The Transnational Effect of Temperature on Remittances

Gonzalo Ares de Parga-Regalado[1,3], Daniel Osuna-Gomez[2], and Luis Sarmiento[2,4,5,6, †]

[1]Bocconi University
[2]Banco de México
[3]LEAP: Laboratory for Effective Anti-Poverty Policies
[4]CMCC: Euro-Mediterranean Center on Climate Change
[5]EIEE: European Institute on Economics and the Environment
[6]Department of Management, Technology, and Economics, ETH Zurich

March 24, 2026

Abstract

We study how temperature shocks on both sides of the US–Mexico migration corridor affect remittances. Using quarterly Mexican municipality-level remittance data and Mexican undocumented migrant network data linking communities of origin to US destinations, combined with weather data, we identify two opposing forces. Extreme temperatures in Mexico increase remittance inflows, consistent with an insurance motive. In contrast, temperature extremes in US destinations reduce remittances. These effects operate through migrant labor markets: heat and cold reduce the earnings of immigrants.

Keywords: Environmental Inequality, Climate change, Migrant Networks, Remittances.
JEL Classification Numbers: F22, F24, J61, O15, Q54, Q56

We would like to thank Thomas Pearson, Marta Prato and participants at the seminar at Banco de Mexico. The data on the Matriculas Consulares de Alta Seguridad was obtained through the Unidad de Transparencia of the Secretaria de Relaciones Exteriores. The opinions and views expressed herein are those of the authors and do not necessarily represent the views of Banco de México, its Board of Governors, or The Government of Mexico. Luis Sarmiento is the corresponding author and can be contacted at lsarmiento@ethz.ch.

1 Introduction

Remittances to low- and middle-income countries reached $656 billion in 2022, comparable to the combined flows of foreign direct investment and official development assistance (World Bank, 2023). Because remittances connect migrants' earnings abroad with the welfare of households at home, economic shocks on either side of the migration corridor can alter their flow through distinct channels. The migration literature highlights two main motives: altruism and self-insurance (Lucas and Stark, 1985; Faini, 1994; Agarwal and Horowitz, 2002). Under altruistic motives, adverse conditions in origin communities increase remittances as migrants provide financial support for their families; under self-insurance motives, hardships in destination locations reduce them as migrants conserve resources for their own security. Existing work documents that international transfers function as an income buffer against negative shocks in migrants' origin communities, including extreme weather events (Yang and Choi, 2007; Randazzo et al., 2023; Smith, 2025), but evidence on destination-side responses remains limited (e.g., Joseph et al. (2018)). Despite growing evidence that climate variability affects virtually every dimension of economic life, from mortality (Carleton et al., 2022) to cognitive performance (Burke et al., 2015b) to long-run growth (Dell et al., 2012; Burke et al., 2015a), it remains unclear how temperature shocks in migrants' origin and destination communities jointly shape remittance flows as a coping response. In this paper, we study this question by leveraging exposure on both sides of the US–Mexico migrant network.

To do so, we use twelve years of quarterly remittance data covering all Mexican municipalities. To study shocks in origin communities, we merge these data with local temperature measures and exploit within-municipality and within-quarter temperature variation in a Poisson pseudo-maximum likelihood (PPML) framework that includes a rich set of fixed effects and weather controls. To analyze shocks affecting migrants in destination communities, we link US county temperatures to administrative data from the *Matrículas Consulares de Alta Seguridad* (MCAS) program, connecting migrants' counties of residence in the United States with their municipalities of origin in Mexico. We then construct a Bartik-style measure of each municipality's exposure to temperature shocks experienced by its migrants in the United States and regress remittance flows on it. Following Borusyak and Hull (2023), we recenter county-level shocks by absorbing municipality-by-time fixed effects before aggregating with migrant shares, ensuring that identification relies on idiosyncratic, within-network temperature variation rather than aggregate climate realizations.

Both identification strategies reveal strongly non-linear relationships between temperature and remittances. For communities of origin in Mexico, we estimate a U-shaped exposure-response. Remittances are minimized at approximately 20°C and increase with deviations in either direction. One additional day per quarter above 30°C raises per-capita remittances by 0.048%, while one additional day below 10°C raises them by 0.25%. Aggregating over the observed temperature distribution, the median municipality receives approximately 1.6% more in remittances than it would due to systematic deviations from the thermal optimum. For destination communities in the US, we find an inverted-U. Remittances are maximized near 17°C in US counties and decline with temperature extremes. One additional share-weighted day above 30°C reduces remittances by 0.21%, while one additional share-weighted day below 5°C reduces them by 0.26%. Aggregating across the temperature distribution, the median Mexican municipality loses approximately 3.6% of its remittances due to non-optimal temperatures in the US. In dollar terms, temperature shocks in destination communities amount to approximately $1.7 billion per year in aggregate foregone remittances, roughly 40% of the annual budget of Mexico's flagship anti-poverty program, Prospera (Parker and Todd, 2017), or similar to

the annual budget of a small Mexican state like Tlaxcala (Congreso del Estado de Tlaxcala, 2024).

The results related to temperature shocks in Mexico are consistent with the idea that remittances serve as an insurance mechanism for migrant households facing negative economic shocks (Yang and Choi, 2007). This is explained by the fact that extreme temperatures can increase household expenditure (Justino et al., 2025; Mohapatra et al., 2012; Sarmiento et al., 2025; Randazzo et al., 2023). In addition, extreme temperatures can reduce the income of rural communities that rely on agricultural production (Michel-Cuello and Aguilar-Rivera, 2022; Estrada et al., 2022), which migrant-sending communities in Mexico tend to be disproportionately dependent on (Riosmena and Massey, 2012). We provide additional suggestive evidence on the effect of temperature extremes on household income. Using the Mexican National Employment Survey (ENOE), we show that municipalities with above-average proportion of MCAS face a drop in their labor supply in response to extreme temperatures -implying a reduction in available income. Taken together, our evidence and previous work indicate that remittances may increase in response to temperature shocks as communities of origin seek to smooth consumption.

Regarding our results on transnational temperature shocks, we use the Current Population Survey (CPS) to provide suggestive evidence on a possible mechanism. We find that temperature shocks in the United States reduce the number of hours worked by Mexican migrants, suggesting that extreme temperatures diminish migrants' labor supply and total income, and therefore their ability to send remittances back home.[1] In fact, we find that Mexican migrants' labor supply is more responsive to extreme temperatures than the labor supply of other foreign-born migrants and U.S.-born workers. In particular, cold temperatures reduce the labor supply of Mexican migrants but have no measurable effect on other workers.[2]

Our main results are robust to different checks. For example, our results are stable across different specifications with different temperature functions or control variables. Our transnational results are not driven by municipalities near the US-Mexico border, and are robust to common concerns related to the use of Bartik-style instruments (e.g., see Jaeger et al. (2018); Adao et al. (2019); Goldsmith-Pinkham et al. (2020)). We also address the binning bias concern raised by Jones et al. (2026) by constructing counterfactual expected temperature bin counts from estimated warming trends; the warm-bin coefficients are essentially unchanged after the correction. Moreover, the warm-temperature effects remain robust to leave-one-out exercises dropping each of the five largest US destination states or the ten most influential counties identified by Rotemberg weight diagnostics (Goldsmith-Pinkham et al., 2020).

We conclude with a back-of-the-envelope exercise based on standard assumptions from the climate economics literature.[3] By combining our reduced-form estimates with climate projections from the CMCC-ESM2 Earth System Model under SSP2-4.5 and SSP5-8.5 scenarios, we provide benchmark quantifications of how climate

[1] While other mechanisms may also play a role, this is the mechanism we can test directly and one that is likely to be important. For migrants, labor earnings represent the primary source of income (American Community Survey), implying that their ability to remit is closely tied to their labor earnings.

[2] This is in line with the fact that Mexican migrants tend to work disproportionately in agriculture and construction (Osuna-Gomez and Medina-Cortina, 2025) -industries that are very exposed to extreme temperatures.

[3] Similar projection exercises combine reduced-form temperature responses with climate model projections (e.g., see (Deschênes and Greenstone, 2011; Burke et al., 2015a; Carleton et al., 2022)). As in these studies, our calculations assume that the estimated relationships remain stable and abstract from longer-run adjustments such as changes in population, migration responses, and other forms of adaptation.

change may affect remittances by mid-century (2050) and end-of-century (2100). Using the delta method, we shift observed daily temperature distributions by the projected warming and recount temperature-bin days for both Mexican municipalities and U.S. counties. In migrants' origin communities, warming produces modest increases in remittances (+1.9% by 2100 under SSP5-8.5) as Mexico's temperature distribution shifts into the extreme heat range. In destination communities, warming in U.S. migrant-receiving counties reduces migrants' capacity to send remittances by −5.0% under SSP5-8.5 by 2100, equivalent to $3.2 billion per year at the 2024 baseline. The negative effects associated with destination communities exceed the increases observed in at origin by a factor of two to three, suggesting that climate change may erode both the level and the insurance value of remittances under current behavioral responses.

We contribute to several strands of the literature. We add to the growing literature that studies migrant networks. Previous work shows that economic cycles (e.g., Monras, 2020; Caballero et al., 2023; Conover et al., 2023), immigration policies (e.g., Bucheli and Fontenla, 2025; Pearson, 2023; Osuna-Gomez and Medina-Cortina, 2025; Ambrosius et al., 2025), changes in migration opportunities (Ares de Parga-Regalado and Prato, 2026), and crime (e.g., Zaiour, 2023) can affect individuals at opposite ends of a migrant corridor. Our paper is the first to study the transnational effects of extreme temperatures. Within this literature, our paper is closely related to Smith (2025), who studies how droughts on both sides of the migrant network affect remittances sent to Mexico. Our analysis differs in several ways. First, droughts affect local economies primarily by generating shocks in the agricultural sector and pushing workers into other industries (Afridi et al., 2022), while extreme temperatures can affect productivity across many sectors (Aguilar-Gomez et al., 2014). Second, while Smith (2025) does not find an effect when the shock occurs in destination communities, we do. Third, we show that both extreme heat and extreme cold reduce remittances, while only high temperatures are strongly correlated with drought conditions (Leeper et al., 2025).

We further contribute to the literature studying the causal effects of temperature, particularly on migrants and migrants communities. Previous work shows that temperature shocks affect migration and remittances in the communities of origin (e.g., Cattaneo and Peri, 2016b; Randazzo et al., 2023; Zhu et al., 2024). The only previous paper to study the effect of temperature on destination communities is Joseph et al. (2018), which uses individual-level data on migrants in the United Arab Emirates (UAE) to document how remitting behavior responds to income shocks that may be driven by extreme heat. Our paper differs from all previous work as we can study the effect of temperature (cold and hot extremes) at both ends of a single migration corridor, allowing us to compare the elasticity of remittances with respect to temperature shocks in both locations.[4] We find that temperature shocks at destinations have a larger remittance response. This distinction is important when considering that global warming is not spatially uniform.[5]

Moreover, we provide new evidence on a mechanism driving the temperature-remittance relationship in destination communities: this is the first paper to study the effect of temperature shocks on migrant labor in

[4]Relative to Joseph et al. (2018), our paper also differs in the external validity of our data. They use data from a single firm covering roughly 10–15 percent of all migrant workers in the UAE, while we study responses for the full sample of the entire US–Mexico migration corridor—the largest migration corridor in the world. Furthermore, to Joseph et al. (2018) is interested in the individual migrant decision to remit as their income changes (i.e. the passthrough rate) due to different income shifters (like temperature shocks), but we study the total effect of temperature on remittance flows in the full network.

[5]National Oceanic and Atmospheric Administration. Kennedy, Caitlyn. 2020. Does "global warming" mean it's warming everywhere?

the United States. Previous work shows that high temperatures affect industries where migrant employment is common, such as agriculture (Malhi et al., 2021) and construction (Tymvios et al., 2016). We complement this literature by directly studying the labor market responses of migrants themselves. Because Mexican migrants are disproportionately concentrated in outdoor occupations (Osuna-Gomez and Medina-Cortina, 2025), they represent a particularly vulnerable group to climate shocks.

