Permanent Climate Change and GDP

Jonathan B. Berk
Stanford University and NBER

Jules H. van Binsbergen
University of Pennsylvania and NBER

Cem Kozanoglu
Stanford University

August 6, 2026

Abstract

We examine internal migration within the United States to assess whether climate change will lead to lower economic growth. Because a mover experiences a persistent change in climate exposure, internal migration offers a useful setting for evaluating whether permanent warming generates large income losses, a question that the existing literature has addressed using temporary weather variation. We find that moves to warmer states and metropolitan areas are associated with marginally lower earnings. The estimated gradients are at least an order of magnitude smaller than even conservative existing estimates of the effect of climate change on economic activity, and a revealed preference for warmer destinations suggests they are an upper bound on a pure productivity penalty. Our findings suggest that research relying on temporary weather shocks to estimate the consequences of permanent warming may overstate the economic costs.

One of the most important questions facing policymakers assessing the importance of climate change is its effect on economic activity. While the existing literature largely agrees that continued warming will reduce economic activity, no consensus has emerged regarding the magnitude of the effect. Estimates range from losses of approximately two percent of gross domestic product (GDP) per degree Celsius at the low end to recent work identifying losses of twenty percent per degree Celsius increase of global temperatures (Bilal and Kanzig, 2026). The distinction is of considerable consequence. At the low end, a 2°C increase in global temperature would be equivalent to foregoing only a few years of economic growth; at the high end, the same increase would reduce GDP by nearly half. The policy implications therefore depend on a parameter that current estimates identify only to within an order of magnitude. The central message of this paper is that even the lower end of this range is difficult to reconcile with the earnings of workers who relocate across climates: the losses associated with such moves are at least an order of magnitude smaller than the low-end estimates imply.

In 2023, Florida had a population-weighted average temperature of 23.9°C, compared with 9.2°C in Wisconsin, a difference of 14.7°C (26.5°F). Applying the low-end gradient of two percent of GDP per degree Celsius, Florida's output per person should fall short of Wisconsin's by roughly a quarter, corresponding to about $15,000 per person. The high-end gradient of twenty percent per degree implies a difference so large that Florida's per capita GDP should be less than one percent of Wisconsin's. Yet the two states have nearly identical output per person, within 1.4 percent of each other ($57,900 in Florida versus $58,700 in Wisconsin). As we document below, this pattern is not peculiar to Florida and Wisconsin: across the United States, the large economic costs that current estimates attribute to a warmer climate are difficult to discern. Many other factors plainly differ across states, but given the magnitude of the effect these estimates in the literature imply and the substantial temperature variation within the country, climate should leave a visible imprint on productivity. That it does not, raises the question of whether the prevailing range of estimates is plausible.

Many existing empirical studies that estimate the effect of climate change on GDP share a common feature: they use temporary variation in temperature to infer the consequences of a permanent change in climate. In most cases these studies explicitly *remove* the long-run trend in the data, identifying the effect from year-to-year fluctuations alone. To see why this is problematic, consider the following thought experiment. Suppose that, at the height of the Roman Empire, the inhabitants of Gaul were required to speak French for a single day. Because they spoke Latin at the time (French would not exist for many centuries) output would collapse that day, as people were temporarily unable to communicate. Crucially,

knowing that they could return to Latin the following day, no one would have any reason to adapt to the imposition; the cost is incurred precisely because the change is known to be transitory. Using the one-day collapse in output to infer what GDP would be centuries later, once French had been learned and fully adopted, would lead one to conclude (erroneously) that speaking French is catastrophic for economic activity. The error lies in mistaking the cost of being unadapted to a temporary shock for the cost of a permanent change to which people have had time to adjust.

The analogy illustrates the difficulty in estimating the effect of climate change on economic activity. Using temporary variation in temperature to infer the consequences of permanent change ignores the possibility that adapting to small increments over a long horizon may be nearly costless, so that a shock which is damaging when it arrives unexpectedly is benign when it is permanent and anticipated. The temporary response therefore need not bear any close relationship to the long-run effect. However, it is this long-run effect that is the object of policy interest. Recovering that long-run effect from realized outcomes is itself difficult, because over the relevant horizons many other things change alongside the climate. In the analogy, the shift from Latin to French coincided with the economic decline of the early medieval period (the dark ages) a correlation that plainly reflects confounding forces rather than any effect of language on output. The research design must therefore isolate the climate component of a permanent change from the many other factors with which it is correlated. This requires not merely a single long-run comparison but a large number of comparable permanent shocks, observed with enough detail to hold constant the most important confounds. This is precisely the structure that internal migration within the United States provides.

Our objective in this article is to bring a new source of variation to bear on this question — one that reflects large permanent differences in climate and is available for a large number of individuals. The approach exploits the cross-sectional variation in climate across the United States described above. When a person permanently relocates from one state to another, they experience a lasting shock to the climate to which they are exposed. Because a large number of Americans have moved between states, we observe many individuals who have undergone such changes, and we can use the resulting earnings changes to assess how permanent differences in climate affect labor productivity.

One key advantage of our methodology is that the permanent climate shocks individuals experience when they move are much larger than either the small permanent shocks caused by the gradual warming of the planet or the temporary weather shocks recently explored in the literature. The climate shocks experienced by the movers in our database span a range of more than 30°C, with some individuals experiencing an increase of 15°C and others a

decrease of the same magnitude. These shocks dwarf the low single-digit changes that other studies have used.

Not surprisingly, because people choose to move, we find that both moves to warmer and colder climates are associated with higher earnings. But the increase is smaller for moves to warmer locations. To the extent that this difference reflects lower individual productivity, it implies a corresponding reduction in output, consistent with the directional results in the existing literature. However, our estimate is an order of magnitude smaller than low end existing estimates of the effect of permanent local temperature differences on output, which cluster around two to three percent per degree Celsius. We find that a one-degree-Celsius increase in destination temperature is associated with a reduction in the earnings gain of roughly 20 basis points. Treating this as a proxy for the effect on output, a 2°C warming would imply a reduction in GDP of less than half a percent. This one-time reduction is in the same ballpark as the uncertainty economists generally associate with estimating GDP to begin with.

Our estimate likely overstates the true productivity effect. We provide evidence that part of the earnings difference reflects an amenity, rather than a difference in productivity. Workers reveal a preference for warmer climates. Holding population size and other characteristics fixed, warmer locations attract more migrants. The apparent desirability of warmer places is, by itself, difficult to reconcile with the large economic damage existing estimates attribute to warmth. Leaving that aside, a revealed preference of this kind should give rise to a compensating differential, so that wages in warmer locations should be lower for reasons unrelated to productivity. Consequently, our estimated gradient likely conflates a productivity effect with an amenity effect, and should therefore be read not as a point estimate but as an upper bound on the magnitude of the climate effect on individual productivity.

Because we compare workers who move to warmer destinations with those who move to colder ones, selection that is common to movers in both directions is differenced out, leaving the component of the earnings change that varies with the temperature change itself. The move is still a choice rather than an experiment, so we do not claim to estimate the full welfare cost of climate change. We provide instead a diagnostic for one channel that is central to the climate-damages literature. If permanent exposure to warmer local temperatures substantially depresses labor productivity (corresponding to >20% drops in GDP), then workers who move across large climate gradients should experience correspondingly large changes in their earnings. They do not.

The remainder of the paper proceeds as follows. In the next section we explain how our article fits into the existing literature. Section 2 sets out the motivating puzzle: within

the United States the cross-sectional relationship between local temperature and economic output is far weaker than what current climate change estimates would imply. We then proceed to test the effect of a permanent change in climate by observing migrants. Section 3 describes the main dataset which lets us observe how a mover's wage changes with origin–destination temperature differences and allows us to control for factors such as age, education, and job sector. Section 4 derives the main estimates in the paper. In Section 6 we discuss the implications and limitations of our findings. Section 7 concludes the article.

1 Background

This paper is motivated by a tension between two observations. The climate-economics literature predicts substantial reductions in GDP from global warming, yet within the United States, which is a climatically diverse but institutionally homogeneous country, the cross-sectional relationship between local temperature and economic output is weak. This section places that tension within the existing literature and explains why internal migration provides a natural setting in which to resolve it.

