$$

and the CES demand system implies

$$Y_{nit} = B_{nit}\, Y_{nt} \left(\frac{P_{nit}}{P_{nt}} \right)^{-\sigma_n},$$

where P_{nt} and Y_{nt} are the cell-level price index and output. The paper's transformed productivity measure is

$$\widetilde{\text{TFP}}_{nit} := \frac{(P_{nit} Y_{nit})^{\sigma_n/(\sigma_n-1)}}{K_{nit}^{\alpha_{Kn}} L_{nit}^{\alpha_{Ln}}}.$$

Since $K_{nit}^{\alpha_{Kn}} L_{nit}^{\alpha_{Ln}} = Y_{nit}/A_{nit}$, this rewrites as

$$\widetilde{\text{TFP}}_{nit} = A_{nit} P_{nit}^{\sigma_n/(\sigma_n-1)} Y_{nit}^{1/(\sigma_n-1)}.$$

Substituting the CES demand expression for Y_{nit} gives

$$\widetilde{\text{TFP}}_{nit} = \left(P_{nt}^{\sigma_n} Y_{nt}\right)^{1/(\sigma_n-1)} B_{nit}^{1/(\sigma_n-1)} A_{nit},$$

so that

$$\widetilde{\text{TFP}}_{nit}^{\sigma_n-1} = P_{nt}^{\sigma_n} Y_{nt}\, B_{nit} A_{nit}^{\sigma_n-1}.$$

Defining the efficient firm-level technology term

$$y_{nit}^{E} := \text{TFP}_{ni}(\tilde{\mathbf{T}}_{rt}, \cdot)^{\sigma_n-1} = B_{nit} A_{nit}^{\sigma_n-1},$$

we obtain the multiplicative decomposition

$$y_{nit} := \widetilde{\text{TFP}}_{nit}^{\sigma_n-1} = e^{o_{nt}}\, y_{nit}^{E}, \qquad o_{nt} := \log(P_{nt}^{\sigma_n} Y_{nt})\,.$$

Averaging within a cell then yields

$$\bar{y}_{nt} = e^{o_{nt}}\, \mathbb{E}_i[y_{nit}^{E} \mid n, t] = e^{o_{nt}} M_{nt}(\tilde{\mathbf{T}}_{rt}).$$

This is the algebraic reason the PPML regression uses $\widetilde{\text{TFP}}_{nit}^{\sigma_n-1}$ as the dependent variable and a cell-level offset constructed from aggregate prices and quantities. In the empirical implementation, the corresponding price normalization is absorbed by the revenue-based offset together with the country-sector-year fixed effects.

With the outcome and offset in place, PPML estimates the conditional mean of y_{nit} directly in

levels. The score equations match fitted sums to observed sums within each conditioning cell, so the estimator lands on $\bar{y}_{nt} = e^{o_{nt}} M_{nt}(\tilde{\mathbf{T}}_{rt})$ rather than on the equal-weighted average that OLS of $\log y^E_{nit}$ recovers. Thus, PPML solves the within-cell aggregation problem that made firm-level OLS inappropriate. Whether PPML also delivers the correct across-cell average is the question addressed next.

D.2.3 Oracle target of regular PPML

Regular PPML, without any additional cross-cell weighting, has an oracle score

$$\Psi^{ppml}(\boldsymbol{\beta}) := \mathbb{E}_{n,t}\left[e^{o_{nt}} \tilde{\mathbf{T}}_{rt} \left(M_{nt}(\tilde{\mathbf{T}}_{rt}) - M^0_{nt} e^{\tilde{\mathbf{T}}'_{rt}\boldsymbol{\beta}}\right)\right], \tag{39}$$

where $M^0_{nt} := M_{nt}(\tilde{\mathbf{T}}^0_{rt})$ is the cell's baseline CES mass. Proposition 7 characterizes what this score identifies.

Proposition 7 (Oracle target of regular PPML) *Suppose that:*

1. $Q := \mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}]$ *is finite and nonsingular;*
2. *the baseline masses* M^0_{nt} *are bounded away from zero and infinity, and the offsets* $e^{o_{nt}}$ *are bounded above;*
3. *in a local neighborhood of* $\tilde{\mathbf{T}}^0_{rt}$*,*

$$\log M_{nt}(\tilde{\mathbf{T}}_{rt}) = \log M^0_{nt} + \tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt} + r_{nt}, \qquad \mathbf{h}_{nt} := \left(\frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \mathrm{Tbin}^b_{rt}}\right)_b,$$

with remainder satisfying $\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|\,|r_{nt}|] = o(\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|^2])$.

Let $\boldsymbol{\beta}^{ppml}$ *solve* $\Psi^{ppml}(\boldsymbol{\beta}) = \mathbf{0}$*, and define*

$$Q_{ppml} := \mathbb{E}_{n,t}\left[e^{o_{nt}} M^0_{nt}\, \tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\right].$$

Then

$$\boldsymbol{\beta}^{ppml} = Q^{-1}_{ppml}\mathbb{E}_{n,t}\left[e^{o_{nt}} M^0_{nt}\, \tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt}\right] + o(1).$$

If, in addition,

$$\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\,\middle|\,\mathbf{h}_{nt}, e^{o_{nt}} M^0_{nt}\right] = Q,$$

then

$$\boldsymbol{\beta}^{ppml} = \frac{\mathbb{E}_{n,t}[e^{o_{nt}} M^0_{nt}\,\mathbf{h}_{nt}]}{\mathbb{E}_{n,t}[e^{o_{nt}} M^0_{nt}]} + o(1),$$

with the division taken componentwise in the scalar denominator.

Proof. Evaluating Equation (39) at $\boldsymbol{\beta}^{ppml}$ and substituting the local expansion $\log M_{nt}(\tilde{\mathbf{T}}_{rt}) = \log M^0_{nt} + \tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt} + r_{nt}$ yields

$$\mathbf{0} = \mathbb{E}_{n,t}\left[e^{o_{nt}} M^0_{nt}\, \tilde{\mathbf{T}}_{rt}\left(e^{\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt} + r_{nt}} - e^{\tilde{\mathbf{T}}'_{rt}\boldsymbol{\beta}^{ppml}}\right)\right].$$

A first-order expansion of the exponentials around the baseline gives

$$\mathbf{0} = \mathbb{E}_{n,t}\Big[e^{o_{nt}} M^0_{nt}\, \tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}(\mathbf{h}_{nt} - \boldsymbol{\beta}^{ppml})\Big] + \mathbb{E}_{n,t}\Big[e^{o_{nt}} M^0_{nt}\, \tilde{\mathbf{T}}_{rt} r_{nt}\Big] + o\Big(\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|^2]\Big)\,.$$

By the remainder condition in hypothesis (3) and boundedness of $e^{o_{nt}} M^0_{nt}$ in hypothesis (2), the middle term is of smaller order, which yields the first expression. The second follows immediately from the additional conditional mean-independence assumption. ■

Proposition 7 shows that PPML recovers each cell's aggregate technology elasticity but averages those elasticities across cells using baseline outcome-mass weights $e^{o_{nt}} M^0_{nt} = \bar{y}^0_{nt}$. In a sample spanning 32 economies with very different scales, such weighting would give more influence to larger economies and sectors simply because their baseline level is larger. Thus, regular PPML has solved the within-cell aggregation problem but not the across-cell weighting problem.