More broadly, we contribute to the literature studying how shocks propagate across migrant networks. Previous work has shown that remittances respond to adverse conditions in origin countries, such as rainfall shocks (Yang and Choi, 2007), and to economic fluctuations in destination economies (Frankel, 2011). Past work also has explored the effect of other negative shocks at destination (Frankel, 2011; Monterroso and Vilán, 2025). We extend this literature by studying temperature shocks along the largest migration corridor in the world. In doing so, we show that extreme temperatures affect remittance flows in a manner similar to other negative shocks, highlighting the role of climate conditions in shaping transnational economic linkages.

Finally, our climate projection exercise contributes to the nascent literature on the consequences of warming for international financial flows. By combining estimated dose-response curves with projections from the CMCC-ESM2 Earth System Model, we quantify how future warming will differentially affect remittances through the origin and destination channels, highlighting a previously unquantified pathway through which climate change may erode household resilience in developing countries (Cattaneo and Peri, 2016b; Carleton et al., 2025).

2 Data

2.1 Data Sources

Remittances: We use Banco de Mexico's quarterly reports on remittances sent from the United States to each Mexican municipality. These reports, available since 2013, record the monetary value of electronic remittances arriving at each municipality each quarter, sourced from credit institutions and legal entities that offer international money transfer services (Dinarte et al., 2021). Although the data exclude remittances transported physically (by foot, car, or plane), Banco de Mexico estimates, through the National Survey of International Travelers, that electronic transfers account for 97-99 percent of the total value of remittances sent to Mexico. The data record remittances arriving at each municipality but do not include the US county of origin. Our analysis covers the period from 2013 to 2024, yielding a panel of 2,453 municipalities observed across 47 quarters.[6]

Migrant networks. We use the Mexican *Matrículas Consulares de Alta Seguridad* (MCAS) to map the bilateral migrant network linking Mexican municipalities to US counties. The MCAS are identification cards issued by Mexican consulates in the United States, which allow migrants to verify their identity and confirm their place of residence abroad (Massey et al., 2010). Because the MCAS serves practical

[6]The exact municipality count varies slightly across analyses depending on data availability: 2,453 municipalities appear in the remittance data; of these, 2,448 have at least one MCAS registration and enter the destination-side estimation; 2,433 have non-zero migrant stocks in the 2012 MCAS cross-section (Panel D of Table 1). Robustness exercises that require additional data (e.g., census variables for the balance test) or that restrict the sample geographically reduce the count further. We report the specific sample size in each table and figure.

purposes—obtaining a driver's license, opening a bank account, and applying for jobs—approximately 75–80 percent of the undocumented Mexican population in the US holds one (Massey et al., 2010).[7] Each MCAS record contains the migrant's municipality of origin in Mexico and their county of residence in the United States. The MCAS does not directly reveal where migrants send remittances; however, migrant surveys indicate that parents—who typically reside in the migrant's community of origin—are the most common recipients (Cervantes, 2015). Similar to previous studies investigating the transnational effects of US shocks on Mexican outcomes (Pearson, 2023; Osuna-Gomez and Medina-Cortina, 2025; Ambrosius et al., 2025), we use the 2012 MCAS stock to construct bilateral migrant shares $s_{mc} = MCAS_{mc}/MCAS_m$, measuring the fraction of municipality m's registered migrants residing in US county c.

Weather: We source hourly weather data for all US counties and Mexican municipalities from the ERA5 reanalysis product by the European Center for Medium-Range Weather Forecasts (ECMWF). ERA5 provides global hourly estimates of air temperature, precipitation, dew point temperature, surface pressure, relative humidity, and wind speed at a resolution of approximately 31 km. We compute population-weighted averages of these variables for each US county and Mexican municipality using the Gridded Population of the World raster (GPWv4), which provides spatially resolved population estimates at approximately 1 km resolution. We aggregate the hourly data to daily means and then to quarterly measures: quarterly mean temperature, total quarterly precipitation, and quarterly averages of relative humidity, surface pressure, and wind speed.

2.2 Descriptive Statistics

Table 1 reports summary statistics for the main variables. Three features of the data are worth highlighting, as they shape the empirical strategy and the interpretation of the results.

Remittances are sparse and skewed. Panel A shows that 23% of municipality-quarter observations record zero remittances, and the distribution of positive values is heavily right-skewed: the mean per-capita remittance is $80 USD, but the median is just $17. The P10 is zero, while the P90 reaches $235. This pattern reflects the wide variation in our outcome variable and migrant network sizes across municipalities. These features motivate our use of Poisson pseudo-maximum likelihood (PPML) estimation with a population offset in the empirical section, which naturally accommodates zeros and does not require a log transformation of the dependent variable.

Migrants span the full range of US climates. Figure 1 shows that Mexican migrants are present in 2,897 out of approximately 3,100 US counties, with the heaviest concentrations along the US–Mexico border, in the Los Angeles and Chicago metropolitan areas, and in other urban centers, as well as substantial populations across the Southeast, Midwest, and Pacific Northwest. Panel D of Table 1 confirms that migrant networks are geographically dispersed: the median municipality has 521 registered migrants spread across 100 US counties, with a Herfindahl index of migrant shares of just 0.071—corresponding to an effective number of approximately 14 destination counties. This dispersal ensures that no single county's weather realization dominates a municipality's exposure.

High-migration municipalities differ systematically at baseline. Municipalities with more MCAS

[7] This figure is calculated as the ratio of MCAS holders to the estimated undocumented population using the "residual method"; see Van Hook et al. (2021). As documented migrants have alternative forms of identification, MCAS registrations primarily reflect undocumented migrants.

Table 1: Descriptive statistics

	N	Mean	SD	P10	Median	P90
Panel A: Remittances (municipality × quarter)						
Quarterly remittances (millions USD)	115,056	4.04	11.34	0.00	0.25	10.98
Remittances per capita (USD)	115,056	79.93	173.03	0.00	17.23	234.95
Population (thousands)	115,056	49.8	143.4	1.6	13.3	93.4
Share with zero remittances	115,056	0.227	—	—	—	—
Panel B: Mexican municipality weather (municipality × quarter)						
Mean temperature (°C)	117,504	20.3	4.5	14.8	19.7	26.9
Quarterly precipitation (mm)	117,504	430	411	37	310	980
Relative humidity (%)	117,504	66.5	14.1	46.3	69.0	82.7
Surface pressure (hPa)	117,504	868	83	773	845	1003
Wind speed (m/s)	117,504	0.97	0.83	0.22	0.71	2.13
Panel C: US county weather (county × quarter)						
Mean temperature (°C)	138,096	13.9	8.9	1.4	14.7	24.9
Quarterly precipitation (mm)	138,096	358	179	125	353	587
Panel D: Migrant network (MCAS 2012, cross-sectional)						
Registered migrants (MCAS)	2,433	1,717	3,865	42	521	4,288
Destination US counties	2,433	152	160	15	100	349
HHI of migrant shares	2,433	0.122	0.146	0.027	0.071	0.271

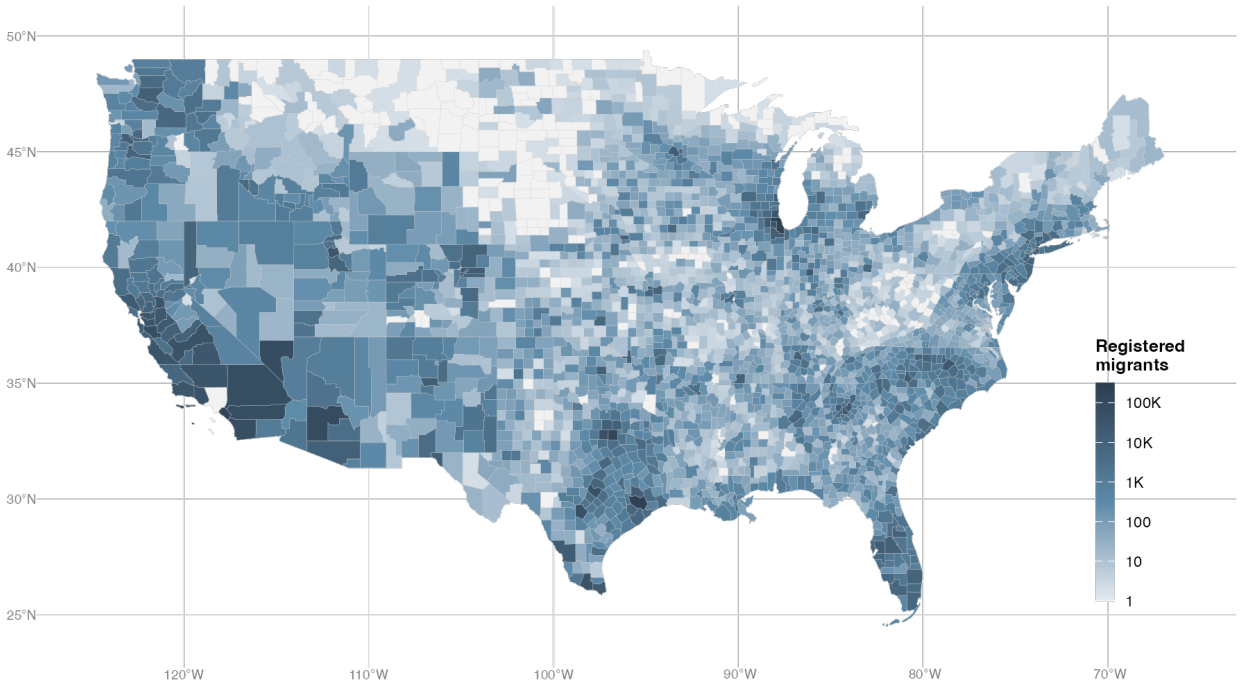
Notes: Panels A–B report statistics at the municipality-quarter level for the 2,453 municipalities in the remittance data over the 2013–2024 period. Panel C reports statistics at the county-quarter level for the 2,877 US counties with at least one MCAS registration. Panel D reports cross-sectional municipality-level statistics from the 2012 MCAS consular registration stock. Remittances are from Banco de Mexico quarterly reports. Weather variables are population-weighted averages from ERA5 reanalysis; precipitation is the quarterly total.

registrations per capita tend to be smaller, poorer, more female, more rural, and less educated than the national average (Appendix Figure A.3). These patterns are consistent with the selective emigration of young working-age males, which reshapes the demographic and economic profile of sending communities. A joint regression of MCAS per capita on eight baseline characteristics from the 2010 Census yields an R^2 of 0.32. We account for these cross-sectional differences through our fixed-effect structure, as discussed in the next section.

Migrants face different—and more variable—temperatures than those in their origin communities. Panels B and C reveal a clear asymmetry. Mexican municipalities average 20.3 °C, with moderate variation (SD = 4.5 °C). The US counties hosting migrants are, on average, 6 °C cooler (13.9 °C) and exhibit nearly twice the temperature variation (SD = 8.9 °C). The P10-P90 range for US counties spans from 1.4 °C to 24.9 °C—meaning that some migrant destinations experience near-freezing quarterly averages, while others are subtropical. This wide temperature distribution, combined with the geographic breadth of migrant networks, ensures that there is substantial variation in the temperature shocks experienced by migrants, even after accounting for municipality-specific seasonal patterns and annual trends.

3 Results

Figure 1: Geographic distribution of Mexican migrants across US counties



Notes: This figure shows the number of Mexican migrants registered in each US county according to the 2012 Mexican Consular Registration (MCAS), on a logarithmic scale. The heaviest concentrations are along the US–Mexico border (Texas, southern California, Arizona), the greater Los Angeles and Chicago metropolitan areas, and other urban centers. Migrants are present in nearly all US counties—2,897 out of approximately 3,100—confirming the geographically dispersed structure of Mexican migrant networks. Counties are excluded for Alaska and Hawaii.