Existing estimates of the economic damage from climate change span a wide range. Meta-analyses and integrated assessment models tend toward the low end, implying losses of roughly two percent of GDP per degree Celsius of warming or less (Tol, 2022; Barrage and Nordhaus, 2024). Reduced-form empirical estimates tend to be larger, with the most recent work estimating a twenty percent decline in GDP per degree Celsius of global mean temperature (Bilal and Kanzig, 2026).[1]

Because labor income constitutes a substantial share of national income, large output damages from temperature should be reflected in labor earnings, particularly if reduced labor productivity is among the operative channels. A large literature documents how heat affects workers' hours, attendance, output, and income, drawing on U.S. time-use and county-income data as well as manufacturing, agriculture, and attendance records in hotter settings (Deryugina and Hsiang, 2014; Graff Zivin and Neidell, 2014; Somanathan et al., 2021; Ireland et al., 2024; Ortiz-Bobea et al., 2021; Casey et al., 2024). The direct heat-productivity

[1]Integrated assessment models imply central losses around 2% of GDP per degree Celsius (Barrage and Nordhaus, 2024); reduced-form weather and subnational panels imply one to five percent (Kalkuhl and Wenz, 2020; Waidelich et al., 2024); growth-effect and nonlinear extrapolations are higher (Burke et al., 2015; Moore and Diaz, 2015); and recent work identifying damages from global temperature reaches twenty to thirty percent per degree (Bilal and Kanzig, 2026; Neal et al., 2025). Burke et al. (2015) estimate a nonlinear relationship between temperature and economic production and project large global income losses under unmitigated warming. Hsiang et al. (2017) estimate sectoral U.S. damages that rise nonlinearly with global mean temperature and include labor supply as one channel. More recent work has continued to debate whether temperature affects output levels, persistent growth, or permanent growth rates (Kahn et al., 2021; Kalkuhl and Wenz, 2020; Nath et al., 2025).

channel is inherently local and operates on the level of output: the relevant exposure is the temperature where workers live and work. The broader climate-economy literature begins from the same premise, treating climate as a distribution of local weather realizations and asking when weather variation can be used to learn about the consequences of climate (Dell et al., 2014; Hsiang, 2016; Carleton and Hsiang, 2016).

These considerations make cross-sectional climate differences a useful diagnostic for the magnitude of long-run local productivity effects. Weather shocks and permanent climate exposure are distinct economic objects: workers and firms have no incentive to adjust to temporary weather shocks, whereas the opposite is true for permanent climate differences. This distinction is emphasized by work separating short-run weather effects from long-run climate relationships ((Tol, 2021; Dong et al., 2025; Gospodinov et al., 2025; Nath et al., 2025)). A strand of recent work finds that long-run and low-frequency temperature relationships are more modest than the effects implied by annual weather shocks (Kalkuhl and Wenz, 2020; Dong et al., 2025; Gospodinov et al., 2025).

Our study connects two strands of evidence. The first shows that heat can have economically meaningful effects in exposed sectors and on extreme days: it reduces and reallocates work hours (Graff Zivin and Neidell, 2014), lowers manufacturing productivity and raises absenteeism (Somanathan et al., 2021), reduces attendance during extreme heat (Ireland et al., 2024), and damages agricultural productivity (Aragon et al., 2021; Ortiz-Bobea et al., 2021). The second uses aggregate temperature-output regressions to infer macroeconomic damages from warming. Our paper asks what the first strand implies for the second, using workers who experience permanent changes in local climate exposure.

Our main contribution concerns internal migrants who move between locations with different climates. Because the decision to move is endogenous, it is unsurprising that most movers earn a sizable premium; to net this out, we study how that premium varies with the origin-to-destination difference in climate. Focusing on movers addresses a central limitation of the existing literature, its reliance on temporary weather variation rather than permanent climate differences, since the translation from weather to climate damages requires strong assumptions about persistence, adaptation, and equilibrium adjustment (Dell et al., 2014; Tol, 2021). However, wages across locations are equilibrium objects that reflect productivity, amenities, prices, and moving costs jointly (Albouy et al., 2016; Sinha et al., 2021; Rappaport, 2007; Kennan and Walker, 2011). To the extent that warmer destinations offer amenities valued by the marginal mover, lower wages in warmer places partly reflect compensation for those amenities, making the wage-temperature gradient we estimate an upper bound on the pure productivity penalty.

Our findings admit two interpretations. Either weather-shock estimates substantially

overstate the damages from permanent warming, because they do not account for adaptation and the other differences between temporary and permanent shocks;[2] or large GDP damages operate only weakly through the local labor-productivity effects of permanent temperature exposure, and are instead generated by other channels.[3] The implication is that if the large damage estimates are to be believed, they require a sharper account of the mechanism. The channels that generate these damages must be identified and measured directly, since they do not appear in the local productivity of workers who move across climates.

2 The Magnitude Puzzle in the United States

We begin by generalizing the Florida-Wisconsin comparison from the introduction and show that its lesson holds across the United States. We focus on cross-sectional differences across U.S. states because, relative to cross-country comparisons, states share a common federal institutional environment, economy, culture, currency, and legal framework. This shared environment removes many of the confounding factors that complicate international comparisons and makes the cross-section a useful descriptive benchmark.

The data for this exercise come from standard public sources, each measured as a state average over 2015–2019 to smooth out business-cycle fluctuations. State output is real gross domestic product from the Bureau of Economic Analysis regional accounts (series SAGDP1), expressed per capita using state population from the National Cancer Institute's SEER county population series. Temperature is the population-weighted average of NOAA nClimDiv county temperatures (the same climate source used throughout the paper and described in Section 3) converted to degrees Celsius. The controls in X_s are drawn from the standard federal statistical sources: population density combines the SEER population with Census Bureau land area; the share of adults aged 25 and over with a bachelor's degree is from the Census Bureau's American Community Survey five-year estimates; the urban population share is from the 2010 Census; and the working-age (20–64) population share is constructed from the SEER age distribution. The sample is the 48 contiguous states; Alaska, Hawaii, and the District of Columbia are excluded.

We estimate the relation between state climate and GDP by running the following state-

[2]This argument is closely related to the point emphasized in the literature that short-run weather effects can be larger than longer-run climate gradients once adaptation and persistence are taken seriously (Kalkuhl and Wenz, 2020; Dong et al., 2025). Direct tests for adaptation within the panel literature have not found strong evidence for it (Burke et al., 2015; Kahn et al., 2021), but these tests identify adaptation from how the response to temporary variation evolves, not from permanent differences in exposure.

[3]Many damage estimates within the climate-economics literature list labor supply and productivity as an important damage channel (Dell et al., 2014; Carleton and Hsiang, 2016; Hsiang et al., 2017; Tol, 2022).

level cross-sectional regression:

$$\log y_s = \alpha + \beta^{CS} T_s + X_s' \gamma + \varepsilon_s, \tag{1}$$

where s indexes states, y_s is the real GDP of the state per capita, T_s is the state's average temperature in °C, and X_s contains the controls shown across the columns of Table 1.

The state table uses the 48 contiguous states and reports GDP per capita under increasingly richer specifications (i.e., additional controls). The first column in the table demonstrates that what was true for Florida and Wisconsin, holds more generally. States with higher temperatures have somewhat lower GDPs, though across specifications, the gradient is around -1.3% per °C. Figure 1 graphically demonstrates the same point.

	(1)	(2)	(3)
Mean temperature	-0.0136	-0.0126	-0.0144
	(0.0053)	(0.0062)	(0.0050)
ln(Population)		0.0406	-0.0069
		(0.0330)	(0.0351)
Population density		0.0001	-0.0000
		(0.0001)	(0.0001)
% Bachelor's degree		0.0137	0.0063
		(0.0041)	(0.0044)
% Urban population			0.0063
			(0.0017)
Working-age share			3.3991
			(2.2314)
Constant	11.1023	9.9581	8.5348
	(0.0732)	(0.4644)	(1.2008)
R^2	0.109	0.471	0.607
N	48	48	48

Table 1: U.S. State Climate and GDP: OLS over the 48 contiguous states (2015–2019 averages); the dependent variable is log real GDP per capita and the regressor of interest is population-weighted mean annual state temperature (°C). Controls are added progressively across columns. HC1 (White) standard errors are in parentheses.

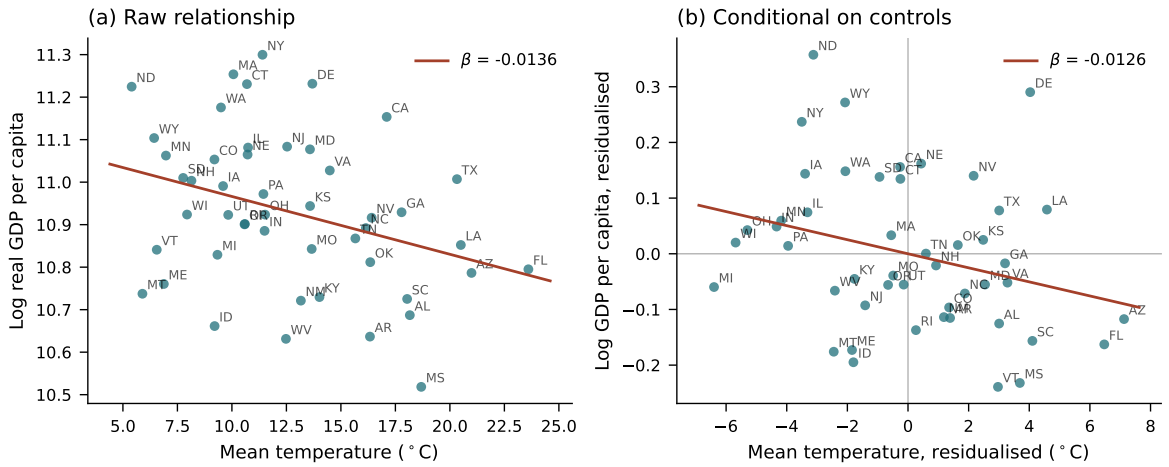


Figure 1: U.S. State Climate and GDP: Each point is one of the 48 contiguous states (2015–2019 averages); warmer states lie to the right. Panel **a** plots the raw relationship between a state's population-weighted mean annual temperature (°C) and its log real GDP per capita (column (1) of Table 1); the dashed line is the OLS fit. Panel **b** is added-variable plot in which both axes are residualized on the partial control set (corresponding to column (2) of Table 1). The controls are log population, population density and bachelor's-degree share, so the slope is the temperature coefficient holding those controls fixed.

