

Revisiting the Global Distribution of Climate Damages

Romain Fillon^{1,2}

¹ *Paris-Saclay Applied Economics, Graduate School of Economics and Management, Université Paris-Saclay, France*

² *Centre International de Recherche sur l'Environnement et le Développement, CIREN, Université Paris-Saclay, France*

October 12, 2026

Abstract

The world's cold economies are portrayed as climate winners. A hump-shaped response of output to annual mean temperature suggests why: they lie below the optimum, so warming should move them there. Measured by daily temperatures in a climate-economy model, that gain disappears: the share of world population in gaining cells falls from 11.7% to 0.1%. The two metrics disagree about the winter that warming removes: the annual mean values it at least 1.8 times what a daily regression does. The annual-mean slope measures not cold days but the cold–temperate gap between countries, and between one country's harsh and mild winters.

Keywords: climate impacts, spatial general equilibrium, temperature distribution, adaptation.

1 Introduction

Quantitative models used to project how climate damages fall across space predict that several of the world’s coldest economies, among them Russia, Canada and the Nordic countries, will be net beneficiaries of warming (Krusell and Smith, 2022; Cruz and Rossi-Hansberg, 2024; Nath et al., 2024). The prediction is consequential, because it concerns rich economies whose policy choices shape the global mitigation path. It follows mechanically from one modelling convention: these models evaluate their damage function at a single summary of local temperature (an annual mean, or the decadal mean of a single month) and a hump-shaped response in the level of that summary places cold economies below the optimum and hot ones above it, so that warming moves the first towards the peak, faster still under polar amplification, and the second away from it. The empirical literature on climate impacts abandoned the annual mean more than a decade ago, because the damage response is sharply non-linear at high frequencies and what harms output is the hot day rather than the average one (Schlenker and Roberts, 2009; Burke et al., 2015), and the global estimates of mortality damages are built on the full distribution of daily temperatures for the same reason (Carleton et al., 2022). Whether that gap matters for the models’ central prediction has not been tested. This paper tests it, and finds that the projected Northern gain is a property of the exposure metric.

The test moves a single lever. I estimate dose-response functions for productivity from bi-lateral trade flows, following Rudik et al. (2022), twice: once with exposure measured as the regional annual mean, as the spatial literature does, and once with exposure measured as the full distribution of daily temperatures. Sample, estimator, controls and fixed effects are the same in both estimations; only the way temperature enters differs. I then propagate each function through a dynamic spatial equilibrium model of trade and migration (Eaton and Kortum, 2002; Redding and Rossi-Hansberg, 2017) along the same emissions pathway, SSP2-4.5. The model is solved in relative changes, so that unobserved fundamentals are never recovered by inversion; that choice closes the extensive margin of adaptation onto currently empty high-latitude cells, which part of the literature’s Northern gains relies on, and it is therefore conservative against the result.

The first result concerns identification. Across 15 specifications that vary the functional form and the winsorisation of the tails, the projected change in the Northern productivity fundamental above 45°N runs from +0.012 to +0.024 under the annual mean, and the intervals exclude zero in 14 of the 15 specifications. Under the daily distribution the same specifications run from −0.001 to +0.003, with a median of +0.002, and none excludes zero. The difference between the two is signed in 15 of 15 specifications. An annual-mean specification does not deliver a coarser estimate of the daily object; it delivers a different one, and the confidence it displays about the North is manufactured by aggregating exposure before estimating.

The second result concerns welfare. Propagated through the model, the annual-mean function gives the population above 45°N a cumulative welfare gain of +0.0115 by 2100; the daily-distribution function gives the same population a loss of −0.0063, a loss that holds in 60 of 60 Monte Carlo draws of the daily dose-response once each is propagated through the model. The loss is not carried by a handful of cells: 99.5% of the Northern population lives in cells whose welfare falls once the daily distribution is imposed. At the world level the share of the population living in cells that gain from warming falls from 11.7% under the annual mean (8% in [Nath et al. \(2024\)](#)) to 0.1% under the daily distribution. The two findings fit together: at daily resolution the estimation leaves the sign of the North's own productivity response open, but welfare also depends on where people and production settle once the climate has changed. Solved with warming confined to the North, the annual-mean gain becomes a loss under daily exposure; warming elsewhere deepens the loss rather than creating it.

The annual mean is a valid measure of exposure only if two places with the same mean temperature suffer the same damage, whatever the mix of hot and cold days behind that mean. They do not, for two reasons, and the paper separates and measures them. The first is Jensen's inequality: the damage function is non-linear, so its value at the mean day is not its average over the year; measured directly, this term is small. The second is that the two specifications do not estimate the same function. A regression on annual means never observes a freezing day or a day at 35°C; it observes that countries with a mean of 12°C produce more than countries with a mean of 0°C, and

its slope absorbs everything that separates the two groups of countries, not the effect of a cold day. The daily specification must allocate the effect across the days a country actually experiences, and what it finds is that hot days are costly and freezing days are costly too. The two specifications therefore put different prices on the same winter. When a country's annual mean rises from 2 to 5°C, the annual curve credits it with what distinguishes countries that already sit at 5°C; the daily curve prices the freezing days that warming removes, and prices them at a fraction of what the annual mean implies. The projected Northern gain is the first of these two quantities, and it is not a measure of what a winter is worth. The result is not an artefact of the bilateral estimator or of the amenity channel. I re-run the contrast inside the estimation pipeline of [Cruz and Rossi-Hansberg \(2024\)](#), on fundamentals inverted from their model, under their fixed effects and controls, and changing only the exposure metric; the annual-mean projection exceeds the daily one by +0.443 in the log fundamental, and the daily projection can no longer be distinguished from zero.

Why has this not been done before? Models with endogenous emissions obtain regional climate by pattern scaling, which maps a global mean temperature path to regional annual means and produces nothing below them; no reduced form yet maps emissions to a regional distribution of daily temperatures, so a model can have endogenous emissions or the daily distribution, not both. Reading the annual-mean convention as the cautious option is a mistake: evaluating a non-linear response at a projected mean imposes that the mapping from a region's mean to its distribution is stationary under warming, and that assumption fails where it has been tested: the coldest winter days in North America warm about twice as fast as the seasonal mean ([Blackport and Fyfe, 2024](#)), and the hottest summer days in Europe warm faster than the summer mean ([Patterson, 2023](#)).

Section [2](#) sets out the model and the estimator, Section [3](#) reports the general-equilibrium results.

2 Estimation of climate damages

2.1 Dynamic spatial climate-economy model

The counterfactual runs in a dynamic spatial equilibrium model of trade and migration in the tradition of [Eaton and Kortum \(2002\)](#) and [Desmet et al. \(2018\)](#). The model is deliberately standard: it is the class of model that [Desmet and Rossi-Hansberg \(2024\)](#) survey and that the spatial literature on climate damages uses, so that nothing in the result comes from a departure in the economic environment, and the object that changes between scenarios is how climate exposure enters it.

There are N cells r and discrete five-year periods t . Cell r carries population $L_t(r)$, fundamental productivity $\bar{Z}_t(r)$, fundamental amenity $\bar{a}(r)$, a fixed stock of structures, and an exposure $\mathcal{E}_t(r)$ to climate. Climate enters through a multiplicative shifter on each fundamental,

$$Z_t(r) = \bar{Z}_t(r) \exp[\mathcal{E}_t(r) f_z(\cdot; \zeta_z)], \quad a_t(r) = \bar{a}(r) L_t(r)^{-\lambda} \exp[\mathcal{E}_t(r) f_a(\cdot; \zeta_a)], \quad (1)$$

where $\bar{Z}_t(r)$ grows at an exogenous rate, λ is the congestion elasticity, f_z and f_a are dose-responses mapping a temperature to a log shift in productivity and in amenities, ζ_z and ζ_a are their coefficient vectors, and $\mathcal{E}_t(r)f$ denotes the exposure of cell r at date t to a dose-response f . The two exposure metrics differ only in this operator. Under the annual mean, $\mathcal{E}_t(r)f = f(\bar{T}_t(r))$: the function is evaluated at the cell's mean temperature. Under the daily distribution, $\mathcal{E}_t(r)f = \int f(T) dF_t(T | r)$: the function is integrated over the distribution of daily temperatures the cell experiences. The equilibrium conditions, the migration and trade structure and the calibration are standard and are set out in [Supplement A](#); the model is solved in relative changes, so that unobserved fundamentals are never recovered by inversion.

2.2 Estimation procedure

Following [Eaton and Kortum \(2002\)](#) and [Rudik et al. \(2022\)](#), I exploit the fact that bilateral trade flows embed information about relative productivities and bilateral migration flows about

relative amenities. The dose-responses are not taken from outside the model. Their functional form is chosen, as a polynomial or a natural spline in temperature bins, and what is estimated is the coefficient vector of that form: ζ_z is recovered from the model’s bilateral trade equation and ζ_a from its bilateral migration equation, both derived from the equilibrium conditions of the model with temperature entering through the flexible form. The shifters of (1) and the counterfactual that propagates them therefore rest on the same structure, and the only free choice in the estimation is the form of f , which the specification grid varies. Each metric is paired with the dose-response estimated under it, so that a scenario is a pair $(\mathcal{E}, f_z)$. In the body f_a is set to zero, since bilateral migration data do not identify it; the choice is a lower bound rather than a neutral one: Supplement D.3 shows the Northern contrast to be negative for every Northern economy at the interior estimate of f_a as well, and deeper there, while the global loss the body reports is the smaller of the two. More details on the estimation are in supplements B and C.

In the taxonomy of Lemoine et al. (2025), this places the exercise in the class of structurally consistent general-equilibrium damage estimates: the object identified is not the local causal response to a transitory weather shock, but the damage function consistent with the spatial equilibrium that observed trade and migration flows reveal. Two features distinguish this from the panel fixed-effects estimates used in global-scale damage estimation (Cruz and Rossi-Hansberg, 2024; Desmet and Rossi-Hansberg, 2024). The dependent variable is a large matrix of bilateral flows rather than a country-level outcome, which draws on a different source of variation. And the fundamentals enter through observed administrative data rather than through model inversion of unobserved productivities and amenities. Relative to Rudik et al. (2022), I cover a larger set of countries, use relative rather than absolute winsorisation thresholds, and estimate the same dose-response under two competing definitions of climate exposure.

The two metrics are the annual mean, which treats $\bar{T}_t(r)$ as a sufficient statistic for local climate with a shape-preserving shift under warming, and the full distribution of daily mean temperatures, which lets the shape of the intra-annual distribution change. Each is estimated over the same grid of 15 specifications (polynomials of degree two to four, natural splines with three or four degrees

of freedom, three symmetric winsorisation thresholds), because tail behaviour drives the result. Akaike weights place all mass on one form per arm, so that weighting amounts to selecting¹. The contrast is reported across the whole grid; the dose-responses of Figure 1 are drawn for one specification, with a band of $B = 300$ Monte Carlo draws from the clustered variance that reports parameter uncertainty conditional on that form.

2.3 Productivity from bilateral trade

At equilibrium, normalising region r 's expenditure on sector- k goods from s by its domestic absorption cancels the importer-side price index (Eaton and Kortum, 2002; Head and Mayer, 2014):

$$\log \frac{X_t^k(rs)}{X_t^k(rr)} = \theta [\mathcal{E}_t(s)f_z - \mathcal{E}_t(r)f_z] + \zeta_y \theta^k \log \frac{y_t(s)}{y_t(r)} + \zeta_C^\top C_{rs} + \rho_t^k + \varepsilon_{rst}^k, \quad (2)$$

where ρ_t^k are product $\times$ year effects, C_{rs} collects time-invariant gravity controls (distance, common language, colonial link, trade agreement), and errors are clustered by origin and destination. The trade elasticity θ is imposed rather than estimated, at the value of the general equilibrium (Desmet et al., 2018). The income ratio enters as a control rather than as a structural term, weighted by the sectoral trade elasticities θ^k of Caliendo and Parro (2015), because it follows the sectoral exporter's cost channel. The climate term is not weighted by θ^k : f_z is a fundamental of the location, common to sectors, and sectoral weights would load the composition of each pair's trade onto the recovered dose-response. I estimate (2) by PPML on BACI–CEPII bilateral flows and GLORIA domestic absorption, 1995–2019, 20 manufacturing and agricultural sectors, 3,511,478 observations.