3.1 Local Temperature Shocks on Remittances

We leverage plausibly exogenous temporal and spatial variation in quarterly temperature to study how local weather conditions affect remittances received by Mexican municipalities. Specifically, we estimate the following Poisson pseudo-maximum likelihood (PPML) model:[8]

$$Rem_{mqy} = \exp\left[f(Temp^{MX})_{mqy} + \delta_{m\times q} + \delta_{m\times y} + \Delta' X_{mqy} + \ln(Pop_{mqy})\right], \tag{1}$$

where Rem_{mqy} denotes the total remittances received by municipality m in quarter q of year y. The term $\ln(Pop_{mqy})$ enters as an offset—i.e., its coefficient is constrained to one—so that the model is equivalent to specifying remittances per capita as the dependent variable. All remaining coefficients are therefore interpretable as approximate percentage changes in per-capita remittances. Municipality-by-quarter fixed effects ($\delta_{m\times q}$) absorb municipality-specific seasonality, and municipality-by-year fixed effects ($\delta_{m\times y}$) absorb municipality-specific annual shocks. We additionally include time-varying controls X_{mqy}: precipitation (quadratic), relative humidity, atmospheric pressure, and wind speed.

[8]Since our dependent variable—total remittances—is a non-negative, count-like measure with a large mass at zero (22 percent of observations), we estimate our baseline specifications using PPML with a population offset.

Identification arises from deviations from typical within-municipality quarterly temperature patterns. Because the fixed effects absorb municipality-specific seasonal cycles and annual trends, the remaining variation consists of short-run temperature surprises—e.g., an unusually warm first quarter in a given municipality relative to its historical first quarter average. Such deviations are largely unanticipated (Deschênes and Greenstone, 2011), making it plausible that they are orthogonal to other short-run determinants of remittances. A remaining concern is that temperature shocks may induce internal migration within Mexico, affecting remittances through compositional changes. However, any such response would bias our estimates toward zero. We examine this directly and find no statistically significant effects of outlier temperatures on internal migration among households with migrant members (Appendix **??**). We further show that our results are robust to alternative standard error clustering strategies (state-level, Conley spatial HAC at various distance thresholds), different weighting schemes (population vs. migrant network weights), and the inclusion of counterfactual temperature controls that address potential binning bias from long-run warming trends (Jones et al., 2026) (Appendix Figure C.10).

The function $f(Temp^{MX})_{mqy}$ captures the potentially non-linear effect of temperature. Motivated by evidence of non-linear temperature effects on economic outcomes (e.g., Burke et al., 2015a; Hsiang et al., 2017; Somanathan et al., 2021; Cohen and Dechezleprêtre, 2022; Wilson et al., 2024), we estimate two specifications. First, we use a *parametric* model with a fourth-degree polynomial:

$$f(Temp^{MX})_{mqy} = \sum_{k=1}^{4} \beta_k \, (Temp^{MX}_{mqy})^k, \tag{2}$$

which allows the exposure-response curve to capture asymmetric effects at the tails of the temperature distribution without imposing the symmetry of lower-order polynomials (see, e.g., Carleton et al., 2022).[9] While the individual polynomial coefficients β_k do not have a standalone interpretation, the exposure-response function $f(\cdot)$ evaluated at any temperature T provides the approximate log-point change in per-capita remittances relative to a baseline. Second, we estimate a flexible *non-parametric* specification following Deschênes and Greenstone (2011), Barreca et al. (2016), and Cohen and Dechezleprêtre (2022):

$$f(Temp^{MX})_{mqy} = \sum_{b=0}^{B} \lambda_b \, D_{bmqy}, \tag{3}$$

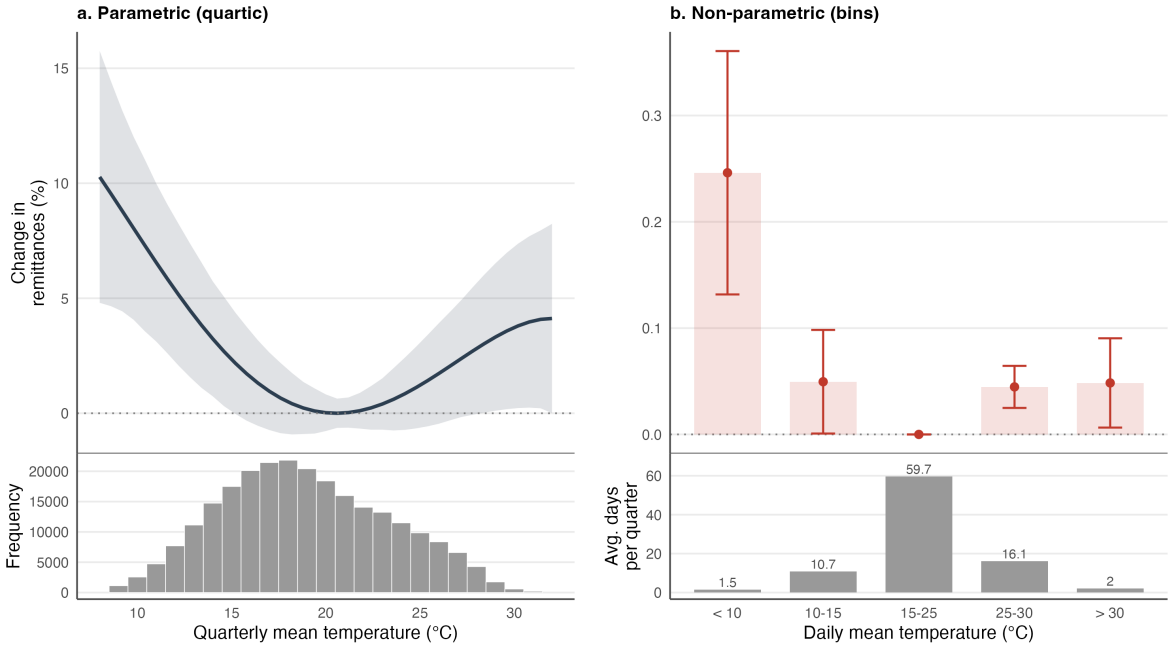
where D_{bmqy} is the number of days in a quarter-year qy for which the municipality-level daily mean temperature falls within bin b. Each coefficient λ_b represents the approximate percentage change in per-capita remittances for each additional day in bin b relative to the omitted 15–25 °C reference category. This approach avoids imposing a specific functional form and can capture threshold effects that a polynomial may miss (Deschênes and Greenstone, 2011; Barreca et al., 2016). It also preserves information from the tails of the daily temperature distribution that is lost when aggregating to quarterly means (Sarmiento et al., 2025).

Figure 2 presents the results of both specifications. Panel (a) shows the quartic exposure-response curve as a function of quarterly mean temperature, recentered at the temperature that minimizes remittances ($Temp^* = 20.4$ °C). Panel (b) displays the non-parametric bin estimates—the effect of one additional day

[9] We also estimate quadratic, cubic, and quintic specifications, showing that the results are robust across polynomial degrees.

in each temperature bin relative to the 15–25 °C reference category. In both panels, shaded areas and error-bars denote 90% confidence intervals. The bottom row shows the relevant temperature distribution for each specification. Coefficient estimates are reported in Table B.1, column (1), and in Table B.2, Panel A.

Figure 2: The effect of local temperature shocks on quarterly remittances



Notes: Panel (a) presents the quartic exposure-response curve from a PPML model of remittances on a fourth-degree polynomial of average quarterly temperatures with a population offset, recentered at its minimum; the histogram below shows the distribution of quarterly mean temperatures across municipality-quarter observations. Panel (b) presents non-parametric bin estimates: the effect of one additional day per quarter in each temperature interval relative to the 15–25 °C reference bin; the bar chart below shows the average number of days per quarter in each bin. Shaded areas denote 90% confidence intervals. Both specifications include municipality-by-year and municipality-by-quarter fixed effects, plus environmental controls (precipitation, relative humidity, atmospheric pressure, wind speed). Standard errors are clustered at the municipality level.

Both specifications confirm a U-shaped relationship between temperature and remittances. Temperatures below and above the optimum are associated with higher remittance inflows. The two panels differ in scale because the quartic measure captures the total effect of shifting the quarterly mean, whereas the bin model captures the marginal effect of a single day. Nonetheless, both approaches yield the same qualitative pattern and are statistically significant at conventional levels across the temperature distribution.

The quartic exposure-response curve implies a temperature that minimizes remittances of $Temp^* = 20.4$ °C, consistent with the typical thermal comfort range reported in the epidemiological literature (see Ormandy and Ezratty, 2012). As we identify the effect from temperature variation within the same quarter and municipality, we can interpret magnitudes using the distribution of such within-cell variation. Appendix Figure A.1 presents the density of demeaned quarterly temperatures by municipality and quarter—that is, how average quarterly temperatures vary within the same municipality during the same quarter across different years (e.g., variation across all Q1s in Tijuana). For warm temperatures, we estimate increases in remittances of 0.03, 0.12, and 0.51 percent for deviations at the 75th, 90th, and 99th percentiles of the within-cell distribution, which are equivalent to shifts of 0.55 °C, 1.11 °C, and 2.34 °C. For cold temperatures, the corresponding increases are 0.04, 0.15, and 0.66 percent for deviations of −0.55 °C, −1.16 °C, and −2.43 °C.

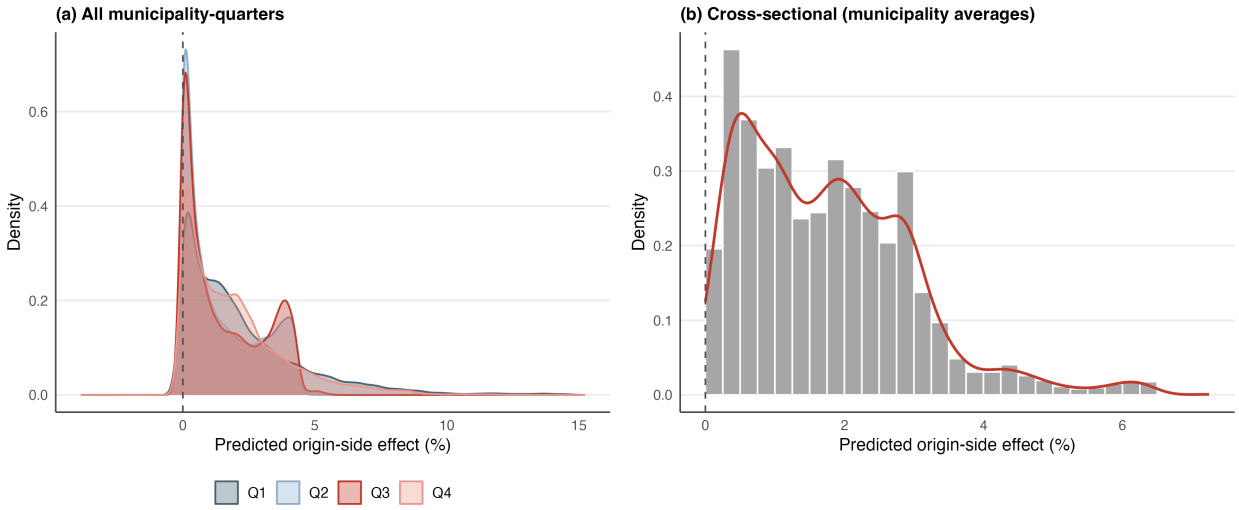
The quartic specification reveals an asymmetry: cold deviations from $Temp^*$ produce larger effects on remittances than equally-sized warm deviations. This is consistent with previous studies examining the relationship between temperature variations and mortality in Mexico (Cohen and Dechezleprêtre, 2022; Bressler et al., 2026).

The non-parametric specification corroborates the U-shaped pattern. One additional day per quarter with average temperatures below 10 °C increases remittances by 0.246 percent, while the effect for the milder cold interval (10–15 °C) is 0.050 percent. For warm temperatures, one additional day in the 25–30 °C and 30–42 °C intervals increases remittances by 0.045 and 0.048 percent, respectively. The extreme heat estimate is conservative: secular warming has predictably increased the number of hot days in Mexican municipalities, attenuating the non-parametric coefficient through the binning bias channel identified by Jones et al. (2026). After absorbing the predictable component of temperature exposure (Section 3.3), the 30–42 °C coefficient nearly doubles to 0.093 percent ($p < 0.01$). As the bottom-right panel of Figure 2 shows, the extreme temperature bins (<10 °C and >30 °C) contain, on average, fewer than 2 days per quarter, whereas the reference category (15–25 °C) accounts for roughly 60 days (see also Appendix Figure A.2). We can estimate the magnitude of these shocks using the within-cell standard deviation of temperature bin counts. For example, the within-cell standard deviation of days below 10 °C is 0.83, implying that a one-standard-deviation increase in extreme cold days raises remittances by approximately 0.20 percent.