Of course, cross-sectional evidence is not always consistent with time-series effects. That is, the fact that people in different states have the same productivity in different climates does not necessarily imply that changing the climate over time in a particular state will have no effect. For example, one possible explanation that could reconcile the cross-sectional evidence above with the time series evidence in the literature is that different jobs are located in different states. That is, jobs that are more productive in warm climates are located in warm states, while jobs that are more productive in cold climates are located in cold states. To accurately estimate a time-series effect, a time-series shock is required.

3 Data and Empirical Setting

When an individual moves from a location in the U.S. to another location in the U.S. she experiences a permanent shock to her personal climate. By measuring the change in her productivity (wages) that results from the move, we can infer how a permanent shock to climate affects productivity. Of course, to accurately assess the effect of climate, it is important to control for other changes that are move-related. For that reason, we will focus our attention on people who move from permanent employment in one location to permanent employment in another location within the same industrial sector.

3.1 Job-to-Job Flows Dataset

The Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program produces the Job-to-Job (J2J) flows dataset, which links employer and household data to observed job changes. An additional product, the Job-to-Job Origin-Destination Table (J2JOD), tracks workers who move as they change jobs, with flows at both state and metropolitan area (metro) levels at a quarterly interval. A metropolitan area is the Core-Based Statistical Area (CBSA) defined by the Census Bureau. The J2JOD has granular level data that allows us to test our hypothesis that a warming climate will lower productivity. Each observation provides the flow of work relocation between a specific origin-destination pair. The flows are broken down by NAICS industry, and the average quarterly earnings of the movers before and after the move are also reported.

More granular versions of J2JOD are also available, where each observation is a demographic subtype of the aggregate origin-destination-industry flow. The different demographics of the workers available are sex, age, race, ethnicity, and education. In addition, movers are separated into categories based on whether they had a short nonemployment gap between the two jobs. Our main analysis restricts the dataset to what the Census Bureau terms stable-to-stable job changes, termed *End-to-End Stable* (EES) in the dataset, where the worker is employed for the full-quarter at both origin and destination so there is no gap in employment. In the appendix, we show that broadening the analysis to include moves with short nonemployment gaps does not materially change the results.

The dataset allows us to observe, for a group of workers in year t, who moved from an origin state o and sector ℓ, to a destination state d, into sector r, their average quarterly wage earnings in the last full-quarter in their origin state-sector (O), $\bar{w}^O_{od\ell rt}$, their average first full-quarter wage earnings in the destination state-sector (D), $\bar{w}^D_{od\ell rt}$, and the number of workers who satisfied all these criteria, $N_{od\ell rt}$.[4] The dependent variable we use is the percent increase in wages the average worker experienced by moving, $y_{od\ell rt}$,

$$y_{od\ell rt} = \log \bar{w}^D_{od\ell rt} - \log \bar{w}^O_{od\ell rt}. \tag{2}$$

Our main specification restricts attention to same-sector moves, $r = \ell$, so in that specification we will drop the redundant subscript and write $y_{od\ell t}$. In our regressions we will weight each observation, $y_{od\ell t}$ by $N_{od\ell t}$, effectively giving each worker equal weight.

[4]Each origin-destination-sector-year cell consists of an average across the quarters in the year, weighted by the number of stable (EES) movers in each quarter, and $N_{od\ell rt}$ is that EES worker count summed across the quarters of the year.

3.2 Climate Measures

Climate data were obtained from the National Oceanic and Atmospheric Administration's (NOAA) nClimDiv dataset at the county level and at a monthly frequency. We use average temperature for all of our main analyses. The main regressor is the destination-origin temperature difference for states and metro areas. The temperature in a state or metropolitan (metro) area is the population weighted average temperature over the counties that make up the area, D_{odt},

$$D_{odt} = T_{dt} - T_{ot}, \tag{3}$$

A few construction details are worth noting. State and metro temperatures are weighted by each county's contemporaneous (same-year) population; because within-state county population shares move slowly, this choice is immaterial in practice. County temperatures are available through 2023; the J2J flows extend to 2024, so the effective estimation window for the temperature regressions is 2000–2023, and the year 2000 (for which only the second through fourth quarters are observed) is retained with year fixed effects absorbing its level. The metro analysis is restricted to metro-to-metro moves, since the LEHD nonmetropolitan "balance of state" areas have no metropolitan temperature; these account for roughly one quarter of metropolitan transitions and are excluded throughout.

We use average temperature as the main climate measure because the object of interest is a persistent difference in local climate exposure, not the realized incidence of any one year's extreme days. Much of the worker-productivity literature emphasizes thresholds and extremes, and this is a potentially important mechanism through which mean climate can matter (Graff Zivin and Neidell, 2014; Somanathan et al., 2021; Ireland et al., 2024). At the same time, the macro climate-damages literature commonly summarizes exposure using average temperature or temperature changes, precisely because these objects provide a compact measure of the local climate gradient (Dell et al., 2014; Carleton and Hsiang, 2016). To connect these approaches, Appendix Table 8 repeats the main state specification using the number of extreme cooling degree days (CDD) and heating degree days (HDD). The resulting coefficients are very similar to the average-temperature specification, suggesting that the small wage-temperature gradients are not driven by the choice to summarize climate with average temperature.

Unless stated otherwise, all of our regressions are weighted by the number of stable job-to-job transitions ($N_{od\ell t} = \text{EES}_{od\ell t}$) and standard errors (and corresponding p-values) are clustered by origin-destination pair (state pair in the analysis that uses states and CBSA pair in the metro analysis). All temperatures are reported in degrees Celsius.

4 Analysis

Our main empirical specification estimates the relationship between the change in wages after a move within job sector and the associated change to average temperature that results from the move (the permanent climate shock):

$$y_{od\ell t} = \alpha + \beta D_{odt} + \tau_t + \kappa_\ell + \varepsilon_{od\ell t}, \tag{4}$$

where τ_t are year fixed effects and κ_ℓ are common-sector fixed effects. In Appendix A.1, we broaden the sample to include people who move to jobs in different sectors and show that the main coefficients do not materially change.[5]

	(1)	(2)	(3)
Temperature difference (°C)	-0.0018***	-0.0026***	-0.0045***
	(0.0005)	(0.0005)	(0.0007)
Log population difference		0.0135***	
		(0.0022)	
Population density difference		0.0000***	
		(0.0000)	
Bachelor's share difference		0.0017***	
		(0.0004)	
Mean moving wage premium	0.096	0.096	0.096
Year FE	Yes	Yes	Yes
Common-sector FE	Yes	Yes	Yes
Same-sector sample	Yes	Yes	Yes
Economic/demographic controls	No	Yes	No
Origin/destination state FE	No	No	Yes
Observations	688,219	688,219	688,219
State-pair clusters	2,550	2,550	2,550
R^2	0.044	0.057	0.103

Table 2: Earnings Changes After Moves. Worker-migrant-weighted regressions of the same-industry earnings change after a move on the destination-minus-origin state temperature difference (°C). Column (1) is the baseline (year and common-sector fixed effects); column (2) adds origin–destination differences in economic and demographic controls; column (3) adds origin and destination state fixed effects, isolating temporary weather variation. Standard errors are clustered by state pair. All three columns are estimated on a common sample. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

We begin our analysis using state level data. Our results are displayed in the first

[5]Approximately half of all moves are within the sector.

column of Table 2. The mean moving wage premium in the state sample is about 0.096 log points, so the average same-sector migrant experiences a wage increase of roughly 9.6 percent by moving.[6] Against that benchmark, the temperature gradient is small. The β coefficient estimate is -0.0018, implying that a move to a destination that is 1°C warmer is associated with a moving premium that is lower by about 0.18%. Putting this another way, the implied effect on productivity of a 2°C increase in temperature is less than 40 b.p.'s. Thus, move-induced climate shocks are largely inconsequential to the overall productivity shock movers experience.

The results are depicted graphically in Figure 2. The temperature-change binned scatter plot shows a small but visible negative gradient in the dense center of the data, while the tails widen for movers crossing unusually large climate gaps in excess of 15°C. In Appendix A.4, we show that the β coefficient is not materially driven by these extreme points. We also investigate the migrants in the extreme bins and find that they are unrepresentative of the broader mover population, with an important role for mining-sector moves, likely oil-related, from warmer states such as Texas to high-wage (colder) destinations such as Alaska and North Dakota.

[6]The mean moving wage premium is the worker-migrant-weighted mean of $y_{od\ell t}$ in the regression sample, not the regression intercept.