The climate term is a difference in exposure between exporter and importer, and it is identified from the cross-section of that difference across pairs, not from its variation within pairs over time. The choice is deliberate and forced. Deliberate, because the counterfactual is a permanent shift in climate, and within-pair variation recovers the response to weather (Deschenes and Greenstone,

¹The identification diagnostics report 12 specifications: the contrast adds a fourth-degree polynomial as a test of flexibility, which carries no Akaike weight in either arm.

2007; Burke et al., 2015). Forced, because bilateral temperature differentials are quasi-permanent: pair $\times$ product effects absorb 97.7% of their variance, and their within-pair standard deviation is 15.2% of the total. A differenced estimator is blind to a globally synchronous shock only when the response is linear: with curvature, a uniform shock shifts comparative advantage towards the partner on the flatter part of the response, and that curvature is what the bin specification recovers from cross-sectional variation in exposure. What cancels in every ratio is the damage component common to all countries, the object aggregate time-series estimators identify (Bilal and Känzig, 2026); the two designs are complementary.

The threat to a cross-sectional design is that cold and warm country pairs differ on dimensions other than climate, and that some of those dimensions are correlated with temperature. Three pieces of evidence bear on it. The gravity controls absorb 4.9–14.6% of the between-pair climate variation without being collinear with the climate term, so they remove the part of the confounding that distance and history induce. Coefficient stability in the sense of Oster (2019) gives, over the 12 specifications on which the diagnostics are computed, a smallest positive selection ratio of $\delta^* = 2.2$ (the cubic polynomial winsorised at 2.5%) and 14.0 for the reference specification: unobservables would have to be at least 2.2 times as important as the gravity controls to explain the response away, against a threshold of one. A negative δ^* comes out in 4 specifications, the case in which selection would have to run against the observed correlation, and the minimum is attained where the coefficient moves by 89.7% with the controls, against 12.9% for the reference. Under exporter $\times$ year effects, which absorb the exporter’s non-climate capability in full (Astier et al., 2026), the hot-side slope keeps its sign at -1.05 per $^{\circ}\text{C}$ with a band that excludes zero; the other schemes leave its precision open. A separate threat applies to daily specifications estimated in panels: Jones et al. (2026) show that extreme-day counts carry warming-induced trends correlated with baseline climate. The bias they describe comes from the variation within a unit over time. The estimator here draws almost nothing from it: replacing each pair’s differentials by their 1995–2019 means moves the Northern contrast by $+0.05$ of its standard error. Adding pair $\times$ product fixed effects, the design they study, turns the cold arm positive; estimating within five-year windows, their remedy,

reduces it (Supplement B.3). One further point applies to the daily arm alone. A confound that separates rich cold countries from poor hot ones loads on the mean of the temperature distribution, which the two arms share. What the daily arm uses in addition is the shape of the distribution at a given mean: two countries with the same annual mean can spend their year in mild winters and mild summers or in freezing winters and hot summers, and that difference is a matter of continentality, not of development. The curvature of the response is identified from that variation, and it is the curvature, not the level, that the contrast of Section 2.5 turns on. The replication of that contrast on fundamentals inverted from the Cruz and Rossi-Hansberg (2024) model, under their sub-region $\times$ period fixed effects (Supplement E), is the check that none of this rests on the bilateral data or on the trade fundamental.

Throughout, a log point is one hundredth of a log; projected changes and welfare effects are reported as changes in logs, and ratios between reported quantities are computed before rounding.

Figure 1 draws the estimated f_z under each metric on common axes. The daily-distribution response is flat over an optimum plateau near 5–10°C and falls above 15–20°C, to 19 log points below the plateau at a daily mean of 30°C; the hot tail is estimated stably across the grid, within –19.3 to –16.4 log points at 30°C relative to each specification’s own maximum; the cold branch is not: few country-sector-years carry mass below 5°C, and the plateau’s location is a parametric extrapolation with wide intervals. The annual-mean response is hump-shaped with its optimum near 12°C and falls away towards cold annual means. The two functions differ most over the cold range, and by an order of magnitude less over 20–30°C. Where they disagree is therefore the winter tail, the part of the distribution the annual mean cannot represent; what that disagreement does to the projected North is the next section.

2.4 Amenities are not identified from bilateral migration

The same logic applies to amenities: the bilateral migration equation of the model, with the pair effect absorbing fundamental amenities and migration costs, is estimated by PPML on the




Figure 1: Productivity dose-response under the two exposure metrics. Grey: the 15 specifications of the grid. Coloured: the specification carrying the Akaike weight (a natural spline with 4 degrees of freedom, winsorised at 0.5% in (a), a natural spline with 3 degrees of freedom, winsorised at 0.5% in (b)), with the median, Q25–Q75 and Q05–Q95 of 300 Monte Carlo draws of its coefficients. Curves are shifted to pass through zero at 0°C. Bars: the arm’s own exposure, share of country-year days in (a) and of countries in (b).

Abel and Cohen (2019) five-year flows. The temperature terms are jointly insignificant in every cell of the amenity grid under either metric (daily: $p = 0.074$ – 0.161 ; annual: $p = 0.138$ – 0.413), and a placebo on lead differentials is not rejected. That null is uninformative: the level of f_a is not identified, since the pair effect absorbs the time-invariant differential, and what survives it, the within-pair change over five years, is too small for the test to have power against the alternatives of interest. I therefore bound f_a rather than rely on its estimate: monotonicity and the restriction that non-market damages do not exceed market ones give the identified set $[0, \hat{f}_z]$, which the migration estimate exceeds by a factor of 10.3. The body runs the counterfactual at the lower bound, $f_a = 0$, so that welfare flows through productivity alone; Supplement D.3 re-runs it at the estimate, a stress test beyond the admissible bound, which enlarges the global loss and deepens the Northern contrast.

2.5 The contrast between metrics

The two dose-responses of Figure 1 are estimated on the same cells with the same controls, so the natural test of whether the exposure metric matters is to project each onto the same warming

and compare. Figure 2 does this for the productivity fundamental of the North above 45°N . For each of the 15 specifications of the grid it reports the projected change by 2100 under SSP2-4.5 when exposure is the annual mean, when it is the daily distribution, and the difference between the two with its 95% interval. Nothing is solved in general equilibrium here: the projection integrates each estimated function against the change in each cell's exposure, population-weighted, and isolates what the estimation delivers before the model reallocates anything.




Figure 2: Projected change in the Northern log productivity fundamental above 45°N by 2100 under SSP2-4.5, in partial equilibrium, for each of the 15 specifications. Top: annual mean (blue) and daily distribution (red), with 95% intervals. Bottom: their difference, with the interval (clustered by origin and destination) carrying the covariance (black) or treating the arms as independent (grey).

Three facts are read off the figure. First, in the top panel, under the annual mean the Northern change runs from +0.012 to +0.024 log points, and the intervals exclude zero in 14 of the 15 specifications. A literature that measures exposure by annual means will find a precise positive

Northern response under almost any form it tries.

Second, still in the top panel, under the daily distribution the same specifications run from -0.001 to $+0.003$, with a median of $+0.002$, with 1 below zero and none excluding it; their intervals together span $[-0.0081, +0.0140]$. The daily arm is therefore consistent with a Northern gain, with a Northern loss and with no response at all, and what it establishes is that the sign of the direct Northern response is not identified at this resolution.

Third, in the bottom panel, the difference between the two arms excludes zero in 15 of 15 specifications, which is a stronger statement than the first two would give on their own.

The two arms are estimated on the same observations, and estimates drawn from the same data move together. A sample that happens to make the annual estimate larger tends to make the daily estimate larger as well, so that the gap between them is more stable across samples than either level; formally, the variance of a difference subtracts twice the covariance, and here the covariance is positive, with a correlation between the two coefficient vectors of 0.50 to 0.75 across the grid. Carrying that covariance, the standard error of the difference is 53 to 74% of what it would be if the two arms were treated as independent. Treating them as independent therefore overstates the uncertainty of the difference, and it is the conservative calculation: under it, the difference still excludes zero in 12 of 15 specifications, the exceptions being the three degree-two polynomials. Discounting the covariance entirely would still find the difference signed across most of the grid.

The same erosion appears in a design that shares neither the data, nor the unit, nor the fundamental of the bilateral estimator. Supplement E re-runs the contrast inside the estimation of Cruz and Rossi-Hansberg (2024) on fundamentals inverted from their model (a cell of one degree across four decadal cross-sections, under their sub-region $\times$ period fixed effects, their controls and their estimation sample) using my own inversion rather than their deposited series, so that inversion, estimation and projection stay inside one reproduction, changing only the exposure metric. Under the annual mean the projected Northern productivity change is significantly positive in six cases out of six; under the daily distribution it is significantly positive in one case out of six, and for

the Northern aggregate the annual-mean projection exceeds the daily one by +0.443 in the log fundamental $[+0.238, +0.651]$, the daily projection being no longer distinguishable from zero. An omitted confound specific to bilateral trade cannot account for a loss of identification in a design built on inverted cell fundamentals.

3 General-Equilibrium Illustration

Section 2.5 established what the two metrics deliver before the economy adjusts: a Northern gain under one, an unsigned response under the other. Whether that difference survives once trade and migration reallocate activity, and what it amounts to in welfare, is a question the estimates cannot answer on their own. This section propagates each dose-response through the model of Section 2.1 along the same emissions pathway and reports the difference.

3.1 Scenarios and data

I simulate the model along three paths that share the same exogenous population and productivity trajectories and differ only in climate. The no-climate-change baseline holds exposure at its 2015 value, so that the climate multiplier in relative changes equals one in every period; Scenario 1 applies the annual-mean pair $(\mathcal{E}^{\text{ann}}, f_z^{\text{ann}})$; Scenario 2 applies the daily-distribution pair $(\mathcal{E}^{\text{daily}}, f_z^{\text{daily}})$; a third scenario applies the daily response to annual-mean exposure, $(\mathcal{E}^{\text{ann}}, f_z^{\text{daily}})$, and serves the decomposition of Section 3.3. The scenarios are solved on one specification per arm: for the daily response, a natural spline with 4 degrees of freedom, winsorised at 0.5%; for the annual one, a natural spline with 3 degrees of freedom, winsorised at 0.5%. The reported object for a cell is the cumulative log change of period utility relative to the baseline at 2100, $\sum_{s \leq t} \log[\dot{u}_s^{\text{scen}}(r)/\dot{u}_s^{\text{base}}(r)]$, aggregated with population weights, and the daily extra is the difference between the two scenarios. It is a level effect on regional utility at the horizon, not a growth effect (Nath et al., 2024).

The grid has 13,986 inhabited 1° cells and 17 five-year periods. Cells outside the Cruz and

Rossi-Hansberg (2024) grid, 1.9% of world population, are excluded (Supplement A). Initial trade shares come from Cruz and Rossi-Hansberg (2024) trade-cost matrix, initial migration shares from Abel and Cohen (2019) and Kummu et al. (2018), and population and productivity paths from SSP2 (KC et al., 2024). Productivity growth is the average of the IAMC models available for each country under SSP2 — OECD Env-Growth (Dellink et al., 2017), IIASA GDP and PIK GDP-32 (Koch and Leimbach, 2023). Local productivity is proxied from observed GDP per capita and set to zero in cells without positive productivity and wage or with at most 1 inhabitant per cell, so that the model is never inverted: activity can reallocate across inhabited cells but not into empty ones. This closes the extensive margin of adaptation on which part of the Northern gains rests, and it is conservative against the contrast examined here. Future temperatures are the ensemble mean of the GFDL-ESM4, IPSL-CM6A-LR, MPI-ESM1-2-HR and MRI-ESM2-0 CMIP6 Earth System Models under SSP2-4.5, bias-adjusted and downscaled to 0.5° against the W5E5 observational dataset by ISIMIP3b (Lange, 2019). What the daily metric takes from them is the distribution of daily temperatures pooled over each five-year period, which is a climatological object rather than a weather trajectory; where regional extremes are misrepresented, the assessed direction of the bias is an understatement of hot extremes (Vautard et al., 2023).

Every magnitude below is a difference between two scenarios run on the same initial state, the same solver and the same iteration path. Whatever the calibrated initial state carries enters both scenarios and cancels to first order in the difference, and I do not report levels. The only uncertainty propagated through the model is the Monte Carlo band on the dose-responses.

3.2 What the metric governs

Table 1 reports the three margins at 2100. At the world level the daily distribution adds losses to the annual mean rather than redistributing them: the daily extra is -0.010 on top of an annual total of -0.008 , and it is negative at every horizon, so that by 2100 it adds losses equal to 124.7% of the annual total. The share of the world population living in cells that gain from warming falls

from 11.7% under the annual mean to 0.1% under the daily distribution.