D.2.4 Why weighted PPML is needed across cells

To target an average effect across cells rather than an outcome-mass-weighted average across cells, we divide out the baseline mass. Because the offset absorbs the common cell-scale term $e^{o_{nt}}$, the corresponding oracle moment can be written as

$$\Psi^o(\boldsymbol{\beta}) := \mathbb{E}_{n,t}\left[\frac{\tilde{\mathbf{T}}_{rt}}{M^0_{nt}}\left(M_{nt}(\tilde{\mathbf{T}}_{rt}) - M^0_{nt} e^{\tilde{\mathbf{T}}'_{rt}\boldsymbol{\beta}}\right)\right], \tag{40}$$

where $\tilde{\mathbf{T}}_{rt}$ denotes the temperature-bin vector after removing the included fixed effects and controls.

Equation (40) is the formal statement of the fix. The numerator $\tilde{\mathbf{T}}_{rt}$ keeps the usual score-type comparison between residual temperature variation and residual outcomes, while the denominator M^0_{nt} renormalizes each cell by its baseline technology mass. This is what makes the technology regression target an average *across cells* rather than an outcome-mass-weighted average across cells.

D.2.5 Oracle target of weighted PPML

Proposition 8 characterizes what the weighted-PPML oracle score identifies.

Proposition 8 (Oracle target of weighted PPML) *Suppose that:*

1. $Q := \mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}]$ *is finite and nonsingular;*
2. *the baseline masses* M^0_{nt} *are bounded away from zero and infinity;*
3. *in a local neighborhood of* $\tilde{\mathbf{T}}^0_{rt}$,

$$\log M_{nt}(\tilde{\mathbf{T}}_{rt}) = \log M^0_{nt} + \tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt} + r_{nt}, \qquad \mathbf{h}_{nt} := \left(\frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \mathrm{Tbin}^b_{rt}}\right)_b, \tag{41}$$

with remainder satisfying $\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|\,|r_{nt}|] = o(\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|^2])$.

Let $\boldsymbol{\beta}^*$ solve $\Psi^o(\boldsymbol{\beta}) = \mathbf{0}$. Then

$$\boldsymbol{\beta}^* = Q^{-1}\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt}\right] + o(1).$$

If, in addition,

$$\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\,\middle|\,\mathbf{h}_{nt}\right] = Q, \tag{42}$$

then

$$\boldsymbol{\beta}^* = \mathbb{E}_{n,t}[\mathbf{h}_{nt}] + o(1), \qquad \frac{\beta_b^*}{\sigma_n - 1} = \mathbb{E}_{n,t}\left[\frac{1}{\sigma_n - 1}\frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \mathrm{Tbin}_{rt}^b}\right] + o(1).$$

Proof. Evaluating Equation (40) at $\boldsymbol{\beta}^*$ and substituting the local expansion in Equation (41) yields

$$\mathbf{0} = \mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\left(e^{\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt}+r_{nt}} - e^{\tilde{\mathbf{T}}'_{rt}\boldsymbol{\beta}^*}\right)\right].$$

A first-order expansion of the exponentials around the baseline gives

$$\mathbf{0} = \mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}(\mathbf{h}_{nt} - \boldsymbol{\beta}^*)\right] + \mathbb{E}_{n,t}[\tilde{\mathbf{T}}_{rt}r_{nt}] + o\Big(\mathbb{E}_{n,t}[\|\tilde{\mathbf{T}}_{rt}\|^2]\Big).$$

By the remainder condition, the middle term is of smaller order, so

$$Q\,\boldsymbol{\beta}^* = \mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt}\right] + o(1),$$

which proves the first statement. Under Equation (42),

$$\mathbb{E}_{n,t}\left[\tilde{\mathbf{T}}_{rt}\tilde{\mathbf{T}}'_{rt}\mathbf{h}_{nt}\right] = Q\,\mathbb{E}_{n,t}[\mathbf{h}_{nt}],$$

so $\boldsymbol{\beta}^* = \mathbb{E}_{n,t}[\mathbf{h}_{nt}] + o(1)$. Dividing by $\sigma_n - 1$ gives the result. ■

Proposition 8 is the key comparability result. The weighted-PPML coefficient targets the same kind of cross-cell average as the misallocation regression: a projection average in general, and the arithmetic average cell-level effect under the corresponding mean-independence condition. In the main text, this is the justification for reading $\hat{\beta}_b/(\sigma_n - 1)$ against $\hat{\lambda}^b_{\sigma^2_{mrpk}}$ directly.

Each assumption has a transparent role. Nonsingularity of Q ensures that the residual temperature bins contain enough independent variation to identify the slope vector. The bounds on M^0_{nt} rule out cells with vanishing or explosive baseline mass, because inverse-mass weighting would otherwise become unstable. The local expansion in Equation (41) says that around the observed baseline, the log CES technology mass can be approximated by a first-order response plus a remainder, so the proposition is a local-comparative-static result rather than a claim that the full climate response is globally linear.

The main caveat is again Equation (42). If cells with steeper technology responses also systematically experience different residual temperature variation, then the coefficient should be read as a projection-weighted average rather than the literal arithmetic mean. We therefore do not need this condition for the core comparability result; we need it only if we want to interpret each coefficient as the simple equal-weight average effect across cells.

Thus, weighted PPML and misallocation OLS target directly comparable cross-cell averages under a single, transparent mean-independence condition on each side.

D.3 Feasible implementation

This subsection turns the oracle estimator into the feasible estimator used in the code. The oracle baseline mass M_{nt}^0 is not observed, so we replace it with a transparent sample analogue, prove that the replacement does not alter the estimand, and then describe the implementation in the empirical code.