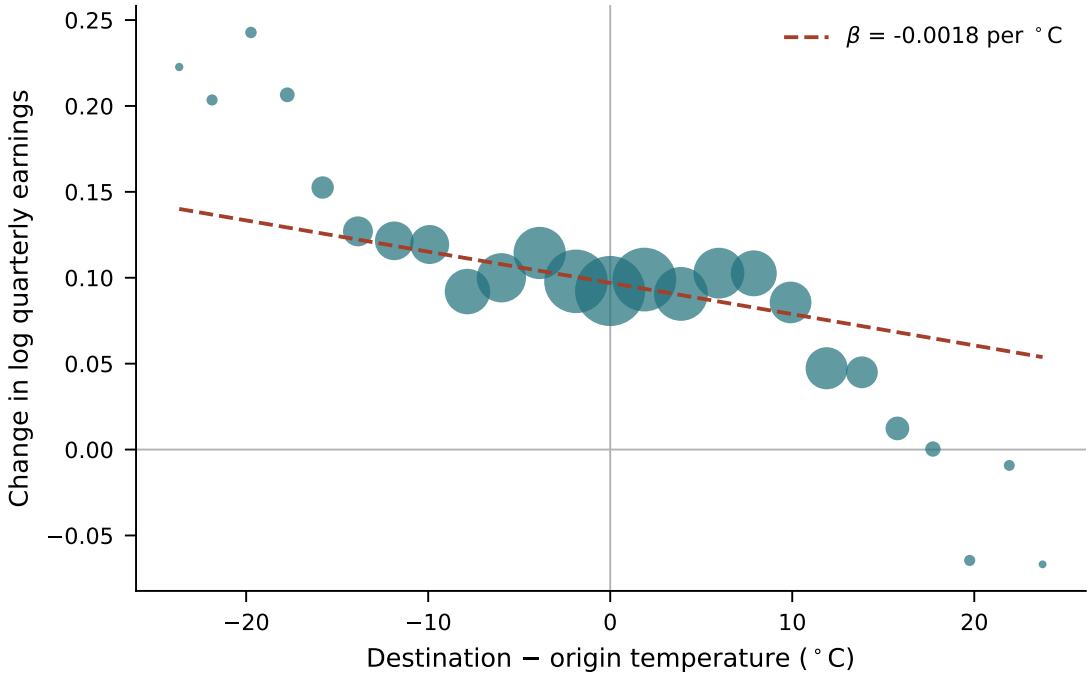


Figure 2: Earnings Changes After Moves. Binned scatter of the within-worker earnings change for U.S. internal migrants who change jobs within the same industry, against the destination-minus-origin difference in population-weighted mean annual temperature (°C). The earnings change is $y_{od\ell t} = \log \bar{w}^D - \log \bar{w}^O$, the log difference between average quarterly earnings on the destination and origin job. Each point is a worker-migrant-weighted mean within a temperature-difference bin (point area $\propto$ number of movers); the dashed line is the regression slope β from column (1) of Table 2 ($\beta \approx -0.0018$ per °C), estimated on the underlying unbinned data so that the line and the table coefficient coincide.

The figure also indicates that within a 25°C range of average temperature differences (−15°C to +10°C) the effect on earnings is remarkably linear, providing important context to the literature emphasizing the importance of nonlinearities in the relationship between average temperature and economic activity. Even taking the nonlinearities observed in the graph as representative of the true effects of large temperature changes on earnings, the effect is relevant only for changes that are an order of magnitude larger than those at which the literature has thus far considered nonlinear responses.

Climate variation is cross-sectional, meaning that any specification that preserves permanent temperature differences across origins and destinations also preserves other permanent differences that are correlated to temperature differences. For example, differences in the competitiveness of labor markets across origin and destination that could contribute to permanent wage differences and that are correlated with temperature differences would be subsumed into the temperature coefficient. To account for this possibility, we include

as control variables the most obvious observable economic and demographic differences — population size and density, and education as proxied by the difference in bachelors degrees. The results are reported in the second column of Table 2. Although the estimated β coefficient with these controls is somewhat higher, it is still very small, -0.0026.

The final column of Table 2 examines the relation between *temporary* weather shocks and productivity by including fixed effects for both the origin and destination. These fixed effects absorb all permanent origin and destination characteristics, including permanent climate differences. The resulting coefficient captures the relationship between wage changes and the annual weather-like component of the temperature difference, rather than the permanent climate gap. We find that the coefficient is negative and larger in magnitude than the coefficient in the first column. That is, there is a significantly negative relation between temporary weather effects and wages, likely indicating the short term disruption that occurs with the arrival of a temporary shock. This coefficient emphasizes our earlier point that using temporary weather shocks can lead to spurious inferences about the effect of climate change on productivity.

Our data also allow us to observe moves at the level of metro areas (CBSAs). The metro sample excludes LEHD nonmetropolitan areas, where agriculture and other outdoor work are likely to be more prevalent. The metro sample covers 76% of worker-migrant flows in the full dataset and the metro counties it draws on account for 86% percent of the U.S. population. The metro setting also helps address a possible omitted-variable concern that institutional, historical, cultural, and policy differences vary significantly at the state level. Further, metro areas are useful because they preserve substantially more geographic and climatic variation than the state-level data does. State averages aggregate together places with meaningfully different climates, such as the San Francisco Bay Area and Los Angeles. In addition, studying these metro areas separately, allows us to control for differences that result from people moving between rural and metropolitan areas.

	(1)	(2)	(3)
Temperature difference (°C)	-0.0014***	-0.0016***	-0.0043***
	(0.0002)	(0.0002)	(0.0005)
Log population difference		0.0147***	
		(0.0011)	
Population density difference		0.0000*	
		(0.0000)	
Bachelor's share difference		0.0025***	
		(0.0002)	
Mean moving wage premium	0.092	0.092	0.092
Year FE	Yes	Yes	Yes
Common-sector FE	Yes	Yes	Yes
Same-sector sample	Yes	Yes	Yes
State economic/demographic controls	No	Yes	No
Origin/destination CBSA FE	No	No	Yes
Observations	5,100,136	5,100,136	5,100,136
Metro-pair clusters	138,229	138,229	138,229
R^2	0.012	0.017	0.033

Table 3: Metropolitan estimates. Worker-migrant-weighted regressions of the within-worker, same-industry earnings change on the destination-minus-origin metro (CBSA) temperature difference (°C). Column (1) is the baseline; column (2) adds state economic and demographic controls; column (3) adds origin and destination CBSA fixed effects (temporary weather). Standard errors are clustered by metro pair. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

The first column of Table 3 repeats the main specification at the metro level. As before, the regressions are weighted by metro worker-migrants, standard errors are clustered by origin-destination CBSA pair and the sample is restricted to same-sector movers. The mean moving wage premium in the metro sample is almost identical to the state sample, 0.092. The effect of the climate shock on productivity, the β coefficient, is slightly smaller than at the state level, approximately -0.0014 per degree Celsius. Figure 3 plots the metro data binned by temperature changes. The pattern is similar to the state-level figure: the dense part of the distribution shows a shallow negative gradient, while extreme climate moves are noisier and more likely to reflect unusual origin-destination labor-market pairs. The coefficient estimate rises slightly, to -0.0016 when the state economic and demographic controls are added (second column of Table 3) [7]. The third column of Table 3 introduces origin-CBSA and destination-CBSA fixed effects. The effect of temporary weather shocks is very similar to the state sample. The metro evidence therefore mirrors the state-level pattern: permanent

[7]Metro Areas spanning multiple states are assigned the state controls for the primary state in the area's definition. For example "New York-Newark-Jersey City, NY-NJ-PA" is assigned to New York.

climate gradients are small, while weather-like temporary variation produces a larger but still modest coefficient.

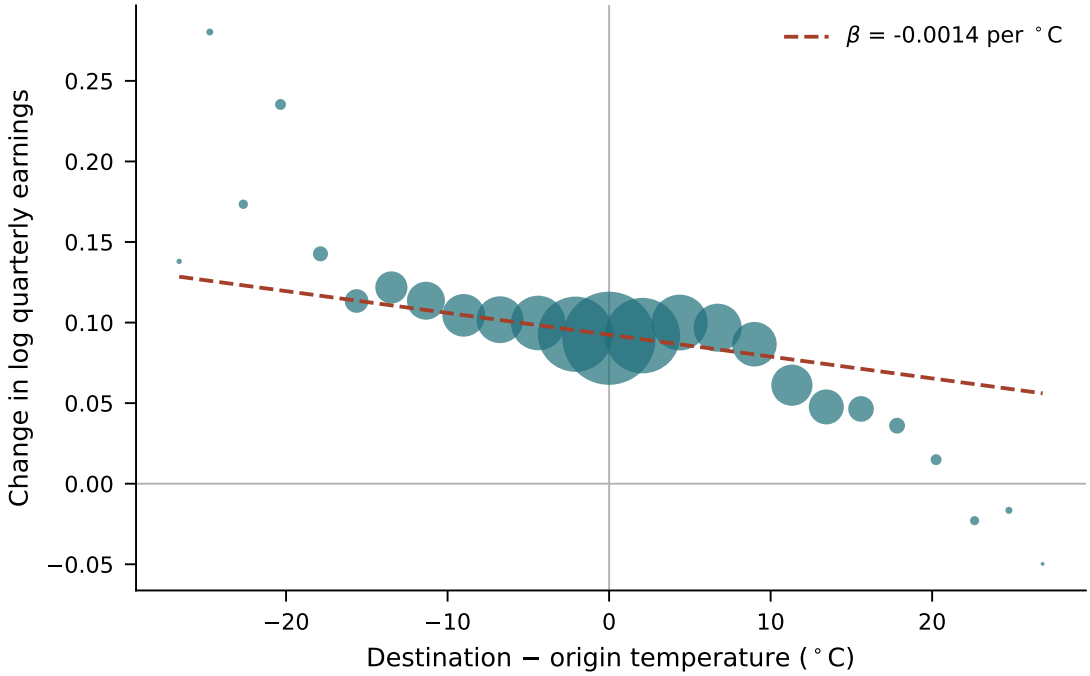


Figure 3: Earnings Changes After Moves at the Metropolitan Level. Binned scatter of the within-worker, same-industry earnings change for internal migrants against the destination-minus-origin difference in population-weighted mean annual temperature (°C), measured across metropolitan areas (CBSAs) rather than states. Each point is a worker-migrant-weighted mean within a bin (point area $\propto$ number of movers); the dashed line is the column (1) regression slope of Table 3 ($\beta \approx -0.0014$ per °C), estimated on the unbinned data. As in the state data, the permanent-climate gradient is small relative to the $\approx +9\%$ average move premium.