Above 45°N the annual-mean gain of +0.0115 becomes −0.0063. The daily extra of −0.0178 is negative for 99.5% of the 662 million people in that aggregate, and in the central run the daily total is negative at every latitude threshold of the grid, and within rounding of zero above 66.5°N, where 2.1 million people live (Section 3.4 reports the draws); the second line under each unit in Table 1 shows that the loss holds across the Monte Carlo band above 45°N and not above 66.5°N. Among the cold economies, the annual mean projects a gain in 6 of 7 and the daily extra is negative in 7 of 7, so that every gain the annual mean projects is removed by the daily distribution. What separates these economies is the size of the gain the daily distribution removes, not whether it is removed: Russia enters with +0.025 and carries a daily extra of −0.028, while Germany enters with +0.004 and needs only −0.011.

Figure 3 shows the geography: the annual-mean panel's beneficiaries are among the daily-total panel's losers, and the daily extra is deepest where the annual-mean gain is largest.

Table 1: General-equilibrium welfare effects by 2100, cumulative population-weighted $\Delta \log U$, amenities at the lower bound $f_a = 0$. Daily total and Annual total are each arm against the no-climate baseline; Daily extra is their difference, computed before rounding. Second line under each unit: 5th–95th percentile range and share of negative outcomes of the daily total across 60 Monte Carlo draws of the daily dose-response, each solved in general equilibrium (Table 16).

	Daily total	Annual total	Daily extra
<i>Panel A. World</i>			
World	−0.017	−0.008	−0.010
5th–95th pct.	[−0.029, −0.005]	negative in 60 of 60 draws	
<i>Panel B. By latitude</i>			
≥ 45°N	−0.006	+0.011	−0.018
5th–95th pct.	[−0.011, −0.003]	negative in 60 of 60 draws	
≥ 55°N	−0.003	+0.023	−0.027
5th–95th pct.	[−0.008, +0.0002]	negative in 56 of 60 draws	
≥ 66.5°N	0.000	+0.033	−0.033
5th–95th pct.	[−0.005, +0.004]	negative in 29 of 60 draws	
<i>Panel C. Cold-climate economies</i>			
Russia	−0.003	+0.025	−0.028
5th–95th pct.	[−0.008, +0.0004]	negative in 55 of 60 draws	
Finland	−0.001	+0.027	−0.028
5th–95th pct.	[−0.006, +0.004]	negative in 38 of 60 draws	
Norway	−0.005	+0.013	−0.018
5th–95th pct.	[−0.009, −0.001]	negative in 59 of 60 draws	
Sweden	−0.004	+0.012	−0.017
5th–95th pct.	[−0.009, −0.0002]	negative in 57 of 60 draws	
Canada	−0.010	+0.011	−0.021
5th–95th pct.	[−0.016, −0.004]	negative in 60 of 60 draws	
Germany	−0.007	+0.004	−0.011
5th–95th pct.	[−0.011, −0.003]	negative in 60 of 60 draws	
France	−0.012	−0.008	−0.004
5th–95th pct.	[−0.017, −0.008]	negative in 60 of 60 draws	
Population in gaining cells (%)	0.1	11.7	—

Figure 4 reports the three margins by population-weighted quartile of 1990 GDP per capita. The daily total loads heaviest on the poorest quartile (−0.024 against −0.011 in the richest) but aggregating exposure understates damages at both ends of the income distribution: in the poor hot economies through the hot tail, and in the rich cold ones through the removal of the annual-mean gain. The band on the daily total is negative at every quartile.




Figure 3: The three maps show the country-aggregate, population-weighted cumulative $\Delta \log U$ by 2100 under the daily total, the annual total and the daily extra, with amenities at the lower bound. The colour scale is common to the three panels and centred at zero.




Figure 4: Welfare effects by 2100 by quartile of 1990 GDP per capita (population-weighted), cumulative $\Delta \log U$, amenities at the lower bound. Bars: daily total, annual total, daily extra; segments on the daily total: 5th–95th percentile range across 60 Monte Carlo draws of the daily dose-response solved in general equilibrium.

3.3 Where the gap comes from

The two scenarios compute different objects. Scenario 1 evaluates a function estimated on annual means at the cell’s annual mean, $f_z^{\text{ann}}(\bar{T})$, Scenario 2 a function estimated on daily bins over the distribution the cell experiences, $\int f_z^{\text{daily}} dF$. Their difference separates into two terms,

$$\underbrace{f_z^{\text{ann}}(\bar{T}) - f_z^{\text{daily}}(\bar{T})}_{\text{different function}} + \underbrace{f_z^{\text{daily}}(\bar{T}) - \int f_z^{\text{daily}} dF}_{\text{different evaluation}}.$$

The second term is Jensen’s inequality, which evaluates one function at a point rather than integrating it. The first term compares two different functions at the same argument, and it is what aggregating exposure before estimating produces. The third scenario, $(\mathcal{E}^{\text{ann}}, f_z^{\text{daily}})$, holds the function fixed and changes only the evaluation. At the Northern aggregate, of a daily extra of -0.0178 , the evaluation term contributes $+0.0034$ and the function term -0.0211 . The Jensen term is of

opposite sign to and 6.3 times smaller in absolute value than the function term. The mechanism is that the annual-mean function assigns to a warming cold country a benefit that the daily function does not find in the same country's record. Finland's population-weighted mean rises from 5.3 to 7.2°C by 2100; the annual function returns a productivity gain of +0.036 in log productivity and the daily function +0.009: both functions leave Finland better off, and they disagree by a factor of 4.2 on how much. The country still spends 10 days a year below -10°C at the end of the projection and its days above 25°C rise from 0.5 to 2.0. Neither tail is visible at the mean.

Put plainly, warming a cold country removes freezing days and adds hot days, and a regression on days finds both to be costly. What it does not find is the value the annual mean places on the freezing days: over the cold range the annual-mean function values the same loss of cold days at least 1.8 times what a regression on days values it at. Measured by the annual mean, the same country simply becomes a warmer country, and is credited with whatever a warmer mean has been worth in the data: across countries or across cells, the income gap between cold and temperate ones; across years, a milder winter in the same place. These are the two sources of variation an annual-mean regression can use. The estimator of Section 2.3 uses the first across countries; the design of Supplement E uses it across cells within a sub-region and period, with 0.6% of the identifying variance of the mean temperature lying within cells across periods; and the Northern gain erodes under both. The second source is the one panel designs use, and the one Jones et al. (2026) show to carry warming-induced trends (Supplement B.3). That credit is not the price of any particular day, and a regression on days cannot reproduce it.

The gap can also be located on the temperature support without solving the model. Integrating each estimated function against the change in each Northern cell's daily exposure and attributing the difference to the temperature bins in which it arises, 70.1% of the Northern contrast comes from the two functions differing at the same temperature rather than from the two metrics integrating against different objects; of that share, 76% lies below 0°C and 8% above 20°C . The two functions disagree about the value of the freezing days that warming removes, and they agree about the cost of the hot days it adds. The two shares are not comparable: the general-equilibrium split

divides a propagated welfare difference, the local one divides the partial-equilibrium contrast of the estimated functions.

3.4 Why the Northern gain does not survive equilibrium

The daily metric does not sign the Northern response in partial equilibrium. The equilibrium loss above 45°N is nonetheless not a coin flip, because the Northern daily total also carries what warming elsewhere does to the North. Two further pairs of runs measure this: in the first, warming outside the North is switched off and only the cells at or above 45°N warm; in the second, warming is applied only outside the North (Supplement D.5). Warming outside the North accounts for 68.9% of the Northern daily total. Warming in the North alone changes Northern welfare by +0.014 under the annual-mean function and by −0.002 under the daily one, so the reversal of sign happens with warming confined to the North. What warming elsewhere does to the North is nearly the same under the two functions and largely cancels in the daily extra: with warming confined to the North, the Northern daily extra is 91.8% of its value when the whole world warms, and warming outside the North accounts for 8.1%.

What the daily distribution adds is an income loss that grows with latitude. Net of the amenity term, which moves only through congestion at $f_a = 0$, the daily extra is −0.002 in the temperate band, −0.019 between 45 and 60°N and −0.037 above 60°N (Table 18). Nearly all of it is produced by warming in the North itself rather than imported from elsewhere, and it comes from the central draw of the dose-response.

Across 60 Monte Carlo draws of the daily dose-response, each propagated through the model, the Northern daily total is negative in 60 of 60 draws, with a median of −0.006 and a 5th–95th percentile range of −0.011 to −0.003 (Table 1). The loss is a property of the Northern aggregate rather than of its coldest members: above 66.5°N the daily total is negative in 29 of 60 draws, and for Finland in 38 of 60, so that the sign of the far North is not determined at daily resolution, in general equilibrium as in the estimation.

4 Conclusion

There are many reasons to expect climate change to be costly in the North, for instance biodiversity collapse, permafrost thaw under roads, pipelines and buildings, or boreal wildfires. This paper has set those channels aside and asked a narrower question, about the one channel the spatial models of climate damages do quantify, the effect of temperature.

Spatial models of climate damages project welfare gains for the world's cold economies. This paper has argued that those gains are a property of the exposure metric. Estimated on the same data with the annual mean and with the distribution of daily temperatures, the productivity response projects a precise Northern gain under the first and a change centred on zero under the second, and the difference between the two is signed in every specification of the grid; propagated through a spatial equilibrium model of trade and migration, the annual-mean gain of the population above 45°N becomes a loss, and the share of the world population living in cells that gain from warming falls from 11.7% to 0.1%. The reason is not that a non-linear function is evaluated at the wrong point. It is that aggregating exposure before estimating recovers a different function, one that absorbs everything separating cold countries from hot ones and assigns to the freezing days that warming removes a value that a function estimated on days does not find. This is a statement about the productivity response, not about daily resolution as such: where the cold arm of the dose-response is real, as for mortality, daily-resolution estimates project gains for cold regions (Carleton et al., 2022), and nothing here contradicts them. What the annual mean does to productivity is to manufacture a cold arm the daily data do not support; what Jones et al. (2026) show for panels is that extreme-day counts can manufacture one too.

The exercise takes the emissions pathway as given, and the limit is a missing object rather than a conceptual one. The emulators embedded in models with endogenous emissions map cumulative emissions to pattern-scaled regional annual means and produce nothing at the daily resolution on which welfare depends. Emulators at monthly resolution now exist (Nath et al., 2022), and they recover the seasonal amplitude that annual means discard; the object this paper requires is one step further, a reduced form from cumulative emissions to the shape of the regional daily distribution.

Until it exists, endogenous-climate and distributional-climate analyses sit on separate branches of the literature, and this paper establishes what the choice between them costs: the branch that buys endogenous emissions at the price of the annual mean buys it at the price of the sign.

References

- Abel, G. J. and J. E. Cohen (2019). Bilateral international migration flow estimates for 200 countries. *Scientific data* 6(1), 82.
- Astier, J., G. Barrows, R. Cael, and H. Ollivier (2026). Quantifying climate damages when regions trade: A structural gravity approach. GRI Working Paper 447, Grantham Research Institute on Climate Change and the Environment, London School of Economics.
- Bilal, A. and D. R. Känzig (2026). The macroeconomic impact of climate change: Global versus local temperature. *The Quarterly Journal of Economics* 141(2), 889–944.
- Blackport, R. and J. C. Fyfe (2024). Amplified warming of north american cold extremes linked to human-induced changes in temperature variability. *Nature Communications* 15(1), 5864.
- Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature* 527(7577), 235–239.
- Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies* 82(1), 1–44.
- Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, 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.
- Cattaneo, C., M. Beine, C. J. Froehlich, D. Kniveton, I. Martinez-Zarzoso, M. Mastrorillo, K. Milllock, E. Piguet, and B. Schraven (2019). Human migration in the era of climate change. *Review of Environmental Economics and Policy* 13(2), 189–206.
- Cruz, J.-L. and E. Rossi-Hansberg (2024). The economic geography of global warming. *Review of Economic Studies* 91(2), 899–939.
- Dellink, R., J. Chateau, E. Lanzi, and B. Magné (2017). Long-term economic growth projections in the shared socioeconomic pathways. *Global Environmental Change* 42, 200–214.
- Deschenes, O. and M. Greenstone (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review* 97(1), 354–385.
- Desmet, K., D. K. Nagy, and E. Rossi-Hansberg (2018). The geography of development. *Journal of Political Economy* 126(3), 903–983.