D.3.1 Feasible raw plug-in weights

The oracle weight depends on an unobserved population object. In the main-text baseline, we replace it with a raw plug-in estimate based on the cell's baseline outcome mass:

$$\widehat{\bar{y}}_{nt,raw}^0 := \frac{1}{N_{nt}} \sum_{i \in (n,t)} y_{nit}, \qquad \widehat{a}_{nit,raw} := \frac{1}{N_{nt}\, \widehat{\bar{y}}_{nt,raw}^0}.$$

In datasets with within-cell sampling rather than a census, the simple mean can be replaced by the corresponding design-weighted cell mean; the argument below is unchanged.

D.3.2 Consistency of the raw plug-in weighted PPML estimator

Proposition 9 (Consistency of the raw plug-in weighted PPML estimator) *Suppose the conditions of Proposition 8 hold and, in addition,*

$$\sup_{n,t} \left| \widehat{\bar{y}}_{nt,raw}^0 - \bar{y}_{nt}^0 \right| \xrightarrow{p} 0, \qquad \inf_{n,t} \widehat{\bar{y}}_{nt,raw}^0 > c > 0 \quad \textit{with probability approaching one.}$$

Then the weighted PPML estimator using weights $\widehat{a}_{nit,raw} = 1/(N_{nt}\widehat{\bar{y}}_{nt,raw}^0)$ has the same probability limit as the oracle estimator. In particular,

$$\widehat{\boldsymbol{\beta}}_{raw} \xrightarrow{p} \boldsymbol{\beta}^*.$$

Therefore, under Equation (42),

$$\frac{\widehat{\beta}_{b,raw}}{\sigma_n - 1} \xrightarrow{p} \mathbb{E}_{n,t}\left[\frac{1}{\sigma_n - 1} \frac{\partial \log M_{nt}(\tilde{\mathbf{T}}_{rt})}{\partial \mathrm{Tbin}_{rt}^b} \right] + o(1).$$

Proof. The generated weight converges uniformly to the oracle weight:

$$\sup_{n,t} \left| \frac{1}{\widehat{\bar{y}}_{nt,raw}^0} - \frac{1}{\bar{y}_{nt}^0} \right| \xrightarrow{p} 0.$$

Hence the feasible sample moment converges uniformly to the oracle sample moment. Standard consistency then implies that the feasible estimator and the oracle estimator have the same probability limit. The final statement follows directly from Proposition 8. ■

A sufficient condition for the uniform consistency requirement is that the cell mean used to construct $\widehat{\bar{y}}_{nt,raw}^0$ is itself uniformly consistent for the baseline expected cell mass. Under census data this follows from a within-cell law of large numbers. Under sampled data, the same condition holds for the corresponding design-weighted within-cell mean.

The intuition is straightforward. The estimated weights only normalize how much influence each cell has in the score. If the estimated denominator is uniformly close to the true denominator, then the feasible weighted score is uniformly close to the oracle score, and the limiting target does not change. The plug-in step therefore does not alter the estimand; it replaces an infeasible normalization with a sample analogue.

The main practical caveat is finite-sample instability in cells with very small or noisy baseline mass. Because the weight is inverse to baseline mass, a near-zero denominator can create very large weights, which is why the positivity condition is stated explicitly. This is less a conceptual concern than a diagnostics concern: if the weight distribution is well behaved, the plug-in argument is innocuous, while a handful of extreme weights warrants inspecting whether they drive the estimates.

D.3.3 Implementation

The baseline estimator in Section 4.3 is implemented in four steps. The steps mirror the logic of the appendix: define the correct CES-based outcome, remove the common cell scale, fix the cross-cell weighting, and estimate the response using the same identifying variation as in the misallocation regression.

1. Compute the firm-level PPML outcome

$$y_{nit} = \widetilde{\text{TFP}}_{nit}^{\sigma_n - 1}.$$

   Using y_{nit} rather than raw firm TFP makes the dependent variable line up with the CES technology mass $M_{nt}(\cdot)$, so the regression speaks directly to the aggregate technology term in the decomposition rather than to an equal-weight average firm response.

2. Construct the cell offset o_{nt} exactly as in Section 4.3, using region-sector-year revenue and, when available, the corresponding cell-level price index. This step strips out the common cell-level scale component, so that changes in aggregate revenue or prices are not confused with changes in the technology object we want to isolate.

3. Construct the baseline cell weights. In the raw plug-in baseline,

$$\widehat{\bar{y}}^{0}_{nt,raw} = \frac{1}{N_{nt}} \sum_{i \in (n,t)} y_{nit}, \qquad \widehat{a}_{nit,raw} = \frac{1}{N_{nt}\, \widehat{\bar{y}}^{0}_{nt,raw}}.$$

   This step fixes the cross-cell weighting problem. The factor $1/N_{nt}$ gives each cell the same total influence, while $1/\widehat{\bar{y}}^{0}_{nt,raw}$ offsets the tendency of PPML to let high-mass cells dominate. Large and productive firms still matter *within* a cell through the CES aggregator, but large cells no longer dominate *across* cells for purely mechanical reasons.

4. Estimate the weighted PPML regression in Equation (11) using the same temperature bins and fixed effects as in the main text. Keeping the same temperature controls and fixed effects as in the misallocation regression aligns the source of identifying variation across the two

exercises. The comparison between the technology and misallocation channels then comes from the different economic margins, not from different conditioning sets.

D.3.4 Stata implementation equivalence

The Stata implementation used in the paper applies this weighting in an algebraically equivalent normalized form. Let x_{nt} collect the temperature bins and fixed effects. Define

$$y_{nit}^{norm} := \frac{y_{nit}}{\widehat{\bar{y}}_{nt,raw}^{0}}, \qquad o_{nt}^{norm} := o_{nt} - \log \widehat{\bar{y}}_{nt,raw}^{0}, \qquad a_{nit}^{code} := \frac{1}{N_{nt}}.$$

Each PPML score contribution then satisfies

$$a_{nit}^{code}\left(y_{nit}^{norm} - e^{o_{nt}^{norm}+x_{nt}'\boldsymbol{\beta}}\right) = \frac{1}{N_{nt}\,\widehat{\bar{y}}_{nt,raw}^{0}}\left(y_{nit} - e^{o_{nt}+x_{nt}'\boldsymbol{\beta}}\right).$$

The implemented regression that normalizes the dependent variable by the cell mean, subtracts $\log \widehat{\bar{y}}_{nt,raw}^{0}$ from the offset, and applies analytic weights $1/N_{nt}$ is therefore exactly equivalent to estimating the original specification with weights $1/(N_{nt}\widehat{\bar{y}}_{nt,raw}^{0})$.