Figure 3 translates these patterns into predicted municipality-level effects by evaluating the quartic dose-response at each municipality's observed quarterly mean temperature and normalizing to zero at the interior minimum ($T^* = 20.4$ °C). The U-shape of the origin curve implies that any departure from the thermal comfort zone increases remittance inflows, consistent with an insurance motive in which households facing adverse weather conditions receive larger transfers from abroad (see also Randazzo et al., 2023, who show that Mexican households use remittances to purchase air conditioning as a thermal adaptation strategy). Appendix **??** provides direct evidence for this channel: using daily person-level data from Mexico's National Survey of Occupation and Employment (ENOE) matched to same-day temperatures, we show that heat significantly reduces individual labor supply, with cumulative exposure to temperatures above 30°C over the current and preceding seven days reducing the probability of working by 1.2 percentage points, with effects driven by formal workers (−1.4 pp) and concentrated in high-migration municipalities (−1.3 pp). Panel (a) shows all municipality-quarter observations colored by calendar quarter. Because Mexico's temperature distribution is asymmetric around the optimum—most of the country cools well below 20 °C in winter but rarely exceeds 30 °C in summer—the seasonal pattern is driven primarily by the cold flank: Q1 (winter, average 18.1 °C) produces the largest predicted increases (median 1.5%), followed by Q4 (fall, 18.7 °C, median 1.5%), while Q2 (spring, 22.5 °C) and Q3 (summer, 21.9 °C), when municipalities are near the optimum, show minimal effects (median <1%). Panel (b) shows the cross-sectional distribution of time-averaged predicted effects: the median municipality experiences a predicted origin-side increase of 1.6%, with the 90th percentile at 3.2% and the most exposed municipalities—those in cold highland regions—reaching over 5%.

The modest magnitudes of the origin-side effects—relative to the destination-side losses documented below—reflect the fact that Mexican quarterly mean temperatures are generally close to the thermal comfort zone. The interquartile range of municipal temperatures spans roughly 17–23 °C, placing most observations on the flat bottom of the U-shape where the curve is nearly level. In contrast, the destination side draws on US county temperatures that range from well below freezing to above 30 °C, placing many municipality-

Figure 3: Distribution of predicted origin-side temperature effects



Notes: This figure shows the distribution of predicted origin-side effects on remittances, computed by evaluating the quartic dose-response curve at each municipality's observed quarterly mean temperature and normalizing to zero at the interior minimum ($T^* = 20.4$ °C). Panel (a) shows all municipality-quarter observations, colored by calendar quarter: Q1 (dark blue) and Q4 (salmon) sit furthest from the optimum, producing the largest predicted increases; Q2 (light blue) and Q3 (red) cluster near the minimum where effects are small. Panel (b) shows the cross-sectional distribution of time-averaged predicted effects across 2,471 municipalities. The median municipality experiences a predicted origin-side increase of approximately 1.6% relative to the temperature optimum.

quarter observations on the steep flanks of the inverted-U.

3.2 Transnational Temperature Shocks and Their Impact on Remittances

We now turn to the destination side: how temperature shocks experienced by Mexican migrants in the United States affect the remittances sent to their municipalities of origin.[10] Estimating this relationship requires linking migrant destinations to origin communities. We use the 2012 Mexican Consular Registration (MCAS) stock to construct pre-determined migrant shares $s_{mc} = MCAS_{mc}/MCAS_m$, which measure the fraction of registered migrants from municipality m residing in US county c. Shares are constructed from a single pre-sample cross-section (2012) to avoid endogeneity between contemporaneous migration decisions and temperature realizations. Combined with county-level weather data, these shares allow us to construct municipality-specific exposure to US temperature shocks in a shift-share (Bartik) framework (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022).

The core identification challenge in shift-share designs is separating idiosyncratic shock variation from compositional differences across units (Borusyak et al., 2022; Adao et al., 2019). In our setting, municipalities whose migrants disproportionately reside in systematically warmer (or colder) US counties may differ in baseline remittance behavior. A standard approach controls for the network mean of shocks—the unweighted average of county-level weather across a municipality's connections—to absorb aggregate climate realizations (e.g.,

[10] Appendix Figure A.1 presents the quarterly temperature distribution decompositions for both Mexican municipalities and US counties.

Cattaneo and Peri, 2016a). However, including network means as regressors alongside the share-weighted exposure indices can induce severe multicollinearity, as the two are mechanically correlated through overlapping county shocks. We address this by recentering the county-level shocks before collapsing them to the municipality level, following the logic of Borusyak et al. (2022). In the first step, we work with the expanded municipality-county-quarter panel and project each county-level shock variable onto municipality-by-time fixed effects:

$$g_{ct} = \alpha_{mt} + \tilde{g}_{ct}, \tag{4}$$

where g_{ct} is a county-level temperature variable for US county c in quarter t, and α_{mt} denotes municipality-by-time fixed effects estimated via OLS. Crucially, the same recentering procedure applies to any function of temperature: in the non-parametric specification, $g_{ct} = D_{ct}^{(b)}$ (the number of days in temperature bin b); in the parametric specification, $g_{ct} = (T_{ct}^{US})^k$ for each polynomial power $k = 1, \ldots, 4$. The residual $\tilde{g}_{ct}$ captures how county c's temperature exposure in quarter t deviates from the average exposure of municipality m's entire migrant network at that time. These fixed effects flexibly absorb the network mean—the aggregate climate realization common to all counties connected to a given municipality in a given period—without requiring its explicit construction as a regressor. Borusyak et al. (2022) shows that shift-share identification relies on the quasi-random assignment of shocks conditional on shares, and that recentering the instrument—removing its expected value given the share structure—ensures that the remaining variation identifies a causal effect. Our municipality-by-time fixed effects accomplish precisely this recentering: they subtract from each county's weather the network-weighted mean so that $\tilde{g}_{ct}$ captures only idiosyncratic deviations.

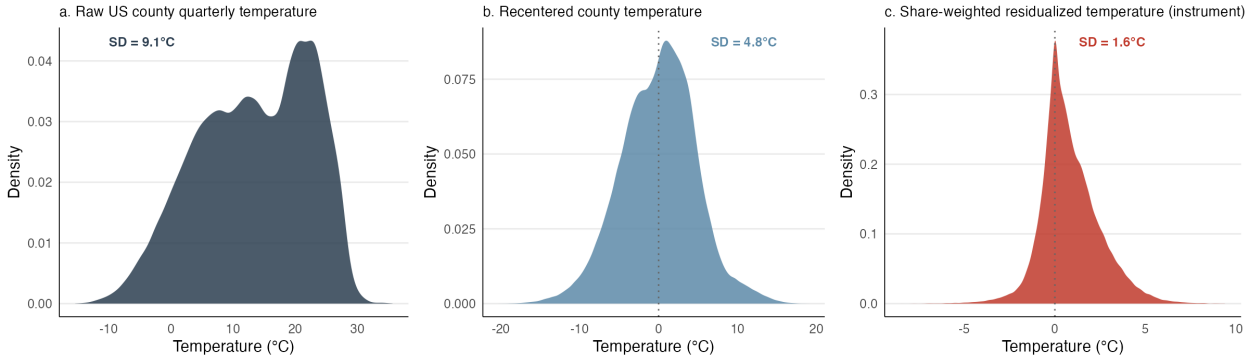
In the second step, we aggregate the residualized shocks to the municipality-quarter level using migrant shares:

$$\tilde{Z}_{mt} = \sum_{c} s_{mc}\, \tilde{g}_{ct}. \tag{5}$$

The exposure index $\tilde{Z}_{mt}$ is a shift-share instrument in which the "shifts" are residualized county-level temperature shocks and the "shares" are pre-determined migrant network weights. By construction, $\tilde{Z}$ is purged of aggregate climate variation common to a municipality's network, isolating the component driven by idiosyncratic, county-specific weather surprises transmitted through the migrant network. Figure 4 illustrates the three stages of the instrument construction. Raw US county temperatures (panel a, SD = 9.1 °C) are first recentered by absorbing municipality-by-time fixed effects (panel b, SD = 4.8 °C), and then collapsed with migrant shares (panel c, SD = 1.6 °C). The progressive narrowing of the distribution shows that the identifying variation is a small, idiosyncratic component of county-level weather—precisely the quasi-random variation that the shift-share design exploits.

As for the origin side, we estimate both a *parametric* and a *non-parametric* specification. In the parametric version, we construct shift-share instruments for each polynomial power of county-level quarterly mean temperature, recentering $(T_{ct}^{US})^k$ for $k = 1, \ldots, 4$ separately and collapsing each with shares to obtain $\tilde{Z}_{mt}^{(k)}$. Note that each polynomial power is recentered independently, so that $\tilde{Z}_{mt}^{(2)} \neq (\tilde{Z}_{mt}^{(1)})^2$; the second-order instrument captures the shift-share average of squared temperature deviations, preserving within-network distributional information that would be lost by squaring the network mean. In the non-parametric version,

Figure 4: Construction of the shift-share temperature instrument



Notes: This figure illustrates the three stages of the shift-share instrument construction. Panel (a) shows the distribution of raw quarterly mean temperatures across US counties hosting Mexican migrants (SD = 9.1 °C). Panel (b) shows the recentered county temperatures $\tilde{g}_{ct} = T_{ct} - \hat{T}_{ct}$, obtained after absorbing municipality-by-year-quarter fixed effects (Eq. 4); the standard deviation drops to 4.8 °C as predictable seasonal and network-level trends are removed. Panel (c) shows the final instrument $\tilde{Z}_{mt} = \sum_c s_{mc}\tilde{g}_{ct}$, obtained by collapsing the recentered county temperatures with pre-determined 2012 MCAS migrant shares (Eq. 5). The standard deviation drops further to 1.6 °C, reflecting the narrow band of idiosyncratic, within-network temperature variation that identifies the destination-side effects.

we recenter each temperature bin's day-count $D_{ct}^{(b)}$ and collapse to obtain $\tilde{Z}_{mt}^{(b)}$. We then estimate the following collapsed PPML specification:

$$Rem_{mt} = \exp\left[f(\tilde{Z}_{mt}) + \Gamma' X_{mt} + \delta_{m\times y} + \delta_{m\times q} + \ln(Pop_{mt})\right], \tag{6}$$

where Rem_{mt} denotes the total remittances received by municipality m in year-quarter t, $\ln(Pop_{mt})$ is the population offset, and $\delta_{m\times y}$ and $\delta_{m\times q}$ are the municipality-by-year and municipality-by-quarter fixed effects. The vector X_{mt} includes share-weighted US precipitation and other environmental controls, as well as Mexican temperature bins and weather controls to account for spatially correlated weather between origin and destination.[11] In the parametric specification, $f(\tilde{Z}_{mt}) = \sum_{k=1}^{4} \beta_k \tilde{Z}_{mt}^{(k)}$, which traces out the quartic exposure-response $f(T) = \sum_k \beta_k T^k$ evaluated at the county level and aggregated through the shift-share structure. In the non-parametric specification, $f(\tilde{Z}_{mt}) = \sum_{b\in\mathcal{B}} \beta_b \tilde{Z}_{mt}^{(b)}$, where we define six US temperature bins—$(-45, 0]$, $(0, 5]$, $(5, 10]$, $(20, 25]$, $(25, 30]$, and $(30, 45]$ °C—with the 10–20 °C interval as the omitted reference category. Each bin coefficient β_b represents the approximate percentage change in per-capita remittances from one additional migrant-share-weighted day in US temperature bin b, relative to a day in the reference range.