- Desmet, K. and E. Rossi-Hansberg (2015). On the spatial economic impact of global warming. *Journal of Urban Economics* 88, 16–37.
- Desmet, K. and E. Rossi-Hansberg (2024). Climate change economics over time and space. *Annual Review of Economics* 16(1), 271–304.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier.
- Hersbach, H., B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* 146(730), 1999–2049.
- Jones, B. F., J. Moscona, B. A. Olken, and C. von Dessauer (2026, January). With or without u? binning bias and the causal effects of temperature extremes. Working Paper 34671, National Bureau of Economic Research.
- KC, S., M. Dhakad, M. Potančoková, S. Adhikari, D. Yildiz, M. Mamolo, T. Sobotka, K. Zeman, G. Abel, W. Lutz, et al. (2024). Updating the shared socioeconomic pathways (ssps) global population and human capital projections. IIASA Working Paper WP-24-003, International Institute for Applied Systems Analysis.
- Koch, J. and M. Leimbach (2023). Ssp economic growth projections: Major changes of key drivers in integrated assessment modelling. *Ecological Economics* 206, 107751.
- Krusell, P. and A. A. Smith, Jr. (2022). Climate change around the world. Working paper.
- Kummu, M., M. Taka, and J. H. Guillaume (2018). Gridded global datasets for gross domestic product and human development index over 1990–2015. *Scientific data* 5(1), 1–15.
- Lange, S. (2019). Trend-preserving bias adjustment and statistical downscaling with isip3basd (v1.0). *Geoscientific Model Development* 12(7), 3055–3070.
- Lemoine, D., C. Hausman, and J. G. Shrader (2025). A guide to climate damages. Technical report, National Bureau of Economic Research.
- Nath, I. B., V. A. Ramey, P. J. Klenow, et al. (2024). *How much will global warming cool global growth?*, Volume 32761. National Bureau of Economic Research Cambridge, MA.
- Nath, S., Q. Lejeune, L. Beusch, S. I. Seneviratne, and C.-F. Schleussner (2022). Mesmer-m: an earth system model emulator for spatially resolved monthly temperature. *Earth System Dynamics* 13(2), 851–877.
- Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics* 37(2), 187–204.

- Patterson, M. (2023). North-west europe hottest days are warming twice as fast as mean summer days. *Geophysical Research Letters* 50(10), e2023GL102757.
- Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics* 9, 21–58.
- Rudik, I., G. Lyn, W. Tan, and A. Ortiz-Bobea (2022). The economic effects of climate change in dynamic spatial equilibrium. Working paper, 12 July 2022.
- Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences* 106(37), 15594–15598.
- Vautard, R., J. Cattiaux, T. Hap  , J. Singh, R. Bonnet, C. Cassou, D. Coumou, F. D’andrea, D. Faranda, E. Fischer, et al. (2023). Heat extremes in western europe increasing faster than simulated due to atmospheric circulation trends. *Nature Communications* 14(1), 6803.

Appendix

Supplement A documents the calibration of the initial equilibrium and the climate inputs. Supplements B and C each treat one estimated dose-response function end to end. Supplement D reports the general-equilibrium detail. A final section E sets the same question in an independent design (Cruz and Rossi-Hansberg, 2024).

A Calibration and Climate Inputs

The counterfactual evaluates damages along an exogenous emissions path; it does not optimise emissions. Its calibration targets a 2015 initial state in levels and geography from published sources, and lets the dynamics run in relative changes from there. This section documents the initial state, the reason the model is not inverted, and what is taken from the climate projections.

A.1 Data

The world is discretised on a 1° grid of 17,048 potential cells, of which 13,986 have positive productivity, positive wage and population above one; the rest (oceans, deserts, ice) carry no economic content. 916 further inhabited cells, holding 142.5 million people or 1.9% of world population, lie outside the Cruz and Rossi-Hansberg (2024) grid: 39 countries and territories, 18.2 million people in all, are outside it entirely, and 100 countries in part. The 7.6 million international migrants of the 2010–2015 window whose origin or destination is one of these cells are dropped with them. Cells are mapped to countries through GHSL population shares within GADM polygons. The 1° resolution is that of the Cruz and Rossi-Hansberg (2024) trade-cost matrix. Fundamental productivity is proxied by GDP per capita (gridded GDP from Kummu et al., 2018, GHSL population, 2015), winsorised at the upper 5% tail; the wage is that proxy times the ILO 2015 country labour share (regional imputation, then global mean, where missing). Initial trade shares follow $\lambda_0(rs) \propto Z(s)^\theta w(s)^{-\alpha\theta} \tau_{rs}^{-\theta}$ with $\theta = 6.5$ (Desmet et al., 2018) and $\alpha = 0.6$ (Desmet and Rossi-Hansberg, 2015), and the Cruz and Rossi-Hansberg (2024) cost matrix, rows summing to one over

sources. At ≈ 100 km resolution this grid is close to autarky: 97.1% of cells spend more than 0.99 of their income on their own output, and the mean import share is 0.4%. GDP per capita is not an Eaton–Kortum fundamental, so the raw allocation (Z, w, L, τ) need not clear exactly. Because the model is solved in relative changes, only λ_0 enters, and any residual adjustment transient is common to the three scenarios and differences out of every reported contrast.

The five-year transition matrix μ_0 combines international bilateral flows from [Abel and Cohen \(2019\)](#) (2010–2015), allocated to cells by within-country population shares, and internal flows reconstructed from the [Kummu et al. \(2018\)](#) net migration stocks (2010–2014) by an outer product of origin and destination shares. Within each country, internal flows are allocated across the cells in which it holds at least half of the population, origins weighted by the country’s population in the cell and destinations by the positive part of the net migration stock that the international flows leave there; the cross-border mass this leaves is 5.73 million over the window, 6.4% of the internal block. Diagonals are zeroed, stayers added, rows normalised to transition probabilities, flows capped by population. Three limitations: the within-country allocation assumes no sub-national distance friction; the internal gross structure is reconstructed from net stocks, not observed; and the same matrix is applied at each of the 16 transitions, so that the refugee-driven flows of the 2010–2015 window (Pakistan, Afghanistan, Syria) recur throughout. [Figure 5](#) shows net losses from the tropics to about 35°N and net gains over the temperate belt.




Figure 5: Net migration by latitude from the calibrated initial matrix.

Spatial models of climate damages recover unobserved fundamentals by inverting the model on an observed equilibrium. The inversion assigns a positive fundamental to every cell, including empty ones, so that under warming activity can relocate into them and the gain is counted as adaptation. This extensive margin is part of what generates the Northern gains in that literature, and two facts argue against relying on it. Observed international net migration over 1990–2020, a period of roughly 0.6°C of warming, shows no clear accelerating northward trend (Figure 6): legal barriers, income gaps and networks have dominated the climate signal at these horizons (Cattaneo et al., 2019). And the cells into which the inverted margin routes activity are largely cells with no activity today (Figure 7); their emptiness has causes (institutions, infrastructure, property rights, labour markets) that warming does not remove. I therefore proxy productivity from observed data, set it to zero in cells without positive productivity and wage or with at most 1 inhabitant per cell, and let adaptation run only on the intensive margins the data record. Re-opening the extensive margin would add cold high-latitude destinations for displaced activity, and so amplify the annual-mean Northern gain; the restriction is conservative against the contrast the paper examines.




Figure 6: Net international migration stock by latitude band and five-year period, 1990–2020, from [Abel and Cohen \(2019\)](#). Countries are assigned to bands by the latitude of their centroid; the five bands sum to zero in each period.




Figure 7: Areas of zero recorded economic activity in 2015 (blue). Model inversion assigns positive fundamental productivity to these cells.

A.2 Model: equilibrium conditions

Workers are freely mobile at a bilateral cost. A worker in s who locates in r obtains $u_t(r) m(sr)^{-1} \epsilon$, where $u_t(r) = a_t(r) y_t(r)$ is spatial welfare, $y_t(r)$ real income per capita, $m(sr)$ a time-invariant and symmetric migration cost, and ϵ an i.i.d. Fréchet shock with dispersion Ω ; symmetry makes the location problem static (Desmet et al., 2018). The share of population of s that locates in r is

$$\mu_t(sr) = \frac{[u_t(r)/m(sr)]^{1/\Omega}}{\sum_n [u_t(n)/m(sn)]^{1/\Omega}}, \quad (3)$$

and population evolves with exogenous births and deaths and these flows. Each cell produces a continuum of varieties under perfect competition with labour and structures, with a labour share α and productivity draws from a Fréchet distribution of scale $Z_t(r)^\theta$ and shape θ ; goods are traded at iceberg costs τ_{sr} with balanced trade. With $x_t(r)$ the unit cost of the input bundle, the share of r 's expenditure on goods from s and the price index are

$$\lambda_t(rs) = \frac{Z_t(s)^\theta (x_t(s) \tau_{sr})^{-\theta}}{\sum_l Z_t(l)^\theta (x_t(l) \tau_{lr})^{-\theta}}, \quad P_t(r) = \Gamma \left[\sum_s Z_t(s)^\theta (x_t(s) \tau_{sr})^{-\theta} \right]^{-1/\theta}. \quad (4)$$

Rents on structures accrue to local residents, so that income per capita is $w_t(r)/\alpha$ and real income is $y_t(r) = w_t(r)/(\alpha P_t(r))$; the goods market clears when $w_t(r) L_t(r) = \sum_s \lambda_t(sr) w_t(s) L_t(s)$.

The model is solved in relative changes, $\dot{x}_{t+1} = x_{t+1}/x_t$. Because migration costs and structures are time-invariant, the system (3)–(4) in changes involves only the observed shares of the previous period, the elasticities, and the changes in fundamentals implied by (1):

$$\dot{\mu}_{t+1}(sr) = \frac{\dot{u}_{t+1}(r)^{1/\Omega}}{\sum_n \mu_t(sn) \dot{u}_{t+1}(n)^{1/\Omega}}, \quad \dot{P}_{t+1}(r) = \left[\sum_s \lambda_t(rs) \dot{Z}_{t+1}(s)^\theta \dot{x}_{t+1}(s)^{-\theta} \right]^{-1/\theta}, \quad (5)$$

with $\dot{Z}_{t+1}(r) = \gamma_{t+1}(r) \exp[\mathcal{E}_{t+1}(r) f_z - \mathcal{E}_t(r) f_z]$ for exogenous growth γ , and for fundamental amenities, $\dot{a}_{t+1}(r) = \dot{L}_{t+1}(r)^{-\lambda} \exp[\mathcal{E}_{t+1}(r) f_a - \mathcal{E}_t(r) f_a]$; the changes in trade shares, wages and utility follow from (4), market clearing and $\dot{u} = \dot{a} \dot{y}$. Each transition is solved as a fixed

point on next-period population, with a nested fixed point on wages and trade shares. Unobserved fundamentals $\bar{Z}$ and $\bar{a}$ and costs m and τ enter only through the observed initial shares λ_0 and μ_0 ; they are never recovered by inversion. That choice closes the extensive margin of adaptation: a cell with no observed activity has no share to scale, and no counterfactual can route activity into it.

B The Productivity Dose-Response

This section collects the evidence behind the productivity dose-response of the body: the sample and first stage, the identification argument and its tests, the robustness of the estimated function across specifications, and the contrast between the two exposure metrics across the same grid.

B.1 Sample and first stage

The estimation combines BACI–CEPII bilateral trade flows (1995–2019) with GLORIA input–output tables for domestic absorption, World Bank GDP per capita, and ERA5 temperatures (Hersbach et al., 2020) aggregated to countries with GHSL-POP population weights. Daily means are computed by the data provider over the UTC calendar day from hourly fields, so that for cells far from the Greenwich meridian the day boundary falls within local daytime. Three exclusions reduce the ISIC Rev.3 nomenclature to 20 manufacturing and agricultural sectors: primary and extractive industries (divisions 02, 05, 10–14), whose productivity is driven by endowments and whose domestic absorption is near zero, which destabilises the normalisation; services (divisions 40 and above), for which $X^k(rs)/X^k(rr)$ is not defined in the GLORIA final-demand matrix; and divisions 18, 30 and 33, recorded by BACI but absent from GLORIA. A fourth exclusion follows from the sectoral elasticities: division 99, goods not elsewhere classified, which the concordance populates almost entirely with waste, scrap and used goods, has no elasticity in Caliendo and Parro (2015) and is dropped. Divisions 18, 30 and 33 carry 11.87% of trade lines and 9.08% of trade by value, measured on the product-level BACI lines that remain once division 99 and non-country entities are removed, and 191 country pairs lose complete year coverage. Before the exclusions the

sample has 3,730,958 observations; 219,480 of them belong to pairs involving one of 21 countries without an ERA5 record and are dropped, and the sample on which the dose-responses are estimated has 3,511,478. Bilateral trade costs are unobserved and proxied by the gravity vector C_{rs} : log distance, common official language, colonial link, and a WTO free-trade-agreement indicator. A spline’s interior knots sit at equally spaced quantiles of the exposure distribution its arm uses (the day-weighted distribution of daily means under one metric, the distribution of annual means under the other) and its boundary knots at the winsorisation thresholds. Monte Carlo draws are $\beta^{(b)} \sim \mathcal{N}(\hat{\beta}, \hat{\Sigma})$; the joint and non-linearity tests across this grid are in Supplement B.4.