Overall, Step 1 defines the correct aggregate-technology outcome, Step 2 removes common scale, Step 3 removes the unwanted cross-cell mass weighting, and Step 4 estimates the response using the same residual temperature variation as in the misallocation regression. This is why the resulting coefficient is the right counterpart to the OLS misallocation coefficient in the main text.

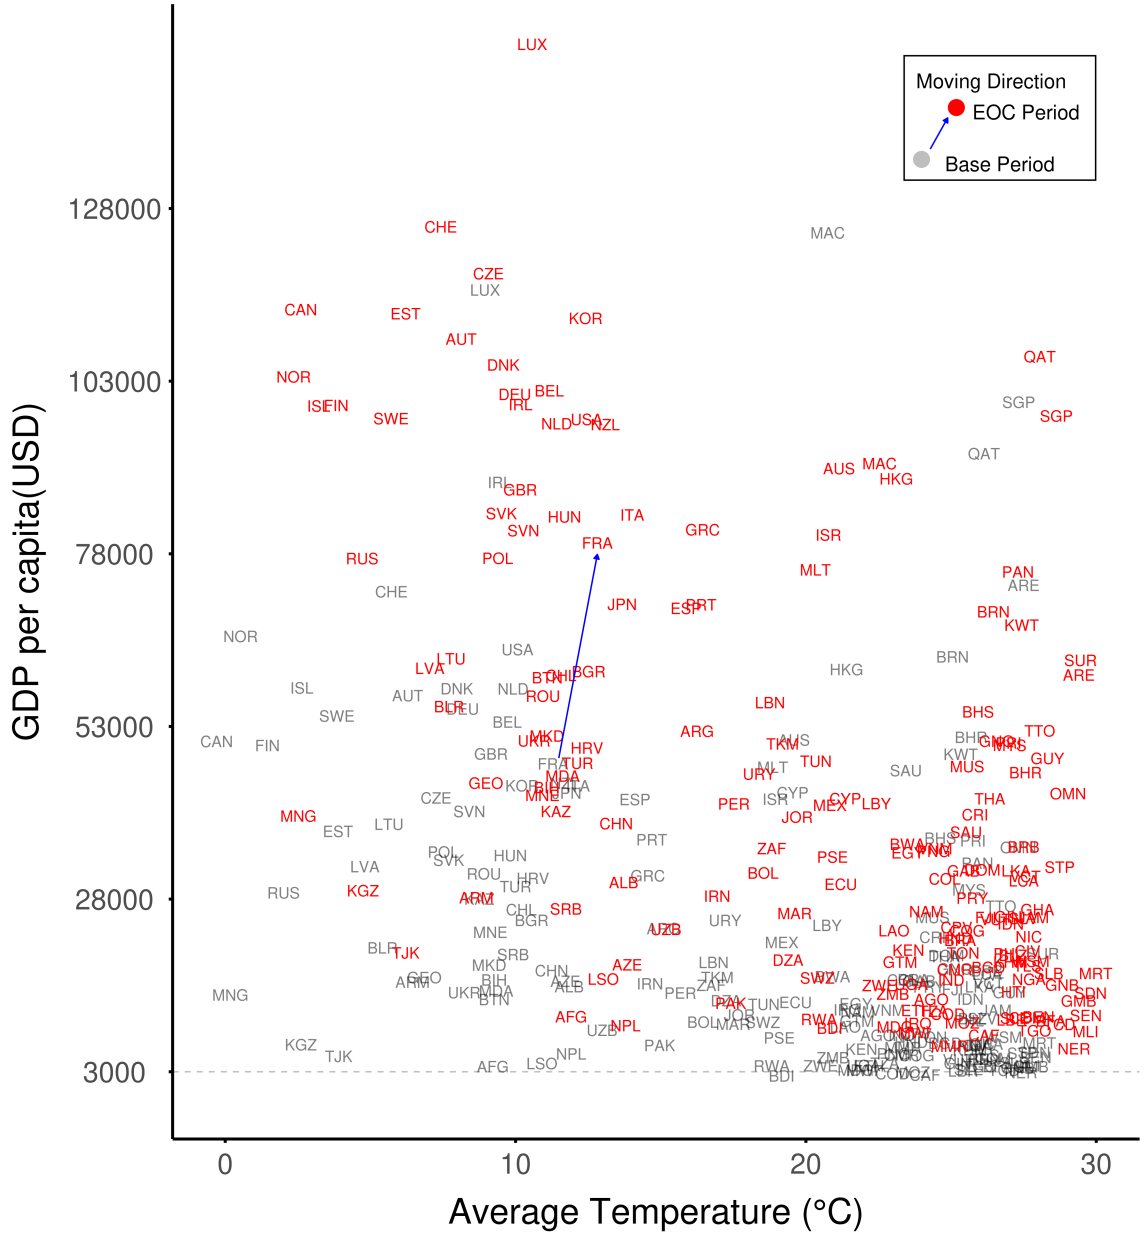
E Additional Figures for Section 4

This appendix provides additional results and figures to complement the main analysis in Section 4.

E.1 Projected Evolution of Income and Temperature under SSP3-4.5

The figure illustrates the projected evolution of income and average temperature from the 2000-2014 baseline to the end of the century (2081-2100) under the SSP3-4.5 scenarios. Among the 172 countries, All 172 countries show a rightward shift (indicating an increase in temperature), and 96% of them also show an upward shift (indicating an increase in per-capita income). In the baseline period, average temperatures are below 5°C in 11 countries, between 5-15°C in 49 countries, between 15-25°C in 62 countries, and above 25°C in 50 countries. Baseline per capita income is below $5000 in 41 countries, between $5000-$30000 in 83 countries, between $30000-$60000 in 38 countries and above $60000 in 10 countries. The blue arrow exemplifies the joint evolution of income and temperature.

Figure E.16: Joint Evolution of Income and Average Temperature from Base Period to End of Century (under SSP3-4.5)



Notes: Grey texts represent the baseline period from 2000 to 2014, and the red texts represent the end of the century (2081 - 2100). End-of-century projection comes from SSP3-4.5 projection. The graph shows the joint evolution of income per capita and average temperature for each country.