5 Taxation

The earnings statistics we use are gross wages reported to state unemployment-insurance systems and carry no adjustment for taxes. Because state income taxes vary between the origin and destination of a move, and the warmest destinations include several states with no income tax, one plausible reason for why people choose to move are lower income taxes. To the extent that temperature is correlated to the state income tax rates, this could affect our results. Therefore, in this section, we repeat the analysis using take home pay, rather than pretax wages.

We compute tax liabilities with the NBER Internet TAXSIM calculator.[8] For each group of migrants we take both pre-move and post-move earnings, annualise them by multiplying by four, and compute the tax owed on that income in the relevant state, filing as a single taxpayer with no dependents and with no income other than wages. Take-home pay is income minus the tax liability, converted back to a quarterly figure by dividing by four. A record's year and quarter refer to the quarter of the move, t, but origin earnings are measured in $t-1$ and destination earnings in $t+1$. We repeat the same aggregation as in our prior estimates with these new quarterly take home earnings.[9]

Calling TAXSIM once per cell would require roughly 47 million evaluations. We instead exploit the fact that, once the filer type is fixed, liability depends only on the state, the year and the level of wage income, and is a continuous piecewise-linear function of income. We therefore evaluate a grid of 1,281 income points for every state and year and interpolate between them, extending the final segment linearly above the top grid point.[10]

Let $\tau^{O}_{od\ell t}$ and $\tau^{D}_{od\ell t}$ denote the effective tax rates a cell faces on its origin and destination earnings. The net earnings change replaces each wage in equation (2) with its take-home counterpart,

$$y^{\text{net}}_{od\ell t} = \log\left[\left(1-\tau^{D}_{od\ell t}\right)\bar{w}^{D}_{od\ell t}\right] - \log\left[\left(1-\tau^{O}_{od\ell t}\right)\bar{w}^{O}_{od\ell t}\right]. \tag{5}$$

We use two different take home numbers, one which accounts only for state taxes, and the other which accounts for federal income tax and the employee share of payroll tax as well. We re-estimate the same-sector baseline and economic/demographic-control specifications of Tables 2 and 3, at both the state and metro level.

[8] TAXSIM version 35; see Feenberg and Coutts (1993) and `https://taxsim.nber.org`.

[9] Due to the progressivity of the tax schedule, our calculation based on average earnings ends up underestimating the true average tax liability. Unemployment-insurance wage records cover a single employer, so annualising one job's quarterly earnings ignores second jobs and spousal income and understates true bracket position. Finally, a metropolitan area spanning several states is assigned the first state in its title, so the metro estimates treat, for example, the New York area as entirely New York.

[10] Validated against direct per-cell TAXSIM calls on a stratified sample of real cell incomes drawn from 60 state–year blocks: the largest error in the total effective rate is 0.013 percentage points for incomes inside the grid. Cells above the top grid point, which carry 0.002% of worker-migrant weight, are extrapolated with a largest observed error of 0.023 percentage points.

	(1)	(2)	(3)	(4)
Temperature difference (°C), gross	-0.0018***	-0.0026***	-0.0014***	-0.0016***
	(0.0005)	(0.0005)	(0.0002)	(0.0002)
Temperature difference (°C), net of state tax	0.0002	-0.0008*	0.0004**	-0.0002
	(0.0004)	(0.0004)	(0.0002)	(0.0002)
Temperature difference (°C), net of all tax	0.0009**	-0.0000	0.0010***	0.0003**
	(0.0004)	(0.0004)	(0.0002)	(0.0002)
Geography	State	State	Metro	Metro
Mean moving wage premium, gross	0.096	0.096	0.092	0.092
Mean moving wage premium, net of all tax	0.089	0.089	0.084	0.084
Year FE	Yes	Yes	Yes	Yes
Common-sector FE	Yes	Yes	Yes	Yes
Same-sector sample	Yes	Yes	Yes	Yes
Economic/demographic controls	No	Yes	No	Yes
Observations	688,219	688,219	5,100,136	5,100,136
Origin–destination pair clusters	2,550	2,550	138,229	138,229
R^2	0.038	0.049	0.011	0.014

Table 4: After-Tax Earnings Changes After Moves Worker-migrant-weighted regressions of the within-worker, same-industry earnings change on the destination-minus-origin temperature difference (°C). Columns (1) and (2) are measured across states and columns (3) and (4) across metropolitan areas; columns (1) and (3) are the baseline specification and columns (2) and (4) add the economic and demographic controls. The "gross" rows reproduce columns (1) and (2) of Tables 2 and 3; the "net" rows replace earnings with take-home pay, taxing each cell at its own income (NBER TAXSIM, see text) before differencing. Standard errors are clustered by origin–destination pair. "Net of state tax" deducts state income tax only; "net of all tax" also deducts federal income tax and the employee share of FICA. Sample size, clusters and R^2 are for the net-of-all-tax specification, and every column reports a total R^2. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

Table 4 reports the results. The gross rows reproduce the baseline same-sector gradients of -0.0018 and -0.0026 at the state level and -0.0014 and -0.0016 at the metro level. Deducting state income tax attenuates every estimate. Accounting for all taxes removes the negative relation between temperature and wages. If anything, the resultant effect is actually positive for three of the four estimates. This suggests that our previous estimates might be driven by movers' preferences for lower tax states that happen to be correlated to temperature rather than a productivity drop.

6 Discussion

Our empirical object is a wage change, not a direct measure of physical output per worker. We use wages as a diagnostic for the labor-productivity channel because labor productivity should ultimately be reflected in earnings, but we fully acknowledge that wages are equilib-

rium objects. They reflect local competition, amenities, prices, job-match quality, moving costs, and sorting. Let the observed wage-temperature gradient, β be written as

$$\beta = \beta_P - \beta_A, \tag{6}$$

where β_P is the effect of climate on productivity and β_A is the compensating value of amenities correlated to climate. If warmer destinations provide positive amenity value for the marginal mover, then $\beta_A > 0$, and the observed wage gradient β is more negative on average than the underlying productivity gradient β_P. In that case, our negative wage-temperature coefficients are upper bounds on the magnitude of the pure productivity penalty.[11]

The climate-amenities evidence makes this interpretation empirically relevant in the U.S. setting. Population growth has historically tilted toward places with nicer weather (Rappaport, 2007). This fact is consistent with the descriptive fact that Americans have moved toward warmer parts of the country over time. If part of the wage penalty we estimate reflects compensation for amenities rather than productivity, then the absence of a large wage penalty is even stronger evidence against a large local productivity penalty from permanent warmth.

If the warmer destinations that attract U.S. movers carry positive amenity value, migration volume should tilt toward warmer places even after accounting for the obvious gravity force of city size. As we show formally in Appendix B, it does. Warmer metro areas draw more in-migration than colder ones of the same population, and within a given origin–destination pair the warmer of the two metros receives the larger flow even when the two are equally sized. This revealed preference implies that $\beta_A > 0$, which makes the small negative wage gradients we report upper bounds on the pure productivity penalty rather than point estimates of it. Consequently, even the small negative gradient we observe may overstate the pure productivity component.

Another concern is that moves to warmer places might partly capture retirement rather than the effect of temperature exposure. Our design already largely addresses this concern because we restrict the dataset to workers who move into a new job, and our main sample focuses on stable-to-stable job changes with full-quarter employment at both origin and

[11]The compensating-differentials and spatial-equilibrium literature has addressed many of these issues. In a Rosen-Roback (Rosen (1979); Roback (1982)) setting, locations differ in productivity and amenities, and wages and housing costs adjust jointly to clear local labor and housing markets. A location with high productivity can support high wages and high rents; a location with attractive amenities can support lower wages and higher rents. Modern climate-amenity papers apply this logic directly to temperature, valuing climate through wages, housing costs, and location choices rather than through wages alone (Albouy et al., 2016; Sinha et al., 2021; Rappaport, 2007). Migration models add a further complication: observed movers are selected workers who move only when expected gains, amenities, job matches, and moving costs jointly justify the move (Kennan and Walker, 2011).

destination. It does not, however, address semi-retirement — workers might move to a warmer climate anticipating eventual retirement with the intention of gradually working less. To assess the importance of semi-retirement, we can estimate the magnitude of the effect as a function of age. If semi-retirement is partly responsible for the negative coefficients, the largest negative coefficients should be concentrated among the oldest age groups and there should be a sharp discontinuity around retirement age. That is what we find.