B.2 Identification

Between-pair variation. Table 2 quantifies the statement of the body that bilateral temperature differentials are quasi-permanent: origin×destination×product effects absorb 97.7–99.7% of the variance of the climate terms Δ_k , and the within-pair standard deviation of Δ_k is at most 15.2% of its total. Product×year effects absorb at most 0.41%. Table 3 shows what the gravity controls do against the remaining between-pair variation: they absorb 4.9–14.6% of it, conditional on product×year effects, without being collinear with the climate term.

Table 2: Variance absorption by fixed effects

Climate term	Pair×prod FE	Prod×year FE	Both FE
Δ_1	0.9969	0.004125	0.9969
Δ_2	0.9942	0.002871	0.9942
Δ_3	0.9768	0.003096	0.9769

R^2 from projecting the bilateral climate differential $\Delta_k = g_k(T(s)) - g_k(T(r))$ on fixed effects alone (OLS). Origin×destination×product fixed effects absorb at least 97.7% of the variance in bilateral climate differentials (97.7%–99.7% across the three terms), leaving insufficient residual variation for PPML identification. Product×year fixed effects absorb at most 0.41%, preserving the between-pair variation that identifies $g(\cdot)$.

Table 3: Absorption of between-pair confounders by gravity controls

Climate term	R^2 (FE only)	R^2 (FE + controls)	Absorbed
Δ_1	0.004125	0.1498	0.1456
Δ_2	0.002871	0.1415	0.1386
Δ_3	0.003096	0.0523	0.0492

OLS projection of Δ_k on gravity controls conditional on product $\times$ year fixed effects. Gravity controls (log bilateral distance, common language, colonial link, WTO FTA, $\theta^k \cdot \log(y_s/y_r)$) absorb 4.9–14.6% of the between-pair climate variation, partially controlling for the main observable confounders.

Selection on unobservables. Table 4 reports the Oster (2019) selection ratio δ^* across the specification grid, with the first-basis coefficient without and with gravity controls. Because δ^* scales with the inverse of the coefficient’s movement, specifications whose coefficient barely moves produce arbitrarily large values of uninformative sign; the informative entry is the least favourable entry on the grid, $\delta^* = 2.2$, whose coefficient moves by 89.7% when the controls are added. Unobservables would have to be 2.2 times as important as the gravity controls to explain the estimated effect away.

Table 4: Oster (2019) selection-on-unobservables robustness for the productivity dose-response function. δ^* : degree of selection on unobservables relative to gravity controls needed to explain away the estimated climate effect ($\beta^* = 0$), computed with $R_{\max} = \min(1, 1.3 \times \tilde{R})$.

Specification	Winsorization	δ^*	$\dot{\beta}_1$	$\tilde{\beta}_1$	$\dot{R}$	$\tilde{R}$	p (no ctrl.)	p (ctrl.)
Polynomial 2	[0.5%, 0.5%]	-24.90	0.1431	0.1538	0.0270	0.0561	0.000	0.000
Polynomial 3	[0.5%, 0.5%]	18.16	0.3812	0.3479	0.0271	0.0565	0.000	0.000
Natural spline 3	[0.5%, 0.5%]	-102.57	0.0838	0.0852	0.0270	0.0558	0.000	0.000
Natural spline 4	[0.5%, 0.5%]	31.92	0.4050	0.3847	0.0278	0.0562	0.000	0.000
Polynomial 2	[1.0%, 1.0%]	-6.91	0.0413	0.0559	0.0252	0.0550	0.000	0.000
Polynomial 3	[1.0%, 1.0%]	14.05	0.2648	0.2346	0.0255	0.0558	0.000	0.000
Natural spline 3	[1.0%, 1.0%]	-37.02	0.0526	0.0553	0.0245	0.0538	0.000	0.000
Natural spline 4	[1.0%, 1.0%]	32.56	0.3195	0.3029	0.0251	0.0540	0.000	0.000
Polynomial 2	[2.5%, 2.5%]	15.23	-0.1178	-0.1042	0.0214	0.0526	0.000	0.000
Polynomial 3	[2.5%, 2.5%]	2.17	0.0378	0.0199	0.0226	0.0545	0.000	0.000
Natural spline 3	[2.5%, 2.5%]	9.97	-0.0245	-0.0202	0.0175	0.0488	0.000	0.001
Natural spline 4	[2.5%, 2.5%]	74.88	0.1804	0.1755	0.0187	0.0496	0.000	0.002

Note. δ^* is computed for the first basis coefficient (dominant linear term). $\delta^* > 1$ means unobservables would have to be more important than the observable gravity controls to explain away the estimated effect; one is the threshold, as in the text. Because δ^* scales with $1/(\dot{\beta}_1 - \tilde{\beta}_1)$, a coefficient that barely moves when controls are added produces an arbitrarily large $|\delta^*|$ of uninformative sign: extreme values in this column indicate maximal coefficient stability, not fragility. Within-pseudo- R^2 is measured relative to a PPML model with product $\times$ year fixed effects only.

The exporter-side cost term. [Astier et al. \(2026\)](#) observe that in gravity-based estimation of local-shock elasticities the exporter-side cost term need not be orthogonal to the shock differential: wages and input prices are determined jointly with productivity and climate, so the residual is correlated with the regressors. The normalisation by domestic absorption cancels the importer's inward resistance and expenditure exactly, both indexed on r alone, but leaves the exporter capability term in place. Writing $Z_t^k(s) = \bar{Z}_t^k(s) e^{\zeta_z^\top b(T_t(s))}$, with $b(\cdot)$ the basis of the dose-response,

and collecting the non-climate exporter terms into $\xi_t^k(s) \equiv \ln \bar{Z}_t^k(s) - \ln x_t^k(s)$,

$$\ln \frac{X_t^k(rs)}{X_t^k(rr)} = \theta \zeta_z^\top \underbrace{[b(T_t(s)) - b(T_t(r))]}_{\Delta b_{rs}} + \theta [\xi_t^k(s) - \xi_t^k(r)] - \theta \Delta \ln \tau_{rs}.$$

The gravity controls target $\Delta \ln \tau_{rs}$, the income ratio targets the wage component of ξ , and ρ_t^k absorbs the product $\times$ year intercept; there is no exporter or importer effect, since one would absorb 99.7% of the differential, so the non-climate fundamental differential stays in the residual. With tildes denoting the parts orthogonal to the controls: under the specification of eq. (2),

$$\hat{\zeta}_z \xrightarrow{p} \zeta_z + [\mathbb{V}(\widetilde{\Delta b_{rs}})]^{-1} \mathbb{C}(\widetilde{\Delta b_{rs}}, \tilde{\xi}_s - \tilde{\xi}_r),$$

and by the additive exporter(+)/importer(−) structure the bias vanishes if and only if $\text{cov}(\tilde{b}(T(s)), \tilde{\xi}(s)) = 0$: climate exposure is uncorrelated with non-climate fundamental productivity, conditional on the controls and product $\times$ year effects.

Three remarks follow. (i) Hot countries plausibly have low $\bar{Z}$ for non-climate reasons, so the covariance is negative in the hot bins and damages there are overstated; this is the channel the Oster bound targets, but Oster’s argument assumes unobservables resemble the observed controls and does not by itself cover a fundamental orthogonal to the gravity set. (ii) No unit effect rescues it: within-country climate barely moves, so $b(T(s))$ is quasi-collinear with an exporter $\times$ year effect, and purging $\bar{Z}(s)$ purges f_z with it. Robustness therefore means bounding the bias, not eliminating it. (iii) The income term: $y = w/P$ is a descendant of the treatment, so dropping it induces omitted-cost bias and including it opens a collider path; I report both. It enters as a control rather than as a structural term, weighted by the sectoral trade elasticity θ^k of [Caliendo and Parro \(2015\)](#) so that it follows sectoral variation in the cost channel, with its own coefficient estimated rather than restricted. The weighting is a property of the control, not of the model of Section 2.1, which carries a single elasticity; the free coefficient absorbs the gap between y and the input bundle cost it proxies, so that this gap does not load on f_z .

Empirical bounds. I bound the bias with estimation-side re-runs only; [Astier et al. \(2026\)](#) address it by instrumenting, and nothing here substitutes for a valid instrument. Table 5 re-estimates f_z under exporter $\times$ year and importer $\times$ year effects, which absorb the exporter (respectively importer) capability term in full, and with and without the income term. Of the 4 schemes, 3 return no band excluding zero on the hot-side slope ($20 \rightarrow 30^\circ\text{C}$) and 1 does, the exporter $\times$ year scheme: 50% of the identifying differential survives the demeaning, the slope is -1.05 per $^\circ\text{C}$, and the 95th percentile of its band is -0.036 . Under importer $\times$ year effects the differential collapses (joint $F < 1$) and the scheme bounds nothing. The baseline and the scheme without the income term keep the sign of the slope but their bands reach zero, with 95th percentiles of $+0.001$ and $+0.061$. Bands rest on 3,000 draws. A second, independent seed returns the same verdict for 4 of the 4 schemes and moves the 95th percentiles to -0.023 , $+0.027$ and $+0.095$; the largest such move across the four schemes is 0.041 , within which the baseline's percentile lies and the other's does not, so the baseline band is classified neither way. The sign of the response, and of the contrast it feeds, does not rest on the exogeneity condition; what the schemes leave open is precision.

Table 5: Bounds on the productivity dose-response under asymmetric fixed effects and without the income control. Relative effect (% , normalised to each scheme’s own maximum) at reference temperatures, and the hot-side slope with its Monte-Carlo band.

	Baseline	Exporter×year	Importer×year	Without income
<i>Relative effect on productivity (% , max = 0)</i>				
5°C	-0.53	-1.41	-0.17	-0.48
10°C	-0.69	-0.00	-1.26	-0.97
15°C	-5.57	-1.83	-6.93	-6.88
20°C	-11.08	-6.84	-10.77	-13.32
25°C	-12.98	-12.76	-6.77	-15.15
30°C	-18.77	-17.36	-10.27	-20.52
Hot-side slope 20→30°C (%/°C)	-0.77 [—]	-1.05 [-2.05, -0.04]	0.05 [—]	-0.72 [—]

Note. Exporter×year fixed effects absorb the exporter capability term in full, so the response is identified off importer climate only; the importer scheme is the mirror. The last column drops bilateral income, which is a bad control: it is itself a consequence of climate. The slope is the point estimate, $(y_{30} - y_{20})/10$ from the rows above, with the $[Q_{05}, Q_{95}]$ Monte-Carlo band in brackets where the band excludes zero; where it contains zero (—) no band is reported, since it would not bound the slope. Under importer×year effects the identifying climate variation has, in addition, collapsed.

A synchronous shock. Consider a uniform shock δ to both partners’ exposure:

$$\frac{\partial}{\partial \delta} [\mathcal{E}_t(s)f_z - \mathcal{E}_t(r)f_z] = \mathcal{E}_t(s)f'_z - \mathcal{E}_t(r)f'_z, \quad (6)$$

which vanishes identically if f_z is linear and otherwise equals the difference in marginal damages between partners: a synchronous shock shifts comparative advantage towards the partner on the flatter part of the response, and that curvature is what the bin specification recovers from cross-sectional variation in exposure. The same non-linearity is what makes the annual mean an insufficient statistic, so the identification strategy and the result rest on one fact.

B.3 Binning bias

Jones et al. (2026) show that, under general warming, the number of days a location spends in each temperature bin trends at a rate that depends on its baseline temperature: a hot location gains hot days faster and loses cold days more slowly than a cold one, so that both tails trend upward in hot places relative to cold ones. In a panel with location fixed effects, identification then rests partly on those trends, and if the outcome also trends with baseline climate the estimated tails are biased, typically towards a U. The same mechanism applies to polynomials of daily temperature summed over the year. Their remedies are counterfactual expected bin counts as controls, or fixed effects interacted with short calendar windows.