The balance test in Appendix Figure A.3 shows that municipalities with larger migrant stocks differ sys-

[11] The weather control structures differ across the two sides. On the origin side, we use quintile bins of relative humidity, atmospheric pressure, and wind speed, which flexibly capture non-linear local weather effects. On the destination side, these variables enter in levels because they must pass through the same shift-share aggregation as the temperature instrument: share-weighted county-level values and unweighted network means, following the recentering approach of Borusyak et al. (2022). Defining within-county quintiles and then aggregating via migrant shares would obscure the ordinal ranking. The temperature bins also differ: the origin side uses four Mexican bins with a 15–25 °C reference, while the destination side uses six US bins with a 10–20 °C reference. These choices reflect the different climate distributions—Mexican municipalities average 20.3 °C with moderate variation, whereas US destination counties average 13.9 °C with nearly twice the variance (Table 1), requiring finer resolution in the tails.

tematically at baseline (smaller, poorer, more rural); because Goldsmith-Pinkham et al. (2020) note that share-level balance is the key requirement for shift-share identification, the municipality-by-time fixed effects that recenter the instrument also absorb these cross-sectional differences.

Inference in shift-share designs requires accounting for the fact that many municipalities share exposure to the same underlying county-level shocks through overlapping migrant networks. Standard municipality level clustering misrepresents dependence because the identifying variation originates at the shock level—county–by-quarter weather realizations—rather than at the municipality level. We, therefore, compute standard errors using the Adao et al. (2019) (AKM) variance estimator, which clusters score contributions at the shock level (county × quarter) and propagates the covariance induced by shared shocks through the share-weighting structure. This approach is valid under the assumption that county-level temperature shocks are as good as randomly assigned conditional on shares—a plausible assumption given the quasi-random nature of weather variation (Deschênes and Greenstone, 2011)—and that shares are pre-determined, which our use of 2012 MCAS data ensures. We show in the appendix that results are robust to alternative approaches for absorbing aggregate climate variation, including explicit network-mean controls and geographic-zone-by-time fixed effects (US state × quarter, Köppen climate zone × quarter) in the recentering step.

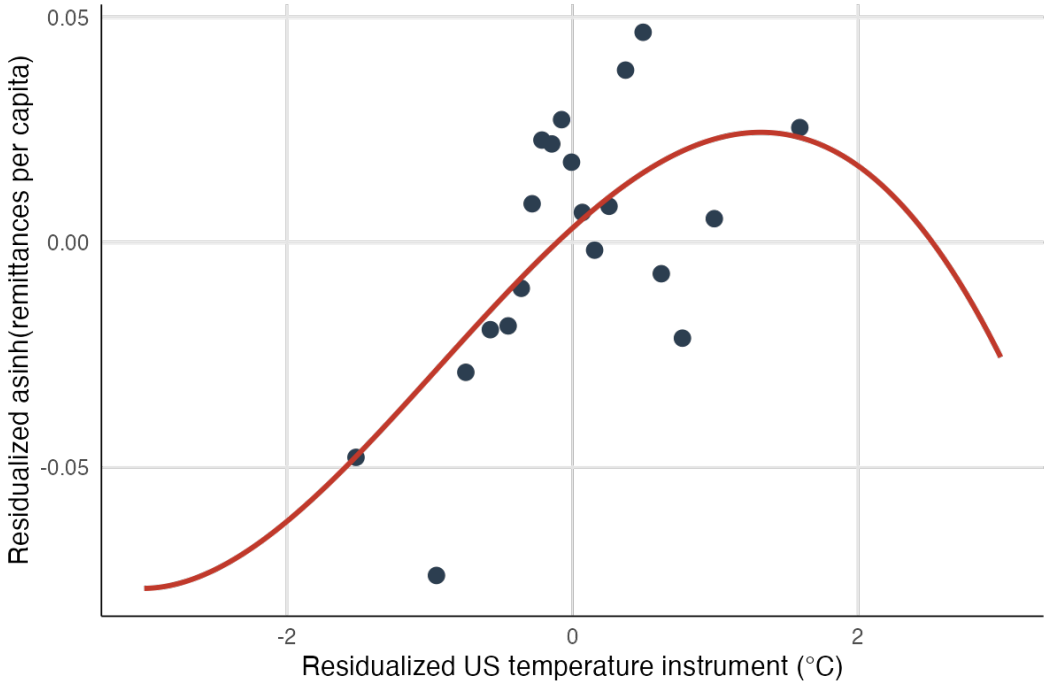
Before presenting the structural estimates, Figure 5 provides model-free visual evidence of the reduced-form relationship. The binscatter plots residualized asinh(remittances per capita) against the residualized share-weighted US temperature instrument after absorbing municipality-by-year and municipality-by-quarter fixed effects. The inverted-U pattern is clearly visible: remittances peak near the demeaned mean and decline with temperature deviations in either direction (quadratic coefficient $\beta_2 = -0.007$, $p < 0.001$).

Figure 6 presents the results of both specifications. Panel (a) shows the quartic exposure-response curve as a function of US county quarterly mean temperature, normalized to zero at the optimum ($T^* \approx 17$ °C). Panel (b) displays the non-parametric bin estimates—the effect of one additional share-weighted day in each US temperature bin relative to the 10–20 °C reference category. In both panels, shaded areas denote 90% confidence intervals. The bottom row shows the relevant temperature distribution for each specification. Coefficient estimates are reported in Table B.1, column (2) and Table B.2, Panel B.

Both specifications reveal an inverted U-shaped relationship between US network temperature and remittances. Temperatures outside the moderate range are associated with lower remittance inflows. The quartic exposure-response curve peaks at approximately 17 °C, close to the center of the non-parametric reference bin (10–20 °C). All six non-parametric bin coefficients are negative and statistically significant at the 1% level (Appendix Table B.2, Panel B), confirming that any deviation from the 10–20 °C optimum—whether cold or hot—reduces remittances.

The quartic curve can be interpreted as a hill: remittances are maximized when US county temperatures are near the peak at 17 °C and decline on both sides. Crucially, the slope of this hill varies—and it is the slope at a given temperature that determines how costly a temperature shock is. Near the peak (15–20 °C), the curve is nearly flat, so a temperature surprise of a few degrees has little effect on remittances. On the flanks—below 10 °C or above 25 °C—the curve is steep, and the same-sized shock produces a much larger decline. To illustrate, consider a uniform 3 °C warming shock in a county that hosts the entirety of a municipality's migrant network. If the county's typical quarterly temperature is 18 °C (near the peak, e.g., coastal California), the shock moves remittances by only −1.7%. If instead the county typically averages 25 °C (e.g., Houston

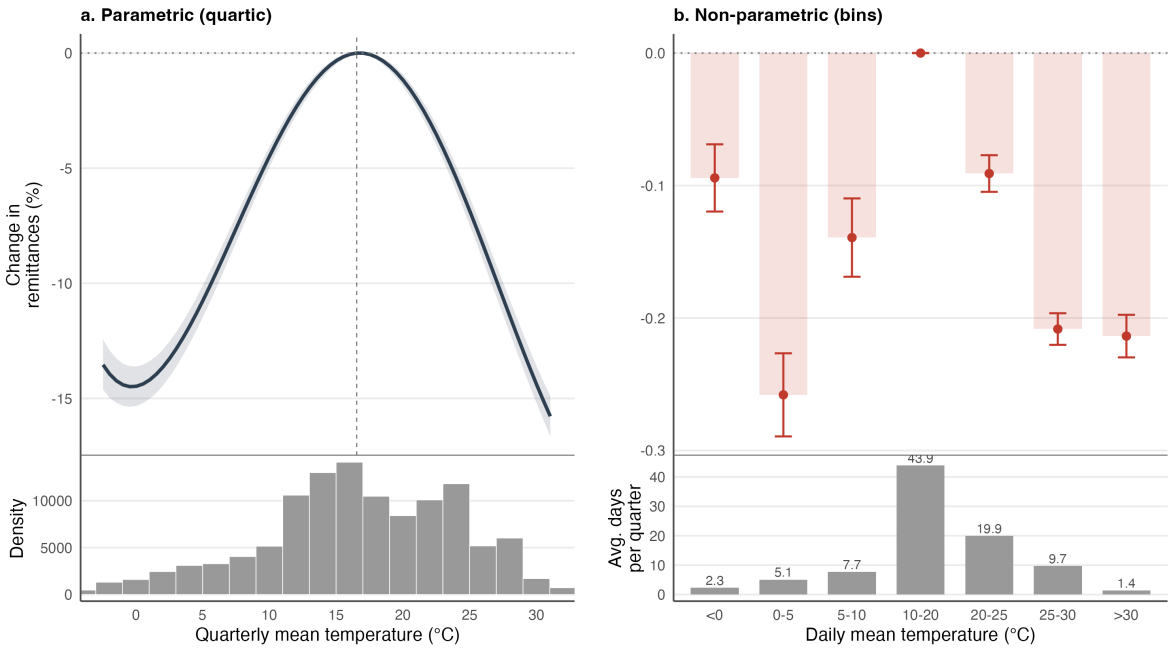
Figure 5: Reduced-form binscatter: remittances vs. temperature instrument



Notes: This figure shows a binscatter of residualized asinh(remittances per capita) against the residualized share-weighted US temperature. Both variables are residualized against municipality-by-year and municipality-by-quarter fixed effects. The asinh transformation preserves all observations including zeros (23% of the sample). Observations are grouped into 20 equal-sized bins; each point represents the within-bin mean. The red curve is a quartic polynomial fit to the underlying (unbinned) data. Sample: 114,351 municipality-quarter observations.

in spring), the same 3 °C shock reduces remittances by −4.8%—nearly three times as much—because the curve is steep and steepening at that point. On the cold side, a 3 °C cooling in a county with a baseline of 10 °C (e.g., the US northeast in early spring) reduces remittances by −4.0%. These county-level effects are transmitted to Mexican municipalities through the shift-share structure: municipalities whose migrants are concentrated in a single county bear the full effect, while those with geographically dispersed networks experience a share-weighted average. In practice, most networks are highly dispersed: the median municipality sends migrants to 100 US counties, with a Herfindahl index of migrant shares of just 0.071—corresponding to an effective number of approximately 14 destination counties (Appendix Figure A.4). The implication is that vulnerability to US temperature shocks depends not only on the magnitude of the shock but also on *where on the curve* the shock lands-municipalities whose migrants reside in already-hot or already-cold counties are disproportionately exposed because they sit on the steep flanks of the dose-response function. Moreover, this vulnerability is seasonal: Appendix Figure B.5 shows that in Q1 (winter), the typical municipality's share-weighted network temperature is around 9.6 °C—on the steep cold flank—while in Q3 (summer), it is around 24.2 °C—on the steep hot flank. In spring and fall, networks are near the flat peak, where shocks are largely innocuous. Beyond position on the curve, the *correlation structure* of temperature shocks across destinations matters for remittance volatility. Appendix Appendix G formalizes this insurance channel by constructing a share-weighted pairwise shock correlation measure ρ_m for each municipality's network. Municipalities whose migrants reside in climatically uncorrelated destinations experience approximately 40% less quarter-to-quarter variation in predicted destination-side effects, consistent with portfolio diversification

Figure 6: The effect of transnational temperature shocks on quarterly remittances



Notes: Panel (a) presents the quartic exposure-response curve from a collapsed PPML model of remittances on shift-share instruments constructed from the four powers of US county quarterly mean temperature, each recentered separately with municipality-by-time fixed effects and collapsed using pre-determined 2012 MCAS migrant shares, with a population offset. The curve is normalized to zero at its maximum; the histogram below shows the distribution of quarterly mean temperatures across US county-quarter observations. Panel (b) presents non-parametric bin estimates: the effect of one additional share-weighted day per quarter in each US temperature interval relative to the 10–20 °C reference bin; the bar chart below shows the average number of share-weighted days per quarter in each bin. Shaded areas denote 90% confidence intervals. Both specifications include municipality-by-year and municipality-by-quarter fixed effects, share-weighted US precipitation and environmental controls, and Mexican temperature and weather controls. Standard errors are clustered at the municipality level.

of weather risk. However, because the dose-response is concave near its peak, geographic diversification also imposes a Jensen's inequality cost on the *level* of predicted effects: spreading exposure across counties at different temperatures yields a worse average outcome than concentrating near the optimum.[12]