E.2 Projection Components

In this appendix we decompose the projected end-of-century misallocation loss into three channels—*shock*, *level*, and *income* effects—and map each channel across the 4,881 regions in the projection sample. Together they account for the 36.41% aggregate TFP loss reported in Section 4.6. The decomposition writes the total effect as

$$\underbrace{\Delta^{\text{Loss}} \ln \text{TFP}_r}_{\text{Total Effect}_r} = \frac{\alpha_{Kn} + \alpha_{Kn}^2(\sigma_n - 1)}{2} \Bigg[\underbrace{\sum_b \left(\lambda^b + \lambda^b_{\text{GDP}_{pc}} \ln \text{GDP}_{pc,r,\text{base}} + \lambda^b_{\bar{T}} \bar{T}_{r,\text{base}} \right) \times \Delta \text{Tbin}^b_r}_{\text{Shock Effect}_r}$$
$$+ \underbrace{\sum_b \lambda_{b,\bar{T}} \text{Tbin}^b_{r,\text{base}} \times \Delta T_r}_{\text{Level Effect}_r}$$
$$+ \underbrace{\sum_b \lambda^b_{\text{GDP}_{pc}} \text{Tbin}^b_{r,\text{base}} \times \Delta \ln \text{GDP}_{pc,r}}_{\text{Income Effect}_r} \Bigg] + \underbrace{\varepsilon_{rt}}_{\text{Residuals}_r}, \tag{43}$$

where Δ denotes the difference between end-of-century and baseline values. The *shock effect* isolates the change in TFP losses from shifts in the daily temperature distribution ΔTbin^b_r, holding baseline income and long-run temperature fixed; the *level effect* isolates the change from rising long-run temperature ΔT_r; and the *income effect* isolates the additional misallocation arising from higher income per capita $\Delta \ln \text{GDP}_{pc,r}$. We aggregate regional TFP losses to the country level using projected GDP shares under SSP3.[37]

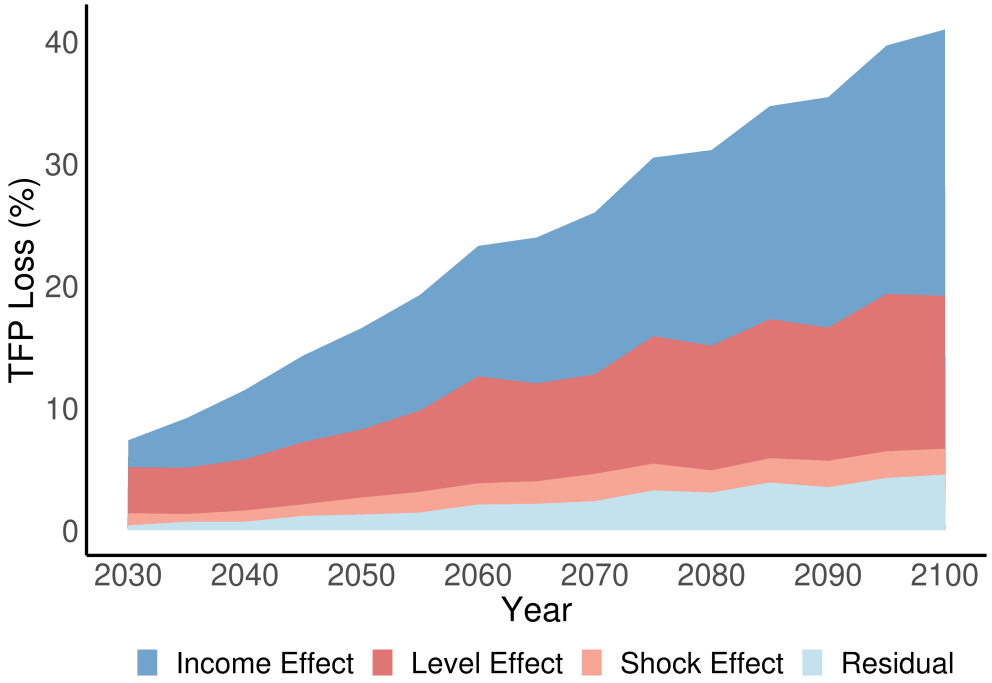
Weighting each country's losses by its projected share of global GDP, the worldwide cost of climate-induced misallocation decomposes as

$$\underbrace{\Delta^{\text{Loss}} \ln \text{TFP}}_{\text{Total Effect}=36.41\%} = \underbrace{\text{Shock Effect}}_{2.13\%} + \underbrace{\text{Level Effect}}_{11.34\%} + \underbrace{\text{Income Effect}}_{19.46\%} + \underbrace{\text{Resid.}}_{3.8\%}.$$

The level and income effects dominate: rising long-run temperature and rising income per capita together account for about 30 percentage points of the 36.41% total, while the direct shift in the daily temperature distribution contributes only 2.13%. Figure E.17 plots the time-series evolution of each component from 2030 to 2100. As a sanity check, holding all marginal effects at baseline values and varying only the daily temperature distribution still delivers a 2.14% global TFP loss—consistent with the shock-effect estimate above.

37. Grid-level SSP3 GDP data from Wang and Sun (2022) provides projected GDP shares.

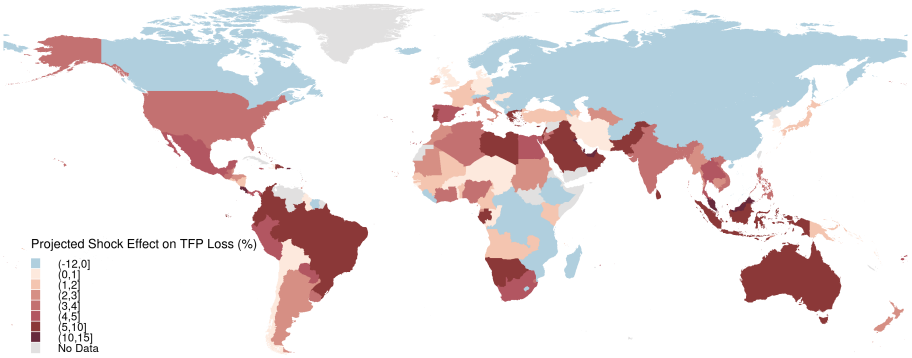
Figure E.17: TFP Loss Decomposition Over Time (SSP3-4.5)



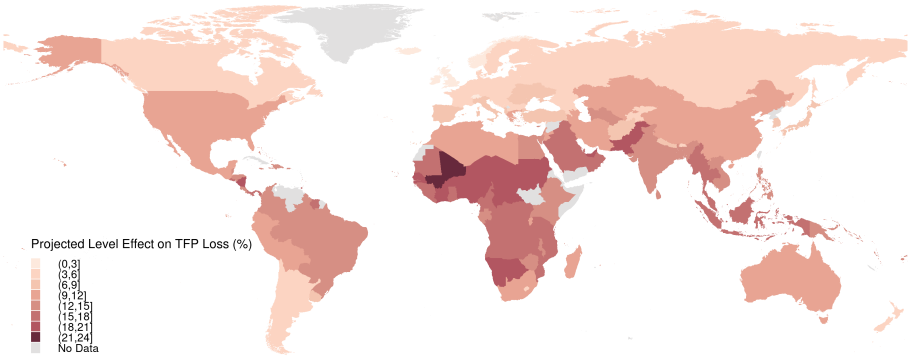
Figures E.18a–E.18c map the spatial distribution of each component across the 4,881 regions in our projection sample.