The ingredients are present in this sample. Over 1995–2019 the population-weighted warming trend is faster in cold countries, the trend of each bin count correlates positively with baseline temperature in both tails and not in the centre, and 0.3–1.9% of the variance of the reference regressors Δ_k is within-pair variation over time; the within-pair trend of the dependent variable correlates with the baseline temperature differential at 0.12. The design of Section 2.3 is cross-sectional: product $\times$ year effects and the gravity controls, no pair effect, so that the bilateral differentials are identified between pairs.

Table 6 re-estimates the two reference specifications under four structures. The first is the published one. The second replaces each pair’s differentials by their 1995–2019 means, which removes every within-pair movement, the channel of the bias, without changing the fixed effects: the cold arm of the daily response moves by -0.24 of its standard error and the Northern contrast by $+0.05$, from $+0.0169$ to $+0.0172$. The third adds pair $\times$ product effects, the panel design of Jones et al. (2026): the daily cold arm becomes positive and significant, $+0.310$ [$+0.143, +0.478$] against -0.178 [$-0.396, +0.040$] in the published design, the projected Northern change turns negative under both arms (-0.0157 and -0.0265), and the contrast reverses to -0.0108 [$-0.0167, -0.0049$]. The fourth interacts those effects with five-year windows, their short-panel remedy: the cold arm falls to $+0.190$ [$+0.087, +0.292$] and the contrast to -0.0023 [$-0.0079, +0.0032$], which includes zero. The signature they describe (a cold-side effect that appears under long-panel effects and re-

cedes when the window shortens) is thus visible in these data, and the published estimate does not rest on it.

Table 6: Binning bias: the productivity dose-response under four fixed-effect structures.

Fixed effects	$\Delta_{\text{North, daily}}$	$\Delta_{\text{North, annual}}$	Contrast D	$f(-10) - f(15)$, daily
Product $\times$ year	+0.0033 [−0.0064, +0.0130]	+0.0202 [+0.0092, +0.0311]	+0.0169 [+0.0069, +0.0269]	−0.178 [−0.396, +0.040]
Product $\times$ year, pair-mean differentials	+0.0045 [−0.0056, +0.0145]	+0.0216 [+0.0106, +0.0326]	+0.0172 [+0.0067, +0.0276]	−0.204 [−0.431, +0.022]
Product $\times$ year + pair $\times$ product	−0.0157 [−0.0272, −0.0042]	−0.0265 [−0.0394, −0.0135]	−0.0108 [−0.0167, −0.0049]	+0.310 [+0.143, +0.478]
Product $\times$ year + pair $\times$ product $\times$ 5 years	−0.0105 [−0.0190, −0.0019]	−0.0128 [−0.0194, −0.0061]	−0.0023 [−0.0079, +0.0032]	+0.190 [+0.087, +0.292]

Note. Each row re-estimates the two reference specifications of Section 2.5 under the structure named in the first column; with pair-mean differentials, each pair’s bilateral temperature differentials are replaced by their means over the estimation window. Δ_{North} is the projected change in log productivity north of 45°N; D is the annual minus the daily change, with the covariance between arms; the last column is the daily response at −10°C relative to 15°C. 95% intervals in brackets, standard errors clustered by exporter and importer.

B.4 Robustness across specifications

Table 7: Robustness: productivity dose-response across specifications, daily distribution versus annual mean

Functional form	Deg.	Winsorization	p (non-linearity)	
			Daily distribution	Annual mean
Polynomial	2	[0.5%, 0.5%]	0.0093	< 0.001
Polynomial	2	[1%, 1%]	0.0105	< 0.001
Polynomial	2	[2.5%, 2.5%]	0.0138	< 0.001
Polynomial	3	[0.5%, 0.5%]	0.0261	< 0.001
Polynomial	3	[1%, 1%]	0.0283	< 0.001
Polynomial	3	[2.5%, 2.5%]	0.0327	< 0.001
Polynomial	4	[0.5%, 0.5%]	0.0501	< 0.001
Polynomial	4	[1%, 1%]	0.0557	< 0.001
Polynomial	4	[2.5%, 2.5%]	0.0651	< 0.001
Natural spline	3	[0.5%, 0.5%]	0.0015	< 0.001
Natural spline	3	[1%, 1%]	0.0013	< 0.001
Natural spline	3	[2.5%, 2.5%]	0.0010	< 0.001
Natural spline	4	[0.5%, 0.5%]	< 0.001	< 0.001
Natural spline	4	[1%, 1%]	< 0.001	< 0.001
Natural spline	4	[2.5%, 2.5%]	< 0.001	< 0.001

Note. Stage 3 PPML estimates with product×year fixed effects and bilateral gravity controls. Standard errors clustered by origin and destination. $N = 3,511,478$ in every cell, identical across the two metrics. Entries are p -values for the non-linearity test: the joint Wald test on the degree- ≥ 2 climate terms alone. Non-linearity is significant at $p < 0.05$ in 27 of the 30 specification–metric cells (15 specifications $\times$ two exposure metrics), and at $p < 0.01$ in 22 of the 30 (7 of 15 under the daily distribution, 15 of 15 under the annual mean). A separate Wald test on all climate terms jointly (linear and higher-order, which is not the quantity tabulated here) rejects at $p < 0.001$ in 30 of the 30 cells. The two metrics are estimated on the same sample, with the same dependent variable, the same fixed effects and the same controls: only the right-hand-side exposure measure differs.

Table 7 reports the joint and non-linearity tests across the grid under both metrics. The temperature terms are jointly significant throughout ($p \leq 0.0001$); linearity is rejected in 27 of the 30 specification–metric cells at 5% and in 22 at 1%; the exceptions at 5% are the three degree-four polynomials under the daily distribution (p between 0.050 and 0.065). The Akaike criterion is uninformative: one specification takes weight one to machine precision under either metric, which at this sample size is what the exponential does whatever the fit, and the criterion is a quasi-likelihood criterion under PPML.

B.5 The contrast across the grid

Table 8 is the grid behind Figure 2 of the body. For each of the 15 specifications (polynomials of degree two to four, natural splines with three or four degrees of freedom, three symmetric winsorisation thresholds, a spline’s interior knots at equally spaced quantiles of the arm’s own exposure distribution and its boundary knots at the winsorisation thresholds) it reports the projected change in the Northern log productivity fundamental above 45°N between 2015 and 2100 under each metric, their difference with its interval, and, in the last two columns, the estimated f_z at 25 and 30°C relative to the specification’s own maximum. The two arms share sample, dependent variable, fixed effects and controls. No specification is selected, ranked or weighted.

Section 2.5 reads the three counts off this table. What the body has no room for is the last two columns: the estimated response itself at two reference temperatures, each relative to its own specification’s maximum, which run from -19.3 to -16.4 log points at 30°C.

The count of signed differences is stronger than a comparison of intervals would be, because the arms are not independent. Estimated on the same cells, their coefficient vectors are correlated (ρ between 0.50 and 0.75 across the grid), and the variance of the difference carries that covariance: the standard error of the difference is 53–74% of the one obtained by treating the arms as independent. The independence calculation is therefore conservative, and under it the difference still excludes zero in 12 of 15 specifications, the exceptions being the three degree-two polynomials. The two counts are reported separately because they are different counts. The last two columns of

Table 8 show that what varies across the grid is not the sign of the hot tail, which is negative in all 15 rows, but its magnitude, from -19.3 to -16.4 log points at 30°C – a 1.2-fold spread in which the flattest retains 85% of the steepest, and the location of the cold-side plateau.

Table 8: Projected change in the northern log productivity fundamental under the annual-mean and daily-distribution exposure metrics, their difference, and the estimated dose-response behind them. The first three columns are the projected CHANGE; the last two are $\hat{f}_z$ itself at two reference temperatures, each relative to its own specification’s maximum, in log points. No specification is preferred: the table reports the grid, not a selection.

Specification	Annual	Daily	Difference	95% CI	$\hat{f}_z(25^\circ)$	$\hat{f}_z(30^\circ)$
Poly 2 [0.5%]	+0.012	+0.0010	+0.011	[+0.003, +0.019]	-10.0	-17.2
Poly 2 [1.0%]	+0.012	+0.0007	+0.011	[+0.003, +0.019]	-10.0	-17.2
Poly 2 [2.5%]	+0.012	-0.0005	+0.013	[+0.004, +0.022]	-10.1	-17.6
Poly 3 [0.5%]	+0.020	+0.0025	+0.018	[+0.007, +0.028]	-11.2	-17.1
Poly 3 [1.0%]	+0.021	+0.0024	+0.019	[+0.008, +0.029]	-11.6	-17.2
Poly 3 [2.5%]	+0.024	+0.0019	+0.023	[+0.011, +0.034]	-13.0	-17.3
Poly 4 [0.5%]	+0.020	+0.0031	+0.017	[+0.008, +0.026]	-10.9	-17.0
Poly 4 [1.0%]	+0.021	+0.0030	+0.018	[+0.009, +0.026]	-11.3	-17.2
Poly 4 [2.5%]	+0.023	+0.0029	+0.020	[+0.011, +0.029]	-13.1	-18.3
NS 3 [0.5%]	+0.020	+0.0013	+0.019	[+0.009, +0.029]	-11.9	-16.9
NS 3 [1.0%]	+0.021	+0.0011	+0.020	[+0.009, +0.030]	-12.3	-16.8
NS 3 [2.5%]	+0.023	+0.0002	+0.022	[+0.010, +0.034]	-13.3	-16.4
NS 4 [0.5%]	+0.020	+0.0033	+0.017	[+0.007, +0.027]	-13.0	-18.8
NS 4 [1.0%]	+0.021	+0.0030	+0.018	[+0.008, +0.028]	-13.2	-18.9
NS 4 [2.5%]	+0.023	+0.0018	+0.021	[+0.010, +0.032]	-14.1	-19.3

Note. Specification labels abbreviate the convention used in the other specification tables of this appendix: *Poly* k and *NS* k are the polynomial and natural-spline entries with *Deg.* k , and the bracketed percentage is the winsorization, applied symmetrically to both tails.

C The Amenity Dose-Response

This section establishes the claim of the body that bilateral migration data do not identify the amenity dose-response f_a , and sets out the bounds under which the counterfactual is run.

C.1 Sample and first stage

At equilibrium the ratio of households moving from r to s to those staying in r satisfies

$$\log \frac{\mu_t(r, s)}{\mu_t(r, r)} = -\frac{\lambda}{\Omega} [\log L_t(s) - \log L_t(r)] + \frac{1}{\Omega} [\log y_t(s) - \log y_t(r)] + \frac{1}{\Omega} [\mathcal{E}_t(s)f_a - \mathcal{E}_t(r)f_a] + \psi_{rs} + \phi_t + \varepsilon_{rst}, \quad (7)$$

where the pair effect ψ_{rs} absorbs fundamental amenities and migration costs, and $\lambda = 0.32$, $\Omega = 0.5$ follow [Desmet et al. \(2018\)](#). Migration costs are unobserved and proxied by the same gravity vector as trade. In the first stage on the 146,004 five-year bilateral flows of [Abel and Cohen \(2019\)](#), 1990–2015, with World Bank population and GDP per capita and ERA5 temperatures, the controls are jointly informative (Wald $F = 53.69$, $p = 6.8 \times 10^{-56}$). The regression is estimated by PPML with origin $\times$ destination and year effects, on 13,929 origin–destination pairs.