Age band	State $\hat{\beta}$	State SE	Metro $\hat{\beta}$	Metro SE
14-18	0.0014	(0.0009)	0.0027***	(0.0004)
19-21	-0.0005	(0.0007)	0.0031***	(0.0003)
22-24	0.0001	(0.0006)	0.0029***	(0.0003)
25-34	-0.0011**	(0.0005)	-0.0010***	(0.0002)
35-44	-0.0021***	(0.0006)	-0.0017***	(0.0002)
45-54	-0.0020***	(0.0006)	-0.0021***	(0.0002)
55-64	-0.0012**	(0.0006)	-0.0030***	(0.0002)
65+	-0.0038***	(0.0008)	-0.0033***	(0.0003)
Exclude 65+ (under 65)	-0.0015***	(0.0005)	-0.0010***	(0.0002)

Table 5: Wage–temperature gradient and Age Each cell is the temperature coefficient $\hat{\beta}$ (per °C) from the baseline specification of equation (4), estimated separately within each age band. The metro columns use the same-sector age panel with year and common-sector fixed effects, matching column (1) of Table 3. The state columns use the LEHD age-margin panel, which carries no sector detail and therefore absorbs year fixed effects only. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

We reestimate (4) over a range of age cohorts and report the results in Table 5. The age pattern shows a clear life-cycle relationship: the youngest workers have insignificant coefficients in the state panel and positive coefficients in the metro panel, while prime-age and older groups have negative coefficients in both panels that become more negative with age, with the steepest gradient in the 65+ band. The final row of Table 5 indicates that while semi-retirement appears to contribute to explaining the negative coefficients, it does not fully explain the effect. Pooling all movers under 65 – dropping the 65+ band entirely – leaves the coefficient negative and statistically significant, at about -0.0015 per °C at the state level and -0.0010 per °C at the metro level, comparable to the prime-age bands. Even after excluding the age group where the gradient is steepest, and where retirement sorting is most likely, the negative relationship persists.

7 Conclusion

This paper studies whether permanent differences in local temperature are associated with large changes in worker earnings. Using U.S. internal migrants, we find that workers moving to warmer places experience slightly lower earnings growth. This result is found to be robust across state and metro specifications, richer controls, same-sector moves. The sign is broadly consistent with the idea that hotter climates can reduce labor productivity but it is inconsistent with the large wage-productivity losses that would be needed to explain the large GDP damage estimates documented in the prior literature.

The existing literature relies on using the effect of temporary weather shocks on GDP to infer the effect of permanent climate change on economic activity. Temporary weather shocks and permanent climate differences are not the same object. Temporary shocks can disrupt production before workers, firms, infrastructure, and local prices have adjusted. Permanent climate exposure, by contrast, allows adaptation and other equilibrium responses. The fact that mover earnings only show a very small temperature gradient cautions against extrapolating large macroeconomic damages from temporary temperature shocks without a clearer account of the mechanisms and adjustment margins involved in equilibrium (Dell et al., 2014; Hsiang, 2016; Tol, 2021; Dong et al., 2025; Nath et al., 2025).

Because we cannot observe a randomized allocation of workers across climates, and movers are selected in ways that may matter for earnings growth, our results cannot and should not be read as a full estimate of the impact of climate change on productivity and GDP. Despite these limitations, our findings are nevertheless informative. Given the large damage estimates found in the climate economics literature, if local labor productivity is an important channel driving these estimates, we should expect to see some trace of it in the earnings of workers who move across large and permanent climate gradients. Movers from cold to warm places, and from warm to cold places, experience changes in local climate exposure that mimic the permanent climate differences that the existing literature models and measures. This makes our setting a useful diagnostic, even if it is not a complete research design to assess total climate damage.

One possible interpretation is that local labor productivity is not the main channel through which climate change damages economic output. This is a reasonable position to take, but since labor productivity is an important driver of economic activity, and it is largely unaffected by climate, to make this argument effectively, the alternative channels generating the damages must be specified and measured directly. The research that purports to find large magnitudes cannot simply rely on the intuitive appeal of heat lowering productivity. Before policy makers rely on those estimates, the absence of large local wage effects has to

be explained.

Our findings do not show that climate change is harmless, nor do we estimate the total welfare cost of climate change. They show something narrower but important: one central channel often invoked in support of large climate damages, local labor productivity, does not appear large enough to support the aggregate estimates in the existing literature.

References

Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff, "Climate Amenities, Climate Change, and American Quality of Life," *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (1), 205–246.

Aragon, Fernando M., Francisco Oteiza, and Juan Pablo Rud, "Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat," *American Economic Journal: Economic Policy*, 2021.

Barrage, Lint and William Nordhaus, "Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model," *Proceedings of the National Academy of Sciences*, 2024.

Bilal, Adrien and Diego R Kanzig, "The Macroeconomic Impact of Climate Change: Global Versus Local Temperature," *The Quarterly Journal of Economics*, 05 2026, *141* (2), 889–944.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel, "Global Non-linear Effect of Temperature on Economic Production," *Nature*, 2015, *527*, 235–239.

Carleton, Tamma A. and Solomon M. Hsiang, "Social and Economic Impacts of Climate," *Science*, 2016, *353*.

Casey, Gregory, Stephie Fried, and Matthew Gibson, "Impact of U.S. Labor Productivity Losses from Extreme Heat," *FRBSF Economic Letter*, 2024. Federal Reserve Bank of San Francisco Economic Letter.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature*, 2014, *52* (3), 740–798.

Deryugina, Tatyana and Solomon M. Hsiang, "Does the Environment Still Matter? Daily Temperature and Income in the United States," *NBER Working Paper*, 2014.

Dong, Jinchi, Richard S. J. Tol, and Jinnan Wang, "The Effects of Climate and Weather on Economic Output: Evidence from Global Subnational Data," *Working paper*, 2025.

Gospodinov, Nikolay, Ignacio Lopez Gaffney, and Serena Ng, "The Economic Impact of Low- and High-Frequency Temperature Changes," *Working paper*, 2025.

Hsiang, Solomon, "Climate Econometrics," *Annual Review of Resource Economics*, 2016.

_ , Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D. J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser, "Estimating Economic Damage from Climate Change in the United States," *Science*, 2017, *356* (6345), 1362–1369.

Ireland, Andrew, David Johnston, and Rachel Knott, "Impacts of Extreme Heat on Labor Force Dynamics," *Working paper*, 2024.

Kahn, Matthew E., Kamiar Mohaddes, Ryan N. C. Ng, M. Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang, "Long-Term Macroeconomic Effects of Climate Change: A Cross-Country Analysis," *Energy Economics*, 2021, *104.*

Kalkuhl, Matthias and Leonie Wenz, "The Impact of Climate Conditions on Economic Production: Evidence from a Global Panel of Regions," *Journal of Environmental Economics and Management*, 2020, *103.*

Kennan, John and James R. Walker, "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 2011, *79* (1), 211–251.

Moore, Frances C. and Delavane B. Diaz, "Temperature Impacts on Economic Growth Warrant Stringent Mitigation Policy," *Nature Climate Change*, 2015, *5*, 127–131.

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow, "How Much Will Global Warming Cool Global Growth?," *Working paper*, 2025.

Neal, Timothy, Ben R. Newell, and Andy Pitman, "Reconsidering the Macroeconomic Damage of Severe Warming," *Environmental Research Letters*, 2025, *20* (4), 044029.

Ortiz-Bobea, Ariel, Toby R. Ault, Carlos M. Carrillo, Robert G. Chambers, and David B. Lobell, "Anthropogenic Climate Change Has Slowed Global Agricultural Productivity Growth," *Nature Climate Change*, 2021.

Rappaport, Jordan, "Moving to Nice Weather," *Regional Science and Urban Economics*, 2007.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Wage-Based Indexes of Urban Quality of Life," in Peter Mieszkowski and Mahlon Straszheim, eds., *Current Issues in Urban Economics*, Baltimore: Johns Hopkins University Press, 1979, pp. 74–104.

Sinha, Paramita, Martha Caulkins, and Maureen Cropper, "The Value of Climate Amenities: A Comparison of Hedonic and Discrete Choice Approaches," *Journal of Urban Economics*, 2021, *126.*

Somanathan, E., Rohini Somanathan, Anant Sudarshan, and Meenu Tewari, "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing," *Journal of Political Economy*, 2021.

Tol, Richard S. J., "The Economic Impact of Weather and Climate," *Working paper*, 2021.

_ , "A Meta-Analysis of the Total Economic Impact of Climate Change," *Working paper*, 2022.

Waidelich, Paul, Fulden Batibeniz, James Rising, Jarmo S. Kikstra, and Sonia I. Seneviratne, "Climate Damage Projections beyond Annual Temperature," *Nature Climate Change*, 2024, *14*, 592–599.

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

A Robustness

In this section we demonstrate that our conclusions are not sensitive to sample definition, transition definition, climate-variable construction, outliers, geography, temperature weights, or inference convention.