Figure E.18: Three Effects Contribution to Projected TFP Loss (SSP3-4.5)

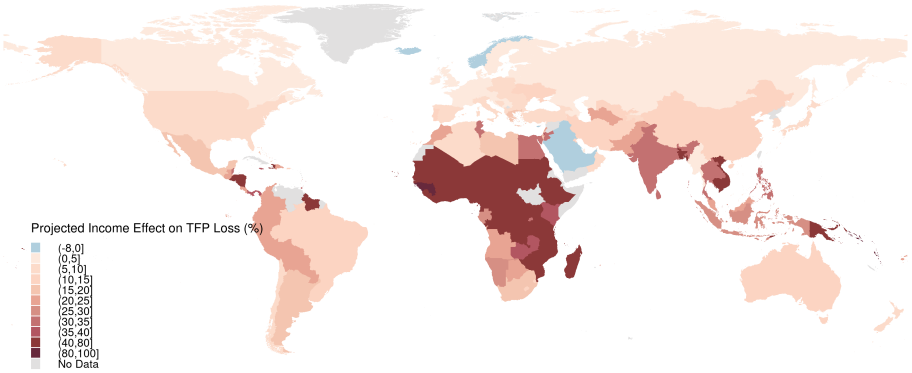
(a) Shock Effects Projection



(b) Level Effects Projection



(c) Income Effects Projection



F Firm Dynamics Model Details

This appendix provides additional details on the household problem, the timing of events, and the perturbation solution method for the model in Section 5.

F.1 Household Budget Constraint and Adaptation-Resource Pricing

The household owns all firms in the region, and all profits and revenues from tax-like wedges are rebated lump sum. At the end of period t, the household consumes C_t and sets a supply schedule R_{t+1} to meet firms' demand for adaptation resources. At the start of $t+1$, before climate shocks are realized, the predetermined quantity R_{t+1} is supplied and the contracted amount is distributed to each firm. The period-t budget constraint is

$$C_t = w_t N + w_t^R R_t + \Pi_t + S_t,$$

where w_t is the wage for unskilled labor, w_t^R is the price of adaptation resources, Π_t denotes total firm profits, and S_t denotes lump-sum rebates.

The household's optimality conditions for the supply of adaptation resources imply that the price of adaptation resources satisfies

$$w_{t+1}^R = \eta \frac{\mathbb{E}_t\left[\Lambda_{t+1} C_{t+1}\right]}{\mathbb{E}_t\left[\Lambda_{t+1}\right]}.$$

Thus, w_{t+1}^R increases with the utility-cost parameter η and with the SDF-weighted expectation of next-period consumption.

F.2 Solution via Risk-Adjusted Perturbation

We solve the model in closed form using a risk-adjusted perturbation around the homogeneous deterministic steady state (HDSS) as in David and Venkateswaran (2019). The key difficulty is that firms' Euler equations feature *risk-adjusted* expectations—products of the stochastic discount factor (SDF) and future payoffs—under non-linear temperature exposure and endogenous heterogeneity. Our approach proceeds in three steps.

Step 1: Risk-adjusted approximation of FOCs. We extend the risk-adjusted approximation in Coeurdacier, Rey, and Winant (2011) beyond log-normal environments by developing a second-order expansion of stochastic variables in conditional moments of shocks (means, variances, and covariances), and then applying it systematically to the firms' optimality conditions. As a simple illustration, the one-period risk-free return is $R_{f,t+1} = \frac{1}{\mathbb{E}_t[\Lambda_{t+1}]}$. Let $\widehat{x}_t \equiv x_t - x^{\text{HDSS}}$ denote deviations from the HDSS and let $\lambda_{t+1} \equiv \log \Lambda_{t+1}$. A second-order expansion delivers the convenient risk-adjusted representation

$$\widehat{r}_{f,t+1} \equiv \log R_{f,t+1} - \log R_f^{\text{HDSS}} \;=\; - \underbrace{\widehat{\mathbb{E}_t \lambda_{t+1}}}_{\substack{\text{Consumption} \\ \text{Smoothing}}} \; \underbrace{- \frac{1}{2} \mathbb{VAR}_t\left(\widehat{\lambda}_{t+1}\right)}_{\text{Precautionary}}$$

which makes clear that the risk-free rate falls when expected marginal utility growth rises (consumption smoothing) and when uncertainty in the pricing kernel increases (a precautionary-savings effect).

Step 2: Bilinear policy-function approximation. We approximate firms' policy rules for next-period capital and adaptation using an affine–bilinear expansion in (i) firm-level state vectors $\mathbf{x}_{it}$ (e.g., persistent and transitory productivity, wedges, and product exposure) and (ii) aggregate state vectors $\mathbf{x}_t$ (e.g., aggregate productivity, aggregate capital, and temperature moments). For example, for log capital $k_{i,t+1} \equiv \log K_{i,t+1}$,

$$k_{i,t+1} = k^{\text{HDSS}} + \mathbf{G}_i^\top \widehat{\mathbf{x}}_{it} + \mathbf{G}_a^\top \widehat{\mathbf{x}}_t + \underbrace{\widehat{\mathbf{x}}_{it}^\top \mathbf{G}_{ia}\, \widehat{\mathbf{x}}_t}_{\substack{\text{heterogeneous response to} \\ \text{aggregate shocks and uncertainty}}} + \mathbf{G}_d^\top \widehat{\mathbf{d}}_t + \mathbf{G}_\Gamma^\top \widehat{\mathbf{\Gamma}}_{micro,t},$$

and analogously for $\widetilde{\beta}_{i,t+1}$. The interaction term $\widehat{\mathbf{x}}_{it}^\top \mathbf{G}_{ia} \widehat{\mathbf{x}}_t$ is essential: it captures how a common temperature realization (or a change in forecast uncertainty) loads differently across firms through heterogeneous effective exposure and wedges, generating endogenous dispersion in firm decisions.

Step 3: Aggregation and coefficient identification. We apply the same second-order (affine-bilinear) expansion to the aggregation equations. Because the policy rules are affine-bilinear, aggregates depend on the cross-sectional distribution of $\mathbf{x}_{it}$ only through a finite set of moments, collected in $\mathbf{\Gamma}_{micro,t}$ (in the spirit of Krusell and Smith (1998)). Finally, we substitute the approximated policy rules and aggregation equations into all equilibrium conditions (firms' Euler equations, factor-market clearing, and the household pricing kernel), and solve for the unknown coefficients by matching terms of equal order (constant, linear, and bilinear). This guess-and-verify procedure yields closed-form expressions for both micro policies and macro equilibrium objects, including MRPK dispersion and the resulting aggregate TFP losses.