C.2 Level, shape, and the limits of the data

Write the climate term of eq. (7) for pair (r, s) at date t as $\Delta f_{rs,t} \equiv \mathcal{E}_t(s)f_a - \mathcal{E}_t(r)f_a = \overline{\Delta f}_{rs} + (\Delta f_{rs,t} - \overline{\Delta f}_{rs})$. The pair effect ψ_{rs} absorbs $\overline{\Delta f}_{rs}$ exactly, so no within-pair estimator recovers the level of f_a . A normalisation is required; I adopt $f_a(15^\circ\text{C}) = 0$, which fixes the level without restricting the shape. What remains is the within-pair change in bilateral temperature differentials over five-year windows, whose standard deviation in the estimation sample is $\sigma_{\Delta T} \approx 0.15^\circ\text{C}$. In the linearised problem the coefficient of eq. (7) is f'_a/Ω , with standard error $\sigma_\varepsilon/(\sigma_{\Delta T}\sqrt{N_{\text{eff}}})$, where σ_ε is the residual dispersion and N_{eff} the number of pair clusters. The smallest slope detectable at the 5% level with 80% power is

$$|f'_a|_{\min} = (z_{0.975} + z_{0.80}) \Omega \frac{\sigma_\varepsilon}{\sigma_{\Delta T}\sqrt{N_{\text{eff}}}} = 2.80 \Omega \frac{\sigma_\varepsilon}{0.15\sqrt{N_{\text{eff}}}}. \quad (8)$$

N_{eff} is the number of origin–destination pairs; the inference clusters by origin and by destination, which leaves fewer independent units, so the minimum detectable effect reported here is a lower bound on the one implied by the error structure actually used. At $\sigma_\varepsilon = 1.62$ (log-scale PPML

residuals over positive flows, highest-weight specification) and $N_{\text{eff}} = 13,929$, this gives $|f'_a|_{\min} = 12.8$ log points per °C. The scale restriction below places any admissible f_a at or below $\hat{f}_z$, whose average slope between the optimum plateau and 30°C is 0.8 log points per °C. The test could detect an amenity slope only if it exceeded the admissible productivity response by a factor of 15.6. Read directly off the estimates, the 95% interval on the amenity slope spans ± 9.0 log points per °C: it contains zero and the entire admissible response. The data do not reject a large amenity response; they cannot distinguish it from none.

Placebo. The identifying assumption is that within-pair changes in temperature differentials are uncorrelated with contemporaneous shocks to flows. Table 9 substitutes one-period-ahead differentials for contemporaneous ones; the smallest p -value on the leads is 0.3456, so the condition is not rejected. This is necessary for the null to be informative, not sufficient: a test with no power passes a placebo whether or not the true effect is zero.

Table 9: Placebo pre-trend test — Migration (daily temperature distribution)

Functional form	Degree	Winsorization	F -stat (leads)	p -value	Within- R^2	No pre-trend
Polynomial	2	[0.5%, 0.5%]	0.376	0.6864	0.7364	✓
Polynomial	2	[1%, 1%]	0.382	0.6827	0.7368	✓
Polynomial	2	[2.5%, 2.5%]	0.519	0.5950	0.7379	✓
Polynomial	3	[0.5%, 0.5%]	0.270	0.8468	0.7368	✓
Polynomial	3	[1%, 1%]	0.281	0.8392	0.7370	✓
Polynomial	3	[2.5%, 2.5%]	0.925	0.4276	0.7378	✓
Natural spline	3	[0.5%, 0.5%]	0.291	0.8320	0.7364	✓
Natural spline	3	[1%, 1%]	0.419	0.7392	0.7366	✓
Natural spline	3	[2.5%, 2.5%]	1.105	0.3456	0.7378	✓
Natural spline	4	[0.5%, 0.5%]	0.992	0.4101	0.7397	✓
Natural spline	4	[1%, 1%]	0.948	0.4348	0.7392	✓
Natural spline	4	[2.5%, 2.5%]	1.039	0.3852	0.7386	✓

Note. One-period-ahead bilateral temperature differentials substituted for contemporaneous differentials in the Stage 3 PPML regression. Fixed effects: origin $\times$ destination and year; standard errors clustered by origin and destination. H_0 : lead coefficients jointly insignificant. The smallest p -value on the leads is 0.3456. ✓ = $p > 0.10$; ~ = $0.05 < p \leq 0.10$; × = $p \leq 0.05$.

C.3 Robustness across specifications

Tables 10 and 11 report the grid under the daily distribution and the annual mean. The smallest joint p -value is 0.074 under the daily distribution and 0.138 under the annual mean; the joint test rejects in none of the 24 cells. The smallest p -value on the degree- ≥ 2 terms alone is 0.027, in five daily cells; it is a minimum over 24 cells of a test with one fewer restriction, and a value near the threshold is what one expects under the null.

Table 10: Robustness: Amenity dose-response function (daily temperature distribution)

Functional form	Deg.	Winsorization	F (joint)	p (joint)	p (non-lin.)
Polynomial	2	[0.5%, 0.5%]	1.824	0.161	0.073
Polynomial	2	[1%, 1%]	2.010	0.134	0.057
Polynomial	2	[2.5%, 2.5%]	2.605	0.074	0.027
Polynomial	3	[0.5%, 0.5%]	2.051	0.104	0.046
Polynomial	3	[1%, 1%]	2.230	0.082	0.035
Polynomial	3	[2.5%, 2.5%]	2.096	0.098	0.044
Natural spline	3	[0.5%, 0.5%]	2.216	0.084	0.245
Natural spline	3	[1%, 1%]	2.252	0.080	0.228
Natural spline	3	[2.5%, 2.5%]	1.996	0.112	0.239
Natural spline	4	[0.5%, 0.5%]	1.724	0.141	0.245
Natural spline	4	[1%, 1%]	1.898	0.108	0.138
Natural spline	4	[2.5%, 2.5%]	1.755	0.135	0.098


Note. Stage 3 PPML estimates on bilateral migration flows, with origin $\times$ destination and year fixed effects. Standard errors clustered by origin and destination. F (joint) and p (joint): joint significance of all climate terms; p (non-lin.): joint significance of the degree- ≥ 2 terms. Across the 12 specifications the smallest joint p -value is 0.074; the smallest p -value for the degree- ≥ 2 terms alone is 0.027. No within-pair R^2 is reported: it is not computed on this specification grid, and the within-pair R^2 quoted elsewhere in this appendix belongs to the placebo specification. No specification is singled out.

Table 11: Robustness: Amenity dose-response function (annual mean temperature)

Functional form	Deg.	Winsorization	F (joint)	p (joint)	p (non-lin.)
Polynomial	2	[0.5%, 0.5%]	1.137	0.321	0.144
Polynomial	2	[1%, 1%]	1.233	0.291	0.123
Polynomial	2	[2.5%, 2.5%]	1.567	0.209	0.077
Polynomial	3	[0.5%, 0.5%]	1.391	0.243	0.131
Polynomial	3	[1%, 1%]	1.547	0.200	0.102
Polynomial	3	[2.5%, 2.5%]	1.718	0.161	0.077
Natural spline	3	[0.5%, 0.5%]	0.956	0.413	0.323
Natural spline	3	[1%, 1%]	1.097	0.349	0.293
Natural spline	3	[2.5%, 2.5%]	1.437	0.230	0.221
Natural spline	4	[0.5%, 0.5%]	1.725	0.141	0.644
Natural spline	4	[1%, 1%]	1.739	0.138	0.648
Natural spline	4	[2.5%, 2.5%]	1.706	0.146	0.613

Note. Stage 3 PPML estimates on bilateral migration flows, with origin $\times$ destination and year fixed effects. Standard errors clustered by origin and destination. F (joint) and p (joint): joint significance of all climate terms; p (non-lin.): joint significance of the degree- ≥ 2 terms. Across the 12 specifications the smallest joint p -value is 0.138; the smallest p -value for the degree- ≥ 2 terms alone is 0.077. No within-pair R^2 is reported: it is not computed on this specification grid, and the within-pair R^2 quoted elsewhere in this appendix belongs to the placebo specification. No specification is singled out.

Figure 8 draws the two arms on common axes: the medians are close to flat and the Q05–Q95 envelopes span zero to large negative values under both metrics, which is what the identification argument predicts, since changing the metric does not create variation where the pair effects have absorbed it.



(a) Daily temperature distribution

(b) Annual mean temperature

Figure 8: Amenity dose-response under the two exposure metrics. Monte Carlo (300 draws; specification sampled by Akaike weight, which is not degenerate here, then coefficients from $\mathcal{N}(\hat{\beta}, \hat{V})$, each draw normalised to zero at 0°C; median, Q25–Q75 and Q05–Q95. Bars: the arm’s own exposure. Axes cut at the common support.

D General-Equilibrium Details

This section reports the GE results that Section 3 summarises: the path between 2015 and 2100, the distribution across income groups, the full country and latitude tables, the sensitivity to the amenity channel, the Monte Carlo band on the contrast, and the decompositions behind Sections 3.3 and 3.4. Unless stated otherwise, every run uses the amenity lower bound $f_a = 0$.

D.1 Path and dispersion

Table 12 reports the three margins at 2030, 2050 and 2100. The annual total reaches -0.001 by 2030 and -0.008 by 2100; the daily extra is negative at each of the three horizons and grows in absolute value along the path, so that the daily distribution adds losses rather than redistributing them. The share of the world population living in cells with cumulative losses by 2100 is 99.9% under the daily total, 88.3% under the annual total and 72.2% under the daily extra. The daily distribution therefore adds losses for a majority of the world population but not for all of it; in the remaining cells the intra-annual distribution shifts in a way that partially offsets annual-mean

damages. Its decisive effect is at the sign margin: the share of the population in cells that gain falls from 11.7% to 0.1%.

Table 12: Aggregate welfare effects, population-weighted cumulative $\Delta \log U$, amenity lower bound $f_a = 0$. Daily extra is the difference of the two totals, computed before rounding, so the last digit of the three columns need not add up.

Horizon	Daily total (daily arm vs. baseline)	Annual total (annual arm vs. baseline)	Daily extra (daily vs. annual)
2030	−0.003	−0.001	−0.002
2050	−0.010	−0.004	−0.005
2100	−0.017	−0.008	−0.010

D.2 Cold economies and latitude

Table 13 reports the Northern aggregate at each threshold of the latitude grid. The contrast grows with latitude rather than sorting economies into categories: the colder the economy, the larger the annual-mean gain and the deeper the daily extra that removes it.

Table 13: Pop-weighted cumulative $\Delta \log U$ by 2100 across latitude bands, amenity lower bound $f_a = 0$.

Threshold	Pop. (M)	Annual total	Daily extra	Daily total
$\geq 40^\circ\text{N}$	1,112	+0.008	−0.015	−0.007
$\geq 45^\circ\text{N}$	662	+0.011	−0.018	−0.006
$\geq 50^\circ\text{N}$	375	+0.013	−0.019	−0.006
$\geq 55^\circ\text{N}$	109	+0.023	−0.027	−0.003
$\geq 60^\circ\text{N}$	18	+0.030	−0.032	−0.002
$\geq 66.5^\circ\text{N}$	2	+0.033	−0.033	0.000

D.3 The amenity channel

Tables 14 and 15 re-run both scenarios at the interior estimate $\hat{f}_a^{q50}$, the within-pair median across specifications of the amenity dose-response that Supplement C shows to be unidentified. The world daily total moves from −0.017 at the lower bound to −0.218, and the annual total from −0.008 to −0.153, because the interior estimate adds a congestion-adjusted amenity penalty across

the heterogeneous tropical population. For the Northern contrast the sign holds and the magnitude does not: the daily extra above 45°N is -0.018 at the lower bound and -0.059 at the interior estimate, negative for every Northern economy under both and 3.3 times deeper at the estimate; the Northern daily total runs from -0.006 to -0.107 on the same comparison. At the interior estimate the annual arm itself projects -0.048 for the aggregate: the annual amenity response carries no non-linearity the data support — the test of its degree-two and higher terms has p between 0.077 and 0.648 across the grid, and the joint test of all its climate terms p between 0.138 and 0.413 — so the position of its optimum, which decides whether cold economies gain, is not identified, and the contrast does not rest on the annual arm projecting a gain. Because $\hat{f}_a$ lies 10.3 times beyond the admissible bound and below the threshold that Supplement C.2 shows the design could detect, this arm is a stress test past the restriction the paper imposes on itself, not a point inside it.

Table 14: Sensitivity of global pop-weighted cumulative $\Delta \log U$ by 2100 to the amenity specification

Specification	Daily total	Annual total	Daily extra
Interior estimate $\hat{f}_a^{q50}$	-0.218	-0.153	-0.065
Lower bound $f_a = 0$	-0.017	-0.008	-0.010

Table 15: Cumulative $\Delta \log U$ by 2100 for cold-climate economies under the two amenity specifications (interior $\hat{f}_a^{q50}$ vs. lower bound $f_a = 0$), in the order of Table 1.