To interpret the non-parametric magnitudes, recall that each coefficient β_b measures the effect of one additional *share-weighted* day in the US temperature bin b. Because the migrant shares s_{mc} sum to one for each municipality, one share-weighted day corresponds to one extra day in bin b experienced uniformly across all US counties where a municipality's migrants reside. The magnitude of the effect for any given municipality therefore depends on both the size of the temperature shock and the geographic concentration of its migrant network. Consider a municipality whose migrants are concentrated in a single US county. If that county experiences a heat wave that shifts 10 days from the 10–20 °C reference range into the above-30 °C bin, the shift-share exposure increases by approximately 10 share-weighted days, implying a reduction in remittances

[12]Jensen's inequality states that for a concave function f, the expected value of the function is less than or equal to the function evaluated at the expected value: $E[f(X)] \leq f(E[X])$. Applied here, the concave portion of the inverted-U dose-response near $T^* \approx 17$°C means that a municipality whose migrants are spread across counties with mean temperatures of, say, 9°C and 25°C will experience a worse average predicted effect than one whose migrants are concentrated in a single county at 17°C, even though both networks have the same average temperature. The magnitude of this gap grows with the dispersion of destination temperatures.

of $10 \times 0.214\% \approx 2.1\%$. For the same shock hitting a municipality whose migrants are dispersed equally across five counties, only one-fifth of the network is exposed, yielding an effect of roughly $2 \times 0.214\% \approx 0.4\%$. On the cold side, the 0–5 °C bin carries the largest marginal effect: one additional share-weighted day in this range reduces remittances by 0.258%. A polar vortex adding 15 days below 5 °C in a county that hosts the entirety of a municipality's network would reduce remittances by $15 \times 0.258\% \approx 3.9\%$. Across all bins, the results are consistent with an economic channel in which temperature extremes reduce migrants' labor productivity and earnings—particularly in outdoor-intensive sectors such as construction and agriculture, which, in turn, lowers the remittances sent home.

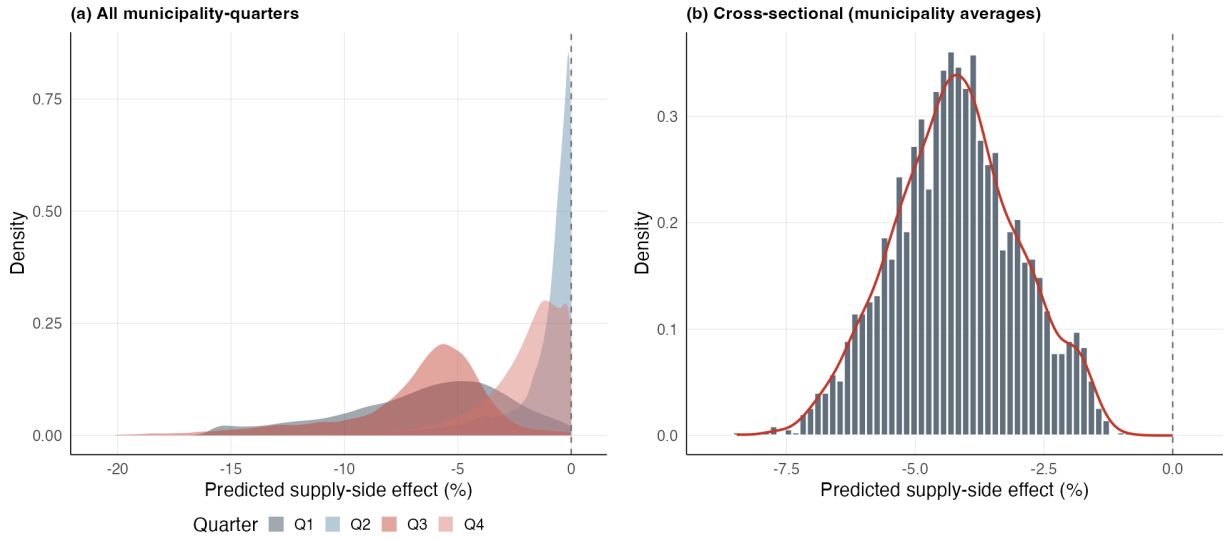
The two specifications are mutually reinforcing. The quartic captures the smooth curvature of the dose-response and locates the peak at $T^* \approx 17$ °C; the non-parametric bins confirm the same inverted-U pattern without imposing a functional form, with the 10–20 °C reference category spanning the quartic's optimum. Both specifications agree on the relative magnitudes: the steepest portions of the quartic curve—below 5 °C and above 25 °C—coincide with the largest non-parametric bin coefficients (-0.258% per share-weighted day for 0–5 °C; -0.208% and -0.214% for 25–30 °C and above 30 °C, respectively), while the bins flanking the optimum (5–10 °C and 20–25 °C) exhibit the smallest effects, consistent with the flat top of the quartic hill. The consistency across specifications—one a smooth polynomial and the other a step function—increases confidence that the inverted-U relationship is a robust feature of the data rather than an artifact of either functional form. Notably, this destination-side inverted U is the mirror image of the origin-side U-shape (Appendix Figure B.6): temperature extremes in Mexico *increase* remittance inflows (consistent with an insurance motive), while temperature extremes in the US *decrease* them (consistent with an earnings channel). The two curves jointly suggest that remittances serve as a climate risk-sharing mechanism whose effectiveness is constrained by conditions at both ends of the migration corridor.

Figure 7 translates these patterns into predicted municipality-level losses by evaluating the quartic dose-response at each municipality's observed share-weighted US temperature and normalizing to zero at the optimum ($T^* \approx 17$ °C). Panel (a) shows all municipality-quarter observations colored by calendar quarter. In Q1 and Q3, when networks sit on the steep flanks of the curve, the average predicted loss reaches -6.4% and -6.8%, respectively; in Q2 and Q4, losses are modest (-1.3% and -2.1%). Panel (b) shows the cross-sectional distribution of time-averaged effects: the median municipality experiences a predicted loss of 3.6%, while the most exposed decile loses over 9%.[13]

In dollar terms, the implied temperature penalty on remittances is economically significant. Average annual per-capita remittances in our sample are $314 USD; the mean predicted loss of 4.2% implies approximately $13 per capita per year in forgone remittances. For the median municipality (population 13,300), this trans-

[13]The approximately normal shape of the destination-side distribution contrasts with the right-skewed, multimodal origin-side distribution in Figure 3. Two structural features explain this difference. First, the Bartik construction averages each municipality's US temperature exposure across many destination counties (the median municipality's migrants reside in 100 counties), and this diversification smooths the distribution toward normality by a central-limit-theorem effect. On the origin side, each municipality has a single local temperature with no averaging, so the distribution directly reflects Mexico's distinct climatic zones (temperate highlands near the optimum, hot lowlands, cold sierra), producing the observed multimodality. Second, the inverted-U functional form means that *every* deviation from the optimum generates a loss, filling out a unimodal distribution. The U-shaped origin curve has a flat bottom near 20–25°C where effects are near zero, creating a mass point at low values, with only municipalities facing genuinely extreme temperatures contributing to the right tail.

Figure 7: Distribution of predicted destination-side temperature effects



Notes: This figure shows the distribution of predicted destination-side effects on remittances, computed by evaluating the quartic dose-response curve at each municipality's share-weighted US temperature and normalizing to zero at the optimum ($T^* \approx 17$ °C). Panel (a) shows all municipality-quarter observations, colored by calendar quarter: Q1 (dark blue) and Q3 (red) sit on the steep cold and hot flanks, respectively, generating the largest negative effects; Q2 (light blue) and Q4 (salmon) cluster near the peak where effects are minimal. Panel (b) shows the cross-sectional distribution of time-averaged predicted effects across 2,433 municipalities. The median municipality experiences a predicted destination-side loss of approximately 4% relative to the temperature optimum.

lates to roughly $40,000 per year in aggregate remittances lost due to non-optimal temperatures. At the national level, the temperature penalty amounts to approximately $1.7 billion per year, or 4.3% of total remittance inflows to Mexico over the sample period. To put this in perspective, Mexico's flagship conditional cash transfer program (*Prospera*, formerly *Progresa/Oportunidades*) reached approximately 6.4 million beneficiary households with an annual budget of roughly $4 billion USD (Parker and Todd, 2017). The estimated temperature penalty on remittances is thus equivalent to roughly 40% of the country's largest anti-poverty program. Unlike *Prospera*, however, these losses are not offset by any compensating transfer; they represent a net reduction in resources flowing to remittance-receiving households, driven entirely by climatic conditions at the migrants' places of work.

3.3 Robustness

We subject our main results to a battery of robustness checks, which are reported in Appendix C.

Specification robustness. Appendix Figure C.8 shows that the shape of the parametric dose-response curve is stable across polynomial degrees 2 through 5 on both sides of the market: the origin-side U-shape and the destination-side inverted-U are qualitatively unchanged, with the quartic and quintic specifications being nearly indistinguishable. Appendix Figure C.8 and Figure C.9 re-estimate the non-parametric bin models under alternative specifications. On the origin side, the extreme cold effect (<10 °C) is stable across all variants: substituting OLS for PPML, removing weather controls, restricting to precipitation-only controls, excluding the COVID-19 pandemic recovery year, and trimming the sample by population or consular registrations. The warm bin coefficients are more sensitive to the inclusion of environmental

controls, consistent with the correlation between temperatures in the 25–42 °C range and humidity and atmospheric conditions in Mexico (Baldwin et al., 2023). On the destination side, the baseline inverted-U pattern is robust to estimation by OLS, weather control ladders, sample restrictions, and the exclusion of 2021. The post-pandemic subsample (2022–2024) confirms that the negative temperature-remittance relationship persists after COVID, which is relevant as the COVID-19 pandemic disrupted labor markets through which temperature shocks affect migrant earnings—via lockdowns, sectoral reallocation, and federal stimulus transfers—while simultaneously triggering an unprecedented surge in remittance flows to Mexico (Dinarte et al., 2021). Appendix Table C.3 investigates how this structural break affected the destination-side estimates. We estimate an interaction model that augments the baseline specification with temperature bin $\times$ $\mathbf{1}$(2020–2021) interactions, allowing the temperature-remittance relationship to differ during the pandemic. The base effects, capturing the structural relationship outside 2020–2021, are negative and highly significant for all six temperature bins. The positive COVID interaction terms indicate that the pandemic *attenuated* the warm-temperature effect, which is consistent with disrupted outdoor employment reducing the importance of the temperature-earnings channel during 2020-2021. Subsample exercises corroborate this interpretation: excluding only 2021 preserves all negative signs, and the post-pandemic period (2022–2024) shows the negative relationship reasserting itself.

Inference. Appendix Table C.4 and Table C.5 report the same point estimates under alternative standard error strategies. On the origin side, we consider state-level clustering and Conley spatial HAC standard errors at 25, 50, and 100 km distance cutoffs to account for spatial correlation in weather shocks across neighboring municipalities. The warm bin (25–30 °C) is robust across all approaches, retaining significance at the five percent level even at 100 km. The extreme cold bin (<10 °C) survives state clustering and Conley SEs through 50 km. On the destination side, the AKM standard errors—which cluster at the shock level (county $\times$ quarter) and account for the shift-share structure—are substantially smaller than those from traditional municipality or state clustering. This pattern reflects the design: the identifying variation originates from county-level temperature shocks transmitted through pre-determined migrant shares, and the AKM estimator correctly concentrates inference on this source of variation. With municipality-level clustering, the warm bins (20–25, 25–30, and 30–45 °C) remain significant at the five percent level, though two of the moderate cold bins lose significance. State-level clustering, the most conservative approach, preserves significance at the ten percent level for all six bins.

Destination-side identification: Appendix Table C.6 reports Rotemberg weight diagnostics (Goldsmith-Pinkham et al., 2020) for the destination-side shift-share design. Panel A reveals that the top 10 counties account for 81% of positive weight for the key (30,45] °C bin, with an effective number of identifying counties of roughly 10. Panel B lists these counties: Harris (Houston), Dallas, Hidalgo (McAllen), Pima (Tucson), Bexar (San Antonio), Maricopa (Phoenix), Travis (Austin), Tarrant (Fort Worth), Fresno, and Webb (Laredo). Negative weights account for less than 8% of total weight magnitude across all bins, within the range considered acceptable for shift-share designs. This concentration of identifying variation is a feature of the design rather than a threat to validity: the counties that contribute the most weight are those that host the largest migrant populations and experience the most temperature variation, making them precisely the labor markets most relevant for aggregate remittance flows.