G Model Estimation Details

This appendix provides the full set of moment conditions used for identification, the fixed-point algorithm used for implementation, and the aggregation formulas used in the quantitative exercises.

G.1 Moment Conditions for Model Identification

This section summarizes the moment restrictions used to identify the region-level parameters and firm-level latent components.

(i) Firm-level Moment Conditions For each region, we assume the permanent component of productivity $\bar{a}_i$ and $\overline{mrpk}_i$ for each firm satisfy the standard firm fixed effect moment conditions. Given the presence of heterogeneous sensitivities, we assume the firm-level a_{it} and $mrpk_{it}$, once netted out the temperature-related and permanent components, should satisfy the Yule-Walker conditions implied by their respective AR(1) structure.[38] Lastly, for each region r', the correlated distortion satisfy $\mathbb{E}^i[\bar{a}_i(\overline{mrpk}_i - \omega^a \bar{a}_i)|r(i) = r'] = 0$. Given β_i, Ψ_{it} and T^*, these conditions jointly

38. These moment conditions are also conditional on region-year fixed effects.

identifies $\left(\omega_r^a,\ \rho_{z,r},\ \rho_{\phi^\tau,r},\ \{\bar{a}_i\}_i\right)$ for each region r. We can also recover all innovations $\varepsilon_{it}^{\phi^\tau}$, ε_{it}^{z} and therefore z_{it} and τ_{it} up to some unobserved random damage $\xi_{it}(T_{rt}-T^*)$.

(ii) Product sensitivity β_i. We assume intrinsic product sensitivity is constant within 3-digit category p: $\beta_i = \beta_{q(i)}$, where $q(i)$ maps firm i to its category. Define the MRPK residual after removing region–year general-equilibrium components and endogenous adaptation terms,

$$\widetilde{mrpk}_{i,t+1} \;\equiv\; mrpk_{i,t+1} - \theta_A\big(a_{t+1} + \tilde{\alpha}_K k_{t+1}\big) - \big(\tilde{\beta} + \Psi_{i,t+1}\big)(T_{t+1} - \mathbb{E}_t[T_{t+1}]).$$

Then β_p is identified by the orthogonality condition implied by equation (**??**), for each category p,

$$\mathbb{E}\Big[\mathbf{1}\{q(i)=p\}\left(\widetilde{mrpk}_{i,t+1} - \beta_p(T_{t+1}-\mathbb{E}_t[T_{t+1}])\right)(T_{t+1}-\mathbb{E}_t[T_{t+1}]) \;\Big|\; \gamma_i, \gamma_{c(r)t}\Big] = 0. \tag{44}$$

Intuitively, after netting out endogenous adaptation and region–year aggregate components, exposure to temperature surprises should operate through intrinsic product sensitivity only, conditional on firm and country–year fixed effects. The regional average product sensitivity can also be recovered by using $\bar{\beta}_r = \mathbb{E}[\hat{\beta}_i | r = r(i)]$.

(iii) Adaptation cost $\eta_r H_r$. To identify regional variation in adaptation costs, we parametrize the inverse cost shifter as $\frac{1}{\eta_r H_r} \;\equiv\; f\Big(\overline{\ln \text{GDP}_r}\Big)$, where $f(\cdot)$ is a third-order polynomial in $\overline{\ln \text{GDP}_r}$. Given this parametrization, equation (**??**) implies a set of orthogonality conditions that link temperature forecast errors to cross-firm heterogeneity in exposure and scale. Specifically, we impose the moments

$$\mathbb{E}\Bigg[\underbrace{\Big(mrpk_{it} - \rho_{\phi^\tau}\phi^\tau_{i,t-1} - \big(\widehat{\beta}_i + \Psi_{it}\big)\Big[(T_{rt} - \mathbb{E}_{t-1}T_{rt}) - \Theta\big(\overline{\beta}_i + \tilde{\beta}\big)\mathbb{VAR}_{t-1}T_{rt}\Big]\Big)}_{\text{MRPK residual}}$$
$$\cdot \underbrace{\Big(\overline{\ln \text{GDP}_r}^{\,k} \cdot (\bar{T}_r - T^*) \cdot \Big[(\bar{T}_r - T^*)\widehat{\beta}_i + \bar{a}_i + \rho z_{it} - \alpha\tau_{i,t+1}\Big]\ (T_{rt} - \mathbb{E}_{t-1}T_{rt})\Big)}_{\text{instruments for adaptation-cost shifter}} \mid \gamma_i, \gamma_{rt}\Bigg] = 0, \tag{45}$$

$\forall\, k \in \{0,1,2,3,4\}$. Intuitively, after netting out the wedge dynamics and the common aggregate temperature component, the remaining cross-firm variation in MRPK responses to temperature surprises should be orthogonal to functions of regional income interacted with the model-implied drivers of endogenous adaptation, within-firm and within-region-year. These moments discipline the coefficients of $f(\overline{\ln \text{GDP}_r})$, and therefore identify $\frac{1}{\eta_r H_r}$ across regions.

(iv) Bliss point T^* and random-damage dispersion σ_ξ^2. The model implies that the bliss point T^* and the dispersion of idiosyncratic climate damages σ_ξ^2 can be jointly identified from region–sector–year variation in (i) TFP-growth volatility and (ii) MRPK dispersion. As in equation (23), after netting out heterogeneous exposure and the components driven by forecast errors and aggregate uncertainty (captured by some $\text{VolCorrection}_{rst}$ and $\text{MisCorrection}_{rst}$), the remaining temperature dependence is governed by the quadratic damage term $\sigma_\xi^2(T_{rt}-T^*)^2$. Define the residual volatility and residual MRPK dispersion, $\widetilde{V}^{\Delta a}_{rst} \equiv \mathbb{VAR}^i_{rst}(a_{it} - a_{i,t-1}) - \text{VolCorrection}_{rst}$, and