A.1 Broadening Beyond Same-Sector Moves

The same-sector restriction gives the cleanest interpretation of the wage change: the comparison is the effect of the climate change within the workers industry and not contaminated by an endogenous choice to switch industries. The downside of this restriction is that it excludes workers who change sectors as they move. We therefore re-estimate the state and metro specifications on the broader samples that include sector switchers, both without sector controls and with destination-sector controls. The results are reported in Table 6. The coefficients are stable across these broader samples.

	State		Metropolitan	
Sample and sector controls	(1)	(2)	(3)	(4)
Same-sector + sector FE	-0.0018***	-0.0026***	-0.0014***	-0.0016***
	(0.0005)	(0.0005)	(0.0002)	(0.0002)
Destination-sector FE	-0.0019***	-0.0027***	-0.0012***	-0.0012***
	(0.0005)	(0.0005)	(0.0002)	(0.0003)
No sector FE	-0.0019***	-0.0028***	-0.0012***	-0.0013***
	(0.0005)	(0.0005)	(0.0002)	(0.0003)

Table 6: Temperature Gradient Adding Sector-Switching Movers. The first row reproduces columns (1) and (2) of Tables 2 and 3 by restricting the sample to same-sector movers with sector fixed effects. The second row expands the sample to include people who move across sectors, and the third row adds destination-sector fixed effects to the broadened sample regression. Columns (1) and (2) are measured across states and columns (3) and (4) across metropolitan areas. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

A.2 Stable Transitions and Short Nonemployment Gaps

The main analysis uses stable-to-stable job changes, measured by LEHD's EES transition count. A concern is that this excludes movers who experience a short nonemployment gap, and that these workers might have different wage-temperature gradients. To address this concern we expand the data to include workers who experience a short nonemployment gap between the two jobs (the LEHD AQHire-S transitions). In the expanded sample both the

worker count and the worker-weighted average earnings are recomputed over the pooled set of stable (EES) and short-gap (AQHire-S) transitions:

$$N^{AS}_{od\ell t} = \text{EES}_{od\ell t} + \text{AQHireS}_{od\ell t}, \tag{7}$$

and the pooled average origin earnings (destination earnings are defined analogously) are the worker-weighted average across the two transition types,

$$\bar{w}^{O,AS}_{od\ell t} = \frac{\text{EES}_{od\ell t}\, \bar{w}^{O,\text{EES}}_{od\ell t} + \text{AQHireS}_{od\ell t}\, \bar{w}^{O,\text{AQ}}_{od\ell t}}{\text{EES}_{od\ell t} + \text{AQHireS}_{od\ell t}}. \tag{8}$$

The estimates in Table 7 change very little when short-gap stable transitions are added. This suggests that the main result is not driven by the narrow stable-to-stable transition definition.

Sample and specification	$\hat{\beta}$	SE	N (cells)
Same-sector sample, stable jobs only	-0.00182***	(0.00046)	688,219
All-mover sample, stable jobs only	-0.00188***	(0.00048)	882,365
All-mover sample, adding short-gap moves	-0.00173***	(0.00050)	942,655

Table 7: Temperature Gradient Adding Short-Gap (AQHire-S) Moves. The data vintage which includes short-gap movers does not contain the sector of employment at the origin, so we display 3 different samples to bridge from the main sample to the sample including movers who had a short gap in their employment. The first row reproduces columns (1) and (2) of Tables 2 and 3, the second row broadens the sample to include movers who swap sectors while the third row broadens it further to include short-gap movers. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

A.3 Alternative Temperature Weights and Measures

The main temperature regressor uses population-weighted annual mean temperature. For robustness, we consider alternative climate-metrics, as annual mean temperature may be too crude if climate damages operate through extreme heat or cold exposure.

We replace average temperature with cooling degree days (CDD), the number of °F the average temperature exceeds 65°F, heating degree days (HDD), the number of °F the average temperature is below 65°F, and the number of hot days, those on which the maximum

temperature exceeds 30°C.[12] In the empirical specification we replace D_{odt} with

$$C_{odt} = \text{CDD}_{dt} - \text{CDD}_{ot}, \qquad H_{odt} = \text{HDD}_{dt} - \text{HDD}_{ot}, \qquad N_{odt} = \text{Hot}_{dt} - \text{Hot}_{ot}, \quad (9)$$

and rerun our analysis. The results are reported in Table 8 qualitatively similar to the baseline temperature results. The two heat measures carry negative gradients and the cold measure a positive one. This suggests that the small wage-temperature gradient is not an artifact of using annual mean temperature rather than a more standard heat- or cold-exposure measure.

Climate measure	$\hat{\beta}$ (raw)	SE	Per SD	N
State				
Mean temperature (°C)	-0.0018793***	(0.0004724)	-0.0101	686,438
Cooling degree days (base 65°F)	-0.0000117***	(0.0000020)	-0.0142	686,438
Heating degree days (base 65°F)	0.0000032***	(0.0000011)	0.0076	686,438
Hot days ($T^{\max} > 30$°C)	-0.0002325***	(0.0000408)	-0.0138	686,438
Metropolitan				
Mean temperature (°C)	-0.0014802***	(0.0002153)	-0.0060	5,103,434
Cooling degree days (base 65°F)	-0.0000098***	(0.0000009)	-0.0098	5,103,434
Heating degree days (base 65°F)	0.0000019***	(0.0000005)	0.0034	5,103,434
Hot days ($T^{\max} > 30$°C)	-0.0002120***	(0.0000171)	-0.0105	5,103,434

Table 8: Temperature Gradient Based on Heat- and Cold-Exposure Measures. We re-estimate the baseline specification from 2 and 3 with four different independent variables. All of these variables are built from daily NOAA nClimGrid county data since only the daily dataset allows us to compute the additional variables. The variables are population-weighted identically to make the rows mutually comparable and because that product covers only the contiguous states, Alaska and Hawaii are excluded. The "Per SD" column displays each coefficient multiplied by the within-sample standard deviation of its own regressor and is the column to read across: every heat, cold and mean-temperature gradient falls in a narrow band around one percent of the move premium. The small gradient is therefore not an artefact of summarising climate with an annual mean. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

To make the temperature, degree-day, and hot-day coefficients comparable, the "Per SD" column of Table 8 reports *standardized* coefficients: each raw coefficient is multiplied by the in-sample standard deviation of its own regressor, so it measures the effect on the

[12] All three are computed from daily NOAA nClimGrid county data and population-weighted as mean temperature is. The 30°C threshold is the one at which Graff Zivin and Neidell (2014) find labor supply falls in climate-exposed industries. Alaska and Hawaii are excluded, since the daily product covers only the contiguous states.

earnings change of a one-standard-deviation larger destination-minus-origin climate difference. On this common scale the four climate measures are of the same order of magnitude. A one-standard-deviation warmer move (about 5.4°C) is associated with roughly a 1.0% lower earnings change in the baseline specification; the cooling-degree-day difference gives a comparable −1.4%, the hot-day difference −1.4%, and the heating-degree-day difference an oppositely-signed +0.8%. A colder destination (more heating degree days) raises earnings, in line with our temperature result. A one-standard-deviation move on the hot-day measure is about 59 additional days a year above 30°C, so even a large change in threshold exposure moves earnings very little.

A.4 Outliers, Trimming, and Alaska/Hawaii

The binned scatterplots show that the largest temperature moves are unusual. Some of these observations are economically meaningful, but they may also reflect special labor markets, especially Alaska, Hawaii, and extractive industries. We therefore re-estimate the main specifications after trimming the absolute temperature difference and after excluding Alaska and Hawaii.

The estimates remain small across these restrictions. The tails help explain some of the visual shape of the binned scatter, but they do not generate a hidden large negative coefficient. If anything, the extreme corridors are useful because they reveal the role of sectoral and geographic selection in the largest climate moves.

	State			Metro		
Cap	% kept	$\hat{\beta}$	SE	% kept	$\hat{\beta}$	SE
$\lvert\Delta T\rvert \leq 5$	67.8%	-0.001695	(0.001060)	83.7%	-0.000086	(0.000484)
$\lvert\Delta T\rvert \leq 10$	90.5%	-0.000418	(0.000541)	95.1%	-0.000356	(0.000232)
$\lvert\Delta T\rvert \leq 15$	99.3%	-0.001604***	(0.000490)	99.6%	-0.001255***	(0.000226)
$\lvert\Delta T\rvert \leq 20$	100.0%	-0.001792***	(0.000462)	100.0%	-0.001332***	(0.000211)
no cap	100.0%	-0.001819***	(0.000459)	100.0%	-0.001354***	(0.000210)

Table 9: Gradient Estimates on the Trimmed Sample. State and metropolitan temperature coefficients (per °C) from the baseline specification re-estimated as the sample is restricted to progressively smaller absolute origin–destination temperature gaps. Each row keeps only cells with $|\Delta T|$ below the stated cap, in degrees Celsius; "% kept" is the share of full-sample worker-migrants that survives the restriction. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

Restriction	State $\hat{\beta}$	SE	N	Metro $\hat{\beta}$	SE	N
Full sample	-0.00182***	(0.00046)	688,219	-0.00135***	(0.00021)	5,100,136
Drop AK+HI as destination	-0.00178***	(0.00047)	672,779	-0.00135***	(0.00021)	5,068,216
Drop AK+HI as origin	-0.00182***	(0.00047)	673,586	-0.00138***	(0.00021)	5,069,789
Drop AK+HI on either side	-0.00179***	(0.00048)	658,445	-0.00138***	(0.00022)	5,039,381

Table 10: Temperature Gradient Excluding Alaska and Hawaii. Baseline temperature coefficients (per °C) at the state and metropolitan level after removing the two non-contiguous states, which sit at the extremes of the climate distribution and carry the extractive corridors discussed in Appendix A.4. Rows drop Alaska and Hawaii as destination, as origin, and on either side; the first row is the unrestricted sample of Tables 2 and 3. Removing them changes the state gradient by at most 0.00004 and the metro gradient by at most 0.00003. *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

A.5 Alternative Clustering of Standard Errors

The baseline tables cluster standard errors by origin-destination corridor. This accounts for serial and sector-level dependence within a migration corridor, but the temperature regressor also contains origin-year and destination-year components that are shared across many corridors. For example, a hot year in a destination state enters all flows into that state in that year, and a hot year in an origin state enters all flows out of that state in that year.