As a stringent test, Appendix Table C.7 simultaneously excludes all ten top-weight counties, renormalizing migrant shares and reconstructing the shift-share instrument from scratch. The non-parametric bin estimates

confirm that all three warm bins (20–25, 25–30, and 30–45 °C) remain negative and highly significant, and the results are also robust to the binning bias correction applied on top of the leave-one-out sample.

The moderate cold bins (0–5 °C and 5–10 °C), however, change sign in the leave-one-out sample, flipping from negative in the baseline to positive or near-zero. This fragility foreshadows and is explained by the binning bias correction discussed below. Secular warming in the United States has progressively and predictably reduced the number of moderate cold days (0–10 °C) in the counties where Mexican migrants concentrate. The major migrant-receiving counties in southern Texas and Arizona have experienced warming trends that shift days out of the 0–10 °C range and into the 10–20 °C reference category. The baseline negative coefficients on the moderate cold bins partly confound this secular trend with the causal effect of weather shocks. Any perturbation to the county composition alters the balance between trend-driven and shock-driven variation, destabilizing these fragile coefficients. Crucially, after the binning bias correction (Figure C.11), the moderate cold bins shrink substantially, while extreme cold (<0 °C) and all warm bins remain robustly negative. The leave-one-out exercise thus provides independent corroboration of the binning bias diagnosis: the moderate cold bins are fragile precisely because they confound secular trends, and the warm-temperature effects that drive the economic conclusions are robust to arbitrary perturbations of the instrument.[14]

Exogeneity: Following Jaeger et al. (2018), we test whether future or past temperature shocks predict current remittances by augmenting the baseline models with one-quarter leads $(t+1)$ or lags $(t-1)$. On the origin side, the contemporaneous extreme cold effect remains stable when leads or lags are included; some lead and lag terms are significant, likely reflecting seasonal autocorrelation in the temperature process.[15] On the destination side, contemporaneous effects for the warm bins are generally robust to the inclusion of temporal leads and lags.

Temporal persistence: We examine temporal dynamics using distributed lag models with four quarterly lags. On the origin side, however, this approach is poorly suited to quarterly data combined with municipality-by-season fixed effects. Because temperature is strongly seasonal, the fourth quarterly lag corresponds to the same calendar quarter one year prior, producing an autocorrelation of $r = 0.85$ between contemporaneous and L4 temperature bins. After municipality-by-quarter fixed effects absorb average seasonal patterns, the remaining identifying variation (year-to-year deviations within a given season) is too collinear across horizons to separately identify lag effects. Diagnostic regressions confirm this: cumulative parametric effects yield implausible optimal temperatures at longer horizons, a hallmark of near-collinearity. We therefore restrict the dynamic analysis to the destination side, where the shift-share structure allevi-

[14] A related concern is that the identifying variation may partly reflect predictable seasonal components of county-level temperature rather than genuinely idiosyncratic weather shocks. Appendix Figure G.22 pre-residualizes each county's temperature by county-by-year and county-by-quarter fixed effects before the Borusyak recentering, isolating only "unexpected" deviations from each county's own seasonal norm. Point estimates are numerically identical to the baseline, because the pre-residualization is algebraically absorbed by the municipality-by-year and municipality-by-quarter fixed effects already in the PPML specification. AKM standard errors widen by a factor of two to four, yet all six non-reference bins and the quartic coefficients retain significance at conventional levels.

[15] The key diagnostic is the stability of the contemporaneous coefficient, not the significance of the leads and lags themselves. Because temperature is serially correlated, adjacent quarters' shocks are mechanically related, producing significant lead and lag terms even under correct specification. Significant lags may also capture genuine persistence. Unlike a difference-in-differences setting where significant pre-trends threaten identification, temperature shocks are repeated stochastic events for which anticipation is implausible.

ates this problem: because each municipality's temperature exposure is a share-weighted average across geographically dispersed US counties, the effective autocorrelation in the constructed instrument is substantially lower than for any single location's raw temperature. Appendix Figure C.12 presents destination-side cumulative impulse response functions. The contemporaneous warm-bin effects are significant at impact but confidence intervals widen at longer horizons, suggesting that destination-side disruptions are largely transitory (panel a). Extreme cold in the US (<0 °C) stands out with a large and persistent cumulative effect of approximately −4.2 percentage points by Q4, consistent with sustained labor market disruptions during prolonged winter episodes (panel b).

Border proximity: One concern is that the origin-side results may be driven by border municipalities, whose remittance patterns could reflect cross-border commuting or other proximity effects rather than temperature-driven transfers. Appendix Figure C.13 maps the distance from each municipality's population centroid to the border. Appendix Figure C.14 progressively excludes municipalities within 50, 100, 200, and 400 km of the US–Mexico border and re-estimate both the quartic polynomial and non-parametric bin specifications. The U-shaped dose-response curve is remarkably stable across all subsamples. In the bin estimates, the moderate heat effect (25–30 °C) remains significant at the one percent level throughout, and cold effects (<10 °C) actually *strengthen* as border municipalities are excluded, rising from 0.25% to 0.46%. The extreme heat bin (>30 °C) loses significance beyond 200 km, likely because the hottest municipalities are disproportionately located near the northern border. On the destination side (Figure C.14), we progressively exclude US counties near the border, renormalizing migrant shares and rebuilding the shift-share instrument at each threshold. The warm-bin effects (25–30 and 30–45 °C) are robust through 100 km but attenuate beyond 200 km. This pattern is expected: the major migrant-receiving counties that drive the heat channel—Harris, Dallas, and Hidalgo in Texas, Maricopa and Pima in Arizona—are located 200–500 km from the border, so excluding them removes the primary source of identifying variation for heat exposure. The cold-bin effects (0–5 and 5–10 °C) strengthen as border counties are excluded, reflecting increased weight on northern counties with greater cold exposure.

Binning bias correction: Appendix Figure C.10 and Figure C.11 address the concern raised by Jones et al. (2026) that within-unit variation in temperature bin counts may be partly predictable from long-run warming trends, thereby biasing non-parametric estimates. We construct counterfactual expected temperature exposure for each geographic unit by estimating linear warming trends per unit-month, shrinking them via empirical Bayes toward month-specific means (to accommodate seasonal heterogeneity in warming rates), de-trending the daily temperature distribution, and shifting the resulting ECDF forward by the estimated trend. Including these expected exposures as controls absorbs the predictable component of temperature variation. For the parametric specifications, the counterfactual enters as a polynomial of the expected quarterly mean; for the non-parametric specifications, it enters as expected bin counts. A joint Wald F-test confirms that the counterfactual controls are highly significant ($p < 0.001$) on both sides. On the origin side (Figure C.10; see also Figure C.10 for the parametric curves), the correction has negligible effects: the quartic dose-response curve and all four non-parametric bins are virtually unchanged. The cold bin (<10 °C) moves from 0.246% to 0.248% per day, and the warm bins strengthen slightly (the >30 °C coefficient rises from 0.048% to 0.098%), confirming that secular warming trends in Mexico do not meaningfully confound the origin-side estimates. On the destination side (Figure C.11), the non-parametric bins provide the clearest test: all three warm bins (20–25, 25–30, 30–45 °C) remain negative and significant, with magnitudes that strengthen under the correction (the (30–45] °C coefficient increases from −0.214% to −0.304%), while the

moderate cold bins (0–5 and 5–10 °C) shrink substantially and lose significance. Extreme cold (<0 °C) strengthens from −0.094% to −0.185%. The quartic dose-response curve shows a steeper warm-side decline after the correction, consistent with the bin results, though the cold side of the polynomial becomes poorly identified: with the predictable moderate-cold variation absorbed, the quartic lacks sufficient signal to anchor the cold tail, producing boundary oscillations characteristic of polynomial extrapolation. This limitation is specific to the parametric specification in the shift-share structure, where the share-weighted counterfactual expected temperature is highly collinear with the actual share-weighted temperature. The non-parametric bins, which are more flexibly identified, provide the definitive test. Secular warming has progressively and predictably reduced moderate cold days (0–10 °C) in US counties; the baseline estimates confounded this trend with a causal effect. Once the predictable component is absorbed, only truly extreme cold retains a large negative effect. We adopt the month-EB corrected estimates as the preferred specification and use them in the climate projection exercises.

3.4 Mechanisms

This section investigates the mechanisms underlying the estimated temperature-remittance relationships by examining outcomes directly on both sides of the migrant corridor.

3.4.1 New migration from Mexico

The results pertaining the communities of origin show that temperature extremes in Mexico increase remittance inflows. A natural question is whether this reflects higher transfers to existing households (the insurance motive) or an increase in the number of migrants (the extensive margin). If temperature shocks induce new migration, the estimated origin-side effects could partly capture additional remitters rather than larger transfers.

We do not expect the extensive margin to drive our results for two reasons. First, our identification exploits quarterly temperature variation, and previous work has found that the effect of temperature on migration is rarely immediate. Mexican migrant households typically require a significant amount of time after an adverse shock-such as a deportation- to save enough money to migrate to the United States (Osuna-Gomez and Medina-Cortina, 2025). Even in contexts with fewer restrictions on cross-border movement, migration responses take at least a year to materialize (Ibáñez et al., 2025; Baylis et al., 2025). Second, if migrants decided to leave their municipalities for another location within Mexico, we would expect this internal migration to cause a bias toward-zero (as affected households leave the municipality of record and new remittances would head to the new municipality). Therefore, our estimates are a low-bound of the true effect of temperature.

To test whether extreme temperatures cause new international migration, we use the number of new MCAS registrations as a proxy for new migration to the United States. We estimate:

$$\ln(\text{NewMCAS}_{my}) = f(\text{Temp})_{my} + \Delta X_{my} + \delta_{sy} + \delta_m + \epsilon_{my} \tag{7}$$

where $\ln(\text{NewMCAS}_{my})$ is the natural logarithm of new MCAS registrations from municipality m in year y, $f(\text{Temp})_{my}$ is the non-parametric temperature function, and ΔX_{my} are time-varying weather controls.

The specification includes state-by-year (δ_{sy}) and municipality (δ_m) fixed effects. Identification comes from within-municipality temperature variation after accounting for state-wide trends.

Table 2 presents the results. None of the temperature bin coefficients are statistically significant, and the point estimates are close to zero across all bins. These results do not imply that temperature shocks never cause migration, but they suggest that in the short-term horizon of our analysis (leveraging deviations within a year), temperature shocks do not generate a measurable migration response and are not likely to drive the coefficients from previous sections.