$\widetilde{V}^{mrpk}_{rst} \equiv \mathbb{VAR}^{i}_{rst}\left(mrpk_{it}\right) - \text{MisCorrection}_{rst}$. Because $a_{it} - a_{i,t-1}$ loads on the idiosyncratic damage realizations in both t and $t-1$, the model predicts $\widetilde{V}^{\Delta a}_{rst}$ scales with $(T_{rt}-T^*)^2 + (T_{r,t-1}-T^*)^2$, while $\widetilde{V}^{mrpk}_{rst}$ scales with $(T_{rt}-T^*)^2$. We therefore impose the following model-consistent moment conditions:

$$\mathbb{E}\Big[\Big(\widetilde{V}^{\Delta a}_{rst} - \sigma^2_\xi\big[(T_{rt}-T^*)^2 + (T_{r,t-1}-T^*)^2\big]\Big)\,(T^2_{rt} + T^2_{r,t-1}) \;\Big|\; \gamma_{rs}, \gamma_{ct}\Big] = 0, \tag{46}$$

$$\mathbb{E}\Big[\Big(\widetilde{V}^{\Delta a}_{rst} - \sigma^2_\xi\big[(T_{rt}-T^*)^2 + (T_{r,t-1}-T^*)^2\big]\Big)\,(T_{rt} + T_{r,t-1}) \;\Big|\; \gamma_{rs}, \gamma_{ct}\Big] = 0, \tag{47}$$

$$\mathbb{E}\Big[\Big(\widetilde{V}^{mrpk}_{rst} - \sigma^2_\xi (T_{rt}-T^*)^2\Big)\,(T_{rt}-T^*)^2 \;\Big|\; \gamma_{rs}, \gamma_{ct}\Big] = 0. \tag{48}$$

Intuitively, once we purge the components of volatility and dispersion accounted for by heterogeneous exposure to temperature shocks and risk, the remaining cross-sectional volatility in productivity growth and dispersion in MRPK should be explained by random climate damages that rise with the squared distance from the bliss point. The first two moments use time-series variation in temperature levels to locate T^* (via how volatility co-moves with temperature), while the third moment disciplines the scale of the quadratic term and helps pin down σ^2_ξ.

Taken together, these moments jointly identify the common parameters $\{T^*, \sigma^2_\xi\}$, as well as the region-specific objects $\{\bar{\beta}_r,\ (\eta_r H_r)^{-1},\ \sigma^2_{\hat{\beta},r},\ \sigma^2_{\bar{a},r},\ \sigma^2_{\varepsilon^z,r},\ \sigma^2_{\varepsilon^{\phi^\tau},r},\ \rho_{\phi^\tau,r},\ \rho_{z,r},\ \omega_{a,r}\}_r$.

G.2 Implementation via a Fixed-Point Algorithm

We implement the estimator using a transparent fixed-point algorithm that circumvents the computational burden of solving the full joint GMM problem in one step. The algorithm iterates between (i) inverting the model to recover firm-level latent objects and region-level moments given a current parameter guess, and (ii) updating the parameters using the moment conditions described above.

1. **Initialization.** Start from an initial guess for regional adaptation costs and intrinsic product sensitivities, e.g. $\frac{1}{\eta_r H_r} = 0$ and $\beta_i = 0$, together with an initial guess for the common parameters $\{T^*, \sigma^2_\xi\}$. Given these values, compute the implied steady-state adaptation $\{\widetilde{\beta}_r\}_r$ and the corresponding regional sensitivity multipliers. This step enables model inversion: we recover firm-level fundamentals and construct all region–sector moments needed for the subsequent updates.

2. **Update the bliss point and random-damage dispersion $\{T^*, \sigma^2_\xi\}$.** Using the objects recovered in Step 1, construct the region–sector–year correction terms $\text{VolCorrection}_{rst}$ and $\text{MisCorrection}_{rst}$. We then implement the moment conditions in equations (46)–(48) via simple two-way fixed-effects regressions to update $\{T^*, \sigma^2_\xi\}$.

3. **Update intrinsic product sensitivity $\{\beta_i\}$.** Given the updated T^*, recompute the endogenous exposure component Ψ_{it} and form the residualized return $\widetilde{mrpk}_{it}$ that nets out the region–year general-equilibrium terms and endogenous adaptation. We then estimate the

moment condition in equation (44) using a high-dimensional fixed-effects specification that allows for 3-digit-category-specific slopes, which updates $\beta_i = \beta_{q(i)}$ for all firms.

4. **Update regional adaptation costs** $\{(\eta_r H_r)^{-1}\}_r$**.** Under the fourth-order polynomial specification of $\frac{1}{\eta_r H_r} \equiv f\left(\overline{\ln \text{GDP}_r}\right)$, using the updated $\{\beta_i\}$, we first net out the intrinsic temperature-shock component from $mrpk_{it}$. We then estimate the moment conditions in equation (45) by projecting the relevant MRPK residual on $\left(\bar{T}_r - T^*\right)\left[\left(\bar{T}_r - T^*\right)\widehat{\beta}_i + \bar{a}_i + \rho z_{it} - \alpha\tau_{i,t+1}\right]\left(T_t - \mathbb{E}_{t-1}[T_t]\right)$ interacted with a polynomial in $\overline{\ln \text{GDP}_r}$. This yields an updated estimate of $f(\overline{\ln \text{GDP}_r})$ and hence $(\eta_r H_r)^{-1}$.[39]

5. **Iterate to convergence.** Update the parameter guesses and repeat Steps 1–4 until $\max\left\{\|\Delta\beta\|_i, \left\|\Delta(\eta H)^{-1}\right\|_r, |\Delta T^*|, |\Delta\sigma_\xi^2|\right\}$ falls below a pre-specified tolerance, ensuring convergence of both firm-level latent objects and region-level moments.

While we do not impose a formal uniqueness proof, in practice the fixed-point iteration converges quickly (typically in fewer than 20 iterations) and is robust to a wide range of reasonable initial guesses. Intuitively, temperature-related terms discipline a specific subset of equilibrium wedges and exposure objects, while the remaining components of firm dynamics are primarily pinned down by standard cross-sectional and time-series moments, which stabilizes the overall mapping from moments to parameters.

G.3 Aggregation Details

We aggregate the world economy using a Cobb–Douglas output index across countries and regions. Let

$$\log Y_t^{\text{Global}} = \sum_c \omega_{ct} \log Y_{ct}, \qquad \log Y_{ct} = \sum_{r \in R(c)} \omega_{rt}^c \log Y_{rt},$$

where ω_{ct} is country c's global output share and ω_{rt}^c is region r's share of country c's output. This implies

$$\log \text{TFP}_t^{\text{Global}} = \sum_r \omega_{rt} \log \text{TFP}_{rt}, \qquad \omega_{rt} \equiv \omega_{ct}\omega_{rt}^c.$$

39. For the small set of regions (less than 2% of observations) in which the fitted $f(\cdot)$ implies negative values, we set $(\eta_r H_r)^{-1} = 0$.