Table 11 reports the coefficient under eight covariance estimators: HC1 i.e. heteroskedastic with a sample size adjustment, the origin-destination corridor used as the baseline, the origin-year and destination-year cells, the origin and destination geographies pooling all years, and a multiway estimator combining origin-destination pair, origin-year, and destination-year clusters. Clustering on the geography alone is the most permissive of these, since it allows the errors of a given location to be arbitrarily correlated across the whole sample period rather than only within a year. The alternative clustering choices change the precision of the estimates but not their economic magnitude. Our paper uses the origin-destination corridor as the baseline because the regressor varies at a corridor-year level.

With alternative standard errors, we can also check whether the coefficients cited in the literature lie in our confidence intervals at more conservative clustering methods. We find that the standard error would have to be roughly eight times its largest value in the state block, and twenty times in the metropolitan block, before the ninety-five percent interval reached the two percent per °C at the low end of the existing literature.

Specification	$\hat{\beta}$	HC1	OD pair	Origin-year	Dest.-year	Origin	Destination	Multiway	N
State									
Baseline	-0.00182	0.00009	0.00046	0.00018	0.00027	0.00062	0.00111	0.00053	688,219
+ econ/demog.	-0.00263	0.00011	0.00047	0.00023	0.00028	0.00084	0.00105	0.00056	688,219
Metropolitan									
Baseline	-0.00135	0.00005	0.00021	0.00011	0.00013	0.00037	0.00044	0.00025	5,100,136
+ econ/demog.	-0.00165	0.00006	0.00022	0.00014	0.00013	0.00048	0.00044	0.00027	5,100,136

Table 11: Alternative Standard-Error Clusterings. The temperature coefficients of Tables 2 and 3 under eight covariance estimators. The point estimate is identical across columns by construction; only the variance changes. "Origin" and "Destination" refer to the state in the upper block and to the CBSA in the lower one. HC1 treats every cell as independent, which the regressor makes untenable: D_{odt} does not vary across sectors within a corridor-year, so the roughly 270 cells in a state corridor carry far less than 270 observations' worth of independent information. "OD pair" is the paper's baseline and clusters by origin–destination corridor (2,550 state pairs; 138,229 metro pairs). "Origin-year" and "Dest.-year" cluster by the relevant geography-year cell, capturing the fact that a given year's temperature at one end of a move is shared across every corridor touching it. "Origin" and "Destination" cluster by geography alone, pooling all years (51 states; 386 CBSAs); these permit arbitrary serial correlation within a location over the 2000–2023 window and are the most conservative estimators reported, though with 51 clusters the state block is also the most exposed to few-cluster bias. "Multiway" combines OD pair, origin-year and destination-year clustering by the Cameron–Gelbach–Miller inclusion–exclusion formula; because each of its three dimensions conditions on either the year or the corridor, it does not nest the geography-level columns and is not the widest estimator here.

B Revealed Preference for Warmer Destinations

In Section 6 we decomposed the wage gradient estimate, $\beta = \beta_P - \beta_A$, where β_A is the compensating value of correlated amenities. The upper-bound interpretation of our small negative gradients rests on $\beta_A > 0$: warmer destinations must, on net, be desirable to the marginal mover. This appendix documents the within-dataset evidence behind that premise.

The logic is the standard spatial-equilibrium one. In a Rosen–Roback world, a location with attractive amenities can clear its labor market at lower wages precisely because workers are willing to move there for the amenity. The migration counterpart is that, holding the economic pull of a place fixed, a desirable amenity should attract larger gross flows. Temperature is the candidate amenity, and metropolitan migration flows are the behavior. The empirical challenge is that the dominant determinant of bilateral flows is gravity—large places send and receive more movers—and city size is correlated with climate. Our specifications are therefore designed to separate warmth from size.

Destination gravity. The first specification aggregates flows to the destination metro and asks whether warmer metros attract more in-migration than colder metros of the same size:

$$\log M_{dt} = \gamma\, T_d + \delta \log P_{dt} + \tau_t + u_{dt}, \tag{10}$$

where M_{dt} is total J2J in-migration to metro d in year t, T_d is the destination's mean temperature (in °C), P_{dt} is its population, and τ_t are year fixed effects. A positive γ says warmth draws movers conditional on size.

Within-pair flows. Equation (10) still compares different metros, which differ in much more than temperature and size. The second specification removes all symmetric corridor characteristics—combined size, distance, historical and cultural ties—by comparing the two directions of the same origin–destination pair:

$$\log F_{odt} = \gamma\,(T_d - T_o) + \delta_1 \log P_{dt} + \delta_2 \log P_{ot} + \mu_{\{o,d\}} + \tau_t + \varepsilon_{odt}, \tag{11}$$

where F_{odt} is the flow from o to d and $\mu_{\{o,d\}}$ is an *unordered* metro-pair fixed effect that is common to both directions. Identification comes entirely from the asymmetry within a pair: does the warmer of the two metros receive the larger flow? Including both metros' populations absorbs the remaining within-pair size confound, so γ is the warmth tilt at fixed size.

This size confound is not incidental. Uncontrolled, only about 46% of corridor-years send more movers toward the warmer metro, because large *cold* metros (New York, Chicago, Boston) win many corridors on population alone. The picture reverses once size is held constant: at equal population the warmer member of a pair receives roughly 5% more movers than the colder member (reported in the caption of Table 12 as the intercept of the net log flow toward the warmer metro on the log size gap). Table 12 collects the estimates: the temperature coefficient is positive and precisely estimated in the gravity specification, in the within-pair specification with and without size controls, and in a variant that absorbs origin-metro rather than pair fixed effects. The point estimates are small—this is migration in equilibrium, not a natural experiment—but their sign is unambiguous and stable.

	(1)	(2)	(3)	(4)
Temperature (°C)	0.0168***	0.0047***	0.0044***	0.0124***
	(0.0036)	(0.0001)	(0.0001)	(0.0010)
Log destination population	0.8594***		0.7079***	0.5183***
	(0.0155)		(0.0080)	(0.0042)
Log origin population			0.6960***	
			(0.0079)	
Year FE	Yes	Yes	Yes	Yes
Metro-pair FE	No	Yes	Yes	No
Origin-metro FE	No	No	No	Yes
Clusters	386	73,649	73,649	73,649
Observations	9,032	1,888,612	1,888,612	1,888,612

Table 12: Warmth as an Amenity. Gravity and within-pair regressions of bilateral metropolitan migration flows on temperature, holding population fixed, testing the premise $\beta_A > 0$ on which the upper-bound reading of Section 6 rests. The sample is all-mover metro-to-metro flows, 2000–2023, aggregated over destination sector. Column (1) is at the destination-metro × year level: the dependent variable is the log of total in-migration and "Temperature" is the destination's own mean temperature. Columns (2)–(4) are at the directed corridor × year level: the dependent variable is the log flow from origin to destination and "Temperature" is the destination-minus-origin difference. Column (2) adds an unordered metro-pair fixed effect common to both directions, so identification comes only from the asymmetry within a pair; column (3) adds both metros' populations to absorb the within-pair size confound; column (4) absorbs origin-metro rather than pair fixed effects. Temperature is in degrees Celsius. Unlike the wage regressions these are unweighted — each cell is one gravity unit — with standard errors clustered on the destination CBSA in column (1) and on the metro pair otherwise. The temperature coefficient is positive and precisely estimated throughout: at equal metro size the warmer member of a corridor receives $e^{0.049} - 1 \approx 5$ percent more movers than the colder member (SE 0.001). *,** and *** indicates significance at the 10%, 5% and 1% confidence level.

Note that temperature is a fixed metropolitan characteristic, so "warmth" is inevitably bundled with correlated consumption amenities (sunshine, lower heating costs, Sun Belt growth) and with the sorting of people and industries. We cannot, and do not, isolate the pure marginal value of degrees Celsius. But β_A in Section 6 is defined as the value of *correlated amenities*, so the bundled object is exactly the one the decomposition requires. Movers are selected, and gross flows reflect job opportunities and moving costs as well as amenities which is why this result should be interpreted merely as corroborating evidence for $\beta_A > 0$. With those qualifications, the direction of migration in our own data is consistent with warmer destinations being, on net, a positive amenity for U.S. movers.