Country	Annual total		Daily extra	
	$\hat{f}_a^{q50}$	$f_a = 0$	$\hat{f}_a^{q50}$	$f_a = 0$
Russia	-0.007	$+0.025$	-0.081	-0.028
Finland	$+0.012$	$+0.027$	-0.071	-0.028
Norway	-0.035	$+0.013$	-0.050	-0.018
Sweden	-0.041	$+0.012$	-0.042	-0.017
Canada	-0.080	$+0.011$	-0.073	-0.021
Germany	-0.065	$+0.004$	-0.041	-0.011
France	-0.108	-0.008	-0.040	-0.004
North $\geq 45^\circ\text{N}$	-0.048	$+0.011$	-0.059	-0.018

D.4 Band on ranked draws

Table 16 solves the model at draws of the daily dose-response. Three hundred coefficient vectors are drawn from the sandwich covariance of the daily arm's estimate. The draws solved are a stratified sample at evenly spaced percentiles of the projected Northern change in partial equilibrium, computed for each draw by the function behind Figure 2; 60 of them are solved in general equilibrium. The partial-equilibrium ordering carries over imperfectly (the Spearman correlation between the partial-equilibrium ordering of the draws and the ordering of their Northern daily totals in general equilibrium is 0.68) so the percentiles and shares reported are those of the general-equilibrium outcomes. The body reports the share of population in gaining cells for the central run, 0.1%. Across the draws the share has a median of 0.3% and reaches 20.0% in the upper tail, above the annual-mean central value.

Table 16: Distribution of the general-equilibrium daily total by 2100 across Monte Carlo draws of the daily dose-response, cumulative population-weighted $\Delta \log U$.

Unit	Daily total						Population in gaining cells (%)	
	median	q05	q95	min	max	neg. draws	median	q95
World	-0.018	-0.029	-0.005	-0.034	-0.0004	60/60	0.3	20.0
North $\geq 45^\circ\text{N}$	-0.006	-0.011	-0.003	-0.013	-0.0002	60/60	0.6	27.6
North $\geq 55^\circ\text{N}$	-0.004	-0.008	+0.0002	-0.009	+0.003	56/60	2.7	60.8
North $\geq 66.5^\circ\text{N}$	0.000	-0.005	+0.004	-0.006	+0.007	29/60	38.8	99.2
Russia	-0.003	-0.008	+0.0004	-0.009	+0.003	55/60	0.8	65.1
Finland	-0.001	-0.006	+0.004	-0.007	+0.006	38/60	13.4	100.0
Canada	-0.010	-0.016	-0.004	-0.021	-0.001	60/60	0.2	1.7
Germany	-0.007	-0.011	-0.003	-0.012	-0.001	60/60	0.0	0.0
France	-0.012	-0.017	-0.008	-0.021	-0.007	60/60	0.0	0.0
Norway	-0.005	-0.009	-0.001	-0.011	+0.002	59/60	5.1	20.0
Sweden	-0.005	-0.009	-0.0002	-0.010	+0.002	57/60	0.6	31.5

Note. Distribution of the daily total across 60 Monte Carlo draws of the daily dose-response (specification a natural spline with 4 degrees of freedom, winsorised at 0.5%), each solved in general equilibrium with amenities at the lower bound and the no-climate baseline common to all runs. Draws are a stratified sample of the Monte Carlo distribution, taken at evenly spaced percentiles of the projected Northern change in partial equilibrium. Percentiles and shares are those of the general-equilibrium outputs.

D.5 Warming switched off by region

Section 3.4 compares the headline runs with runs in which warming is switched off in part of the world. Under the first mask, the 4,581 cells at or above 45°N keep their projected exposure and the other 9,405 cells keep their 2015 exposure in every period; the second mask does the reverse. The model computes damage from the change in exposure between periods, so a cell held at its 2015 exposure receives no damage, as in the scenario without warming. Nothing else is switched off: goods trade, migration, wages and congestion operate in every cell, and each run is a full general-equilibrium solution under a partial warming. Each mask is solved under the daily and the annual-mean function, at the central draw of the dose-response and at the amenity lower bound $f_a = 0$; the run without warming is the headline one. These runs carry no Monte Carlo band. The model is not linear, so the full-warming value need not equal the sum of the two masked values; Section 3.4 reports the difference, the interaction, with the shares.

D.6 The two decompositions

The two tables below carry quantities that the body reports and that no other table does. The first is a difference between runs already described above. The second reports the income component by latitude band and, above 45°N, its share with warming confined to the North, from the masked runs of Supplement D.5.

Table 17: Evaluation and function terms of the Northern daily extra (Section 3.3). The third scenario applies the daily response to annual-mean exposure, so the evaluation term is the part of the extra that the change of exposure alone produces and the function term the remainder.

	Cumulative $\Delta \log U$
Northern daily extra	−0.0178
evaluation term	+0.0034
function term	−0.0211
Ratio, function to evaluation	6.3

Note. The two terms are of opposite sign to one another. In partial equilibrium the same split attributes 70.1% of the Northern contrast to the shape of the response, of which 76% lies below 0°C and 8% above 20°C.

Table 18: The income component of the daily extra by latitude band (Section 3.4), at the central draw of the dose-response. Each entry is the income component for the cells of the band; the last column is its value with warming confined to cells at or above 45°N, as a share of its value when the whole world warms.

Latitude band	Income component	Northern warming only (%)
Below 30°N	−0.012	—
30–45°N	−0.002	—
45–60°N	−0.019	95.8
Above 60°N	−0.037	98.5

Note. The share is not reported below 45°N, where the cells do not warm under that mask.

E The Cruz–Rossi–Hansberg Pipeline

The bilateral estimator identifies the productivity dose-response from a cross-section of trade flows between countries, and its vulnerability is the confounding specific to that cross-section.

This section re-runs the contrast in a design that shares neither its data, nor its unit, nor its fundamental: the estimation of [Cruz and Rossi-Hansberg \(2024\)](#) on model-inverted cell fundamentals, under their fixed effects — sub-region $\times$ period for productivity, blocks of 2×2 cells and sub-country $\times$ period for amenities — with a single lever moved, the exposure metric. The design has no cell fixed effect: identification is across cells within a sub-region and period, with 6% of the identifying variance of the productivity regressors and 25% of the amenity regressors lying within cells across periods (0.6% and 2.4% for the mean temperature).

I invert the [Cruz and Rossi-Hansberg \(2024\)](#) model to recover productivities $\bar{a}_t(r)$ and amenities $\bar{b}_t(r)$ on their 1° grid (17,048 cells, four decadal cross-sections 1990–2005, 68,192 cell-periods), following [Desmet et al. \(2018\)](#); I use my own inversion rather than the deposited fundamentals so that inversion, estimation and projection are internal to one reproduction. Re-estimating the authors’ own specification (a ten-year mean of the coldest-month temperature) on my inversion reproduces their bliss points (1.9°C against 1.6°C for productivity, 8.6°C against 8.6°C for amenities). The daily arm follows [Schlenker and Roberts \(2009\)](#): the fundamental is regressed on the shares of days in each temperature bin, $\sum_k \hat{\delta}_k s_k(r)$, under the authors’ fixed effects and geographic controls, and the $\hat{\delta}_k$ are re-estimated, not plugged in from the annual function. The two metrics agree in their first moment: the cell mean of the ERA5 daily distribution reproduces the authors’ own decadal average temperature with correlation 0.996, slope 0.97 and bias -0.03° . The authors cluster spatially (Conley, 550 km). Every interval reported here clusters at the country level. The projection is linear in the estimated coefficients, so the interval on the difference between arms carries their covariance (correlation 0.67 to 0.89) rather than treating them as independent.

Figure 9 plots the re-estimated damage functions.

In the day-response $\hat{\delta}(T)$, an additional day is most valuable near 12.2°C for productivity and 13.5°C for amenities, and its value declines above 20–25°C. As a cold cell warms, its distribution grows an upper tail into that declining region, which an evaluation at the mean cannot register. Across baseline temperatures the daily projection lies below the annual-mean one, and in the temperate band the two disagree in sign at the cell level. The upper tail is data-poor in the estimation




Figure 9: Productivity (top) and amenities (bottom), under the daily-distribution metric (left) and the annual-mean metric (right). Left: the effect on the change in the log fundamental of shifting one unit of day-share into temperature T , relative to the reference bin at 25.9°C; points are bin coefficients, the line a weighted natural-spline fit. Right: the effect on the same object of a one-degree rise in the annual mean. Columns differ and their vertical scales are not comparable.

sample, the mirror image of the cold-branch caveat of Supplement B; I hold $\hat{\delta}(T)$ flat beyond 40.0°C, the hottest bin with at least 1% of cell-periods, and $1.1 \times 10^{-5}\%$ of the cold economies' projected day-mass lies beyond it: clamping moves the Northern aggregate by 3.8×10^{-8} . Holding $\hat{\delta}(T)$ flat beyond 30°C instead moves the Northern difference by $+8.2 \times 10^{-4}$, and 0.28% of the day-mass of cells above 45°N lies above 30°C by 2100: the upturn of the fitted response between 30 and 40°C, which rests on a few bins with little mass, plays no part in the Northern result.

The shape of the annual-mean response in Figure 9 differs from the authors' published one,

and the difference is a property of the fitting, not of the data. The authors fit to their twenty bin slopes a logistic function whose slope parameter is fixed, so that the estimated response is a smooth step by construction: flat at the cold end, flat at the hot end, and a single transition whose centre is the bliss point. The annual arm here smooths the same kind of bin slopes with a natural spline, which imposes no shape, and what it shows is that the marginal effect of warming is largest at annual means between 0 and 10°C and weakens towards the coldest cells, below −10°C, where few cells lie and the slope is poorly estimated under either specification. The two responses also sit on different supports: in the authors’ winter-month metric the Siberian cells lie between −40 and −20°C, on the flat upper plateau of the logistic, whereas in the annual mean the same cells lie between −15 and 0°C, where the spline turns down. What the two fits agree on is what the projection uses: the marginal effect is positive throughout the North and negative in the tropics, and its zero crossing (at 20.0°C in annual-mean terms, 1.9°C in the authors’ winter-month terms) separates the same cells. I keep the spline rather than adopt the authors’ logistic for the annual arm. Imposing their functional form would reproduce the appearance of their figure, but it would do so by ruling out what the bin slopes show: that the marginal effect of warming weakens in the coldest cells. The spline reports that feature and the logistic cannot; where the two agree, on the sign and the crossing, the choice is immaterial, and where they differ, the spline follows the data.

Table 19 reports the projected change in the log fundamental between 2015 and 2100 under SSP2-4.5 for the five largest cold economies and the aggregate above 45°N. The projection runs to 2100 rather than the authors’ 2200, uses SSP2-4.5 rather than RCP 8.5, and is a partial-equilibrium projection of the fundamental without their endogenous innovation.

Table 19: Projected change in the log fundamental, 2015–2100(SSP2-4.5), in the Cruz–Rossi-Hansberg pipeline: annual-mean versus daily-distribution metric and their difference. Partial-equilibrium projection, four-period sample, 95% intervals clustered by country.

Region	Annual mean	Daily	Difference
<i>Productivity</i>			
Russia	+0.778 [+0.310, +1.258]	+0.247 [−0.140, +0.635]	+0.531 [+0.311, +0.754]
Canada	+0.790 [+0.287, +1.307]	+0.171 [−0.248, +0.591]	+0.618 [+0.359, +0.883]
Finland	+0.670 [+0.274, +1.075]	+0.379 [+0.027, +0.728]	+0.291 [+0.071, +0.515]
Germany	+0.441 [+0.104, +0.789]	+0.095 [−0.212, +0.394]	+0.347 [+0.103, +0.592]
France	+0.297 [+0.015, +0.584]	−0.187 [−0.453, +0.071]	+0.484 [+0.264, +0.712]
North $\geq 45^\circ\text{N}$	+0.559 [+0.186, +0.943]	+0.116 [−0.184, +0.414]	+0.443 [+0.238, +0.651]
<i>Amenity</i>			
Russia	+0.063 [+0.000, +0.126]	+0.033 [−0.006, +0.072]	+0.031 [−0.012, +0.074]
Canada	+0.066 [−0.001, +0.132]	+0.028 [−0.011, +0.066]	+0.038 [−0.007, +0.084]
Finland	+0.054 [+0.001, +0.107]	+0.036 [+0.006, +0.066]	+0.018 [−0.016, +0.052]
Germany	+0.039 [−0.003, +0.081]	+0.015 [−0.008, +0.038]	+0.024 [−0.004, +0.053]
France	+0.028 [−0.006, +0.062]	−0.004 [−0.028, +0.020]	+0.032 [+0.005, +0.060]
North $\geq 45^\circ\text{N}$	+0.047 [−0.001, +0.096]	+0.019 [−0.006, +0.044]	+0.028 [−0.003, +0.060]