Table 2: Effect of Mexican temperature on new MCAS registrations

	<10 °C	10–15 °C	15–20 °C	20–25 °C	25–30 °C	>30 °C
	0.0012	0.0003	0.0003	0.0000	0.0000	0.0003
	(0.0011)	(0.0004)	(0.0003)	(0.0000)	(0.0006)	(0.0008)
Municipalities	2,445	2,445	2,445	2,445	2,445	2,445

Notes: This table presents non-parametric estimates of equation (7), where the dependent variable is the log number of new MCAS registrations per municipality-year. The reference category is 20–25 °C. All specifications include municipality fixed effects, state-by-year fixed effects, and weather controls (precipitation, relative humidity, atmospheric pressure, and wind speed). Standard errors clustered at the municipality level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

3.4.2 Temperature and labor outcomes in Mexico

To provide suggestive evidence on the mechanisms driving the response of shocks in communities of origin, we estimate the effect of daily temperature on individual labor outcomes using person-day data from the Mexican National Survey of Occupation and Employment (ENOE, 2013–2024). The specification is:

$$y_{it} = \sum_{k} \sum_{\ell=0}^{7} \beta_{k\ell} \, \mathbf{1}\left[T_{m(i),t-\ell} \in B_k\right] + \mathbf{X}'_{m(i)t}\boldsymbol{\gamma} + \alpha_i + \mu_c + \nu_y + \omega_d + \varepsilon_{it} \tag{8}$$

where y_{it} is the labor outcome for person i on day t; $T_{m(i),t-\ell}$ is the mean temperature in municipality $m(i)$ on day $t-\ell$; B_k denotes temperature bins (<10°C, 10–15°C, 25–30°C, >30°C) with 15–25°C as the omitted category; $\mathbf{X}_{m(i)t}$ includes precipitation and relative humidity; α_i are person fixed effects; μ_c are calendar-month fixed effects; ν_y are year fixed effects; and ω_d are day-of-week fixed effects. Time-invariant individual characteristics (e.g., gender, education) and variables collinear with person and year fixed effects (e.g., age) are absorbed by α_i and ν_y. Standard errors are clustered at the municipality level. We report cumulative effects $\hat{\beta}_k = \sum_{\ell=0}^{7} \hat{\beta}_{k\ell}$ with delta-method standard errors. The sample comprises 25 million person-day observations across 1.6 million individuals in 209 municipalities. This specification follows Graff Zivin and Neidell (2014), identifying temperature effects from within-person deviations relative to the typical monthly climate.[16]

Table 3 reports cumulative effects on whether the individual worked on a given day (a dummy equal to one if the individual worked in a given day, zero otherwise). The full-sample results show that both cold (<10°C: −0.8 pp) and heat (>30°C: −1.2 pp) reduce the probability of working, with the heat effect stronger and more precisely estimated. Splitting by employment type, formal workers drive the heat response (−1.4 pp at >30°C

[16]Results are robust to a more demanding specification that replaces additive month and year fixed effects with municipality-by-month-of-year interactions.

vs. −0.6 pp for informal workers), consistent with schedule rigidity limiting adaptation: formal workers on fixed schedules cannot shift work to cooler hours, making full-day absences more likely on extreme heat days. The migrant-network split reveals that heat effects are concentrated in municipalities with above-median US migrant populations (MCAS 2012 stock): >30°C reduces work probability by 1.3 pp in high-migration municipalities but has no effect in low-migration ones. This pattern seems to be aligned with the fact that migrant communities in Mexico tend to be more rural and more dependent on agriculture Riosmena and Massey (2012). Finally, the age gradient shows that moderate heat sensitivity (25–30°C) increases with age, consistent with physiological vulnerability.

Table 3: Temperature Effects on Daily Work Probability: Full Sample and Heterogeneity

	Full	**Formal**	**Informal**	**High MCAS**	**Low MCAS**	**Young (≤30)**	**Prime (31–50)**	**Old (51+)**
<10°C	−0.008**	−0.004	−0.008	−0.006	−0.018**	−0.011**	−0.002	−0.013
	(0.004)	(0.004)	(0.006)	(0.004)	(0.008)	(0.006)	(0.004)	(0.008)
10–15°C	−0.002	−0.005**	+0.003	−0.001	−0.004	−0.000	−0.003	−0.004
	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
25–30°C	−0.006***	−0.008***	−0.004*	−0.008***	+0.003	−0.005**	−0.007***	−0.010***
	(0.002)	(0.002)	(0.003)	(0.002)	(0.005)	(0.002)	(0.002)	(0.003)
>30°C	−0.012***	−0.014***	−0.006	−0.013***	−0.002	−0.013***	−0.014***	−0.014***
	(0.002)	(0.003)	(0.004)	(0.002)	(0.010)	(0.003)	(0.003)	(0.005)
R^2	0.454	0.503	0.462	0.457	0.438	0.448	0.489	0.499
N (millions)	25.0	11.6	6.7	22.1	2.9	10.2	9.9	4.9

Notes: Dependent variable: worked on day t (0/1). Cumulative effects summed over lags L0–L7. Reference bin: 15–25°C. All specifications include person, month-of-year, year, and day-of-week fixed effects, with precipitation and relative humidity as controls. Standard errors clustered at the municipality level in parentheses. MCAS split: municipalities above/below median 2012 consular registration stock. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3.4.3 Migrant labor outcomes in the United States

The destination-side results show that temperature extremes in US counties reduce remittances sent to Mexico. The proposed mechanism is that heat and cold may depress migrants' earnings, reducing the surplus available for transfers. There are different possible reasons behind this response: temperature shocks might reduce total labor supply, by reducing the number of hours the migrant can work or by changing their unemployment status, or they may affect the productivity of each hour they work.

To assess the effect of temperature on migrants' labor supply and earnings, we use individual-level data from the CPS Basic Monthly Survey (2005–2024), which we prefer over the American Community Survey for three reasons.[17] First, the CPS reference week (the Sunday-through-Saturday interval containing the 12th of each month) records *actual* hours worked during a specific week, whereas the ACS asks about "usual hours per week in the past 12 months," conflating temperature conditions across seasons. Second, the CPS 4-8-4

[17]We focus on Mexican-born men because (i) their labor force participation is nearly universal, avoiding selection issues; (ii) they constitute the majority of undocumented migrants; and (iii) they are disproportionately employed in climate-exposed industries (agriculture and construction).

rotation design yields a short panel (up to 8 observations per person over 16 months), supporting person fixed effects via the longitudinal identifier that absorbs time-invariant worker heterogeneity (occupation, skill, residential sorting). Third, county-level geographic identification is available for approximately 41% of observations, enabling temperature measurement at the county-day level rather than the coarser PUMA geography available in the ACS.

We restrict our analysis to the county-identified sample and compute temperature exposure over the 21 days preceding each reference week (days -21 to -1). This pre-period window captures the cumulative effects of temperature on labor outcomes. Equation 9 presents the econometric strategy to identify the effect. In it, q_{ict} is the outcome for individual i in county c and year-month t; $\text{PreBinDays}_{k,ct}$ counts the days (out of 21) in temperature bin k (reference: 10–20 °C); α_i absorbs time-invariant individual characteristics; $\delta_{c,m}$ absorbs county-specific seasonality; and $\delta_{s,y,m}$ absorbs state-level monthly shocks. Standard errors are clustered at the county level.

$$q_{ict} = \sum_k \beta_k \, \text{PreBinDays}_{k,ct} + \gamma_1 \overline{\text{Rain}}_{ct} + \gamma_2 \overline{\text{RH}}_{ct} + \alpha_i + \delta_{c,m} + \delta_{s,y,m} + \varepsilon_{ict} \tag{9}$$

Table 4 reports results for three outcomes: log weekly hours (conditional on employment), employment (LPM on the full labor force), and imputed weekly earnings estimated by Poisson pseudo-maximum likelihood (unconditional, with zeros for non-workers). For the earnings measure, we compute each person's average hourly wage from months when they appear in the Outgoing Rotation Group (MISH 4 or 8) and multiply by actual hours in all months, setting earnings to zero for non-workers. We estimate column (3) via PPML (Santos Silva and Tenreyro, 2006), which naturally accommodates the mass of zeros, yields semi-elasticity coefficients directly, and is consistent under heteroskedasticity. With person fixed effects, the wage level is absorbed; the coefficient captures the dollar-weighted hours response to temperature.

The results reveal a clear margin decomposition. Warm pre-period temperatures significantly reduce hours worked (column 1): an additional day above 20 °C in the three weeks before the reference week reduces weekly hours by 0.23–0.31%, with a monotonic heat gradient. Employment, however, is unaffected across all temperature bins (column 2), establishing that the destination-side channel operates through fewer hours per employed worker, rather than job losses. The null employment result also implies that the conditional hours estimates are not contaminated by temperature-driven selection into work.

Column (3) reports PPML estimates for imputed weekly earnings. Because earnings are constructed by multiplying a person-level average wage (from ORG months) by actual hours, these estimates should be interpreted with caution: they mechanically reflect the hours margin scaled by a fixed wage, not an independent measure of income. With this caveat, the warm-side coefficients corroborate the hours results, showing earnings reductions of 0.30–0.55% per additional warm pre-period day. The positive cold-extreme coefficient (+1.7%) lacks a clear counterpart in hours or employment and likely reflects noise in a thin tail of the temperature distribution.

Figure 8 presents the quartic polynomial response function estimated from equation (9) with mean pre-period temperature replacing the bins, separately for three nested populations. The temperature-hours relationship is flat for all workers (left panel), shows a slight negative curvature for foreign-born males

Table 4: Effect of US pre-period temperature on Mexican-born male labor outcomes (CPS)

	(1) Ln(hours)	(2) Employment	(3) Earnings (PPML)
<0 °C	−0.0031	0.0018	0.0174**
	(0.0037)	(0.0026)	(0.0077)
0–5 °C	0.0030	−0.0014	0.0042
	(0.0023)	(0.0020)	(0.0039)
5–10 °C	−0.0002	−0.0001	0.0023
	(0.0012)	(0.0011)	(0.0019)
20–25 °C	−0.0023***	−0.0001	−0.0030*
	(0.0008)	(0.0004)	(0.0016)
25–30 °C	−0.0024**	−0.0004	−0.0039
	(0.0011)	(0.0007)	(0.0026)
>30 °C	−0.0031*	−0.0003	−0.0055
	(0.0016)	(0.0011)	(0.0034)
Observations	134,247	173,177	125,872
R^2	0.531	0.721	—
Person FE	✓	✓	✓
County × month FE	✓	✓	✓
State × year × month FE	✓	✓	✓

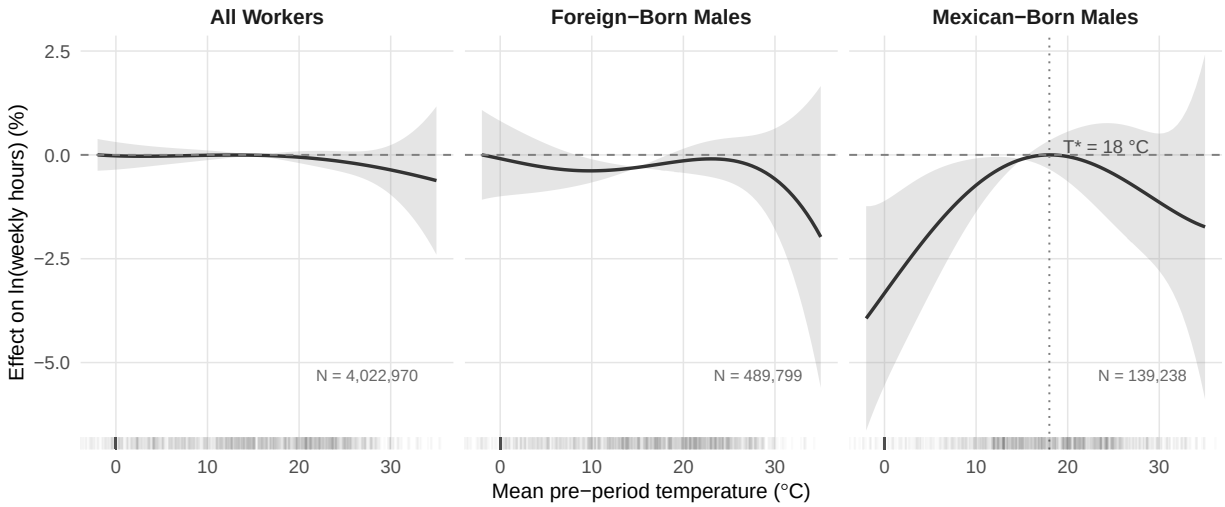
Notes: CPS Basic Monthly, 2005–2024, Mexican-born males aged 18–65 with county identification. Temperature bins count days (out of 21) in the pre-period (days −21 to −1 before the reference week). Reference category: 10–20 °C. Column (1): log actual weekly hours, conditional on employment. Column (2): employment indicator (LPM) on the full labor force. Column (3): imputed weekly earnings (person-average hourly wage from ORG months × actual hours; zero for non-workers) estimated by Poisson pseudo-maximum likelihood; coefficients are semi-elasticities. All specifications include person, county × calendar-month, and state × year × month fixed effects, weighted by WTFINL. Standard errors clustered at the county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

(center), and reveals a pronounced inverted-U for Mexican-born males (right), peaking at $T^* \approx 18$ °C. The cold flank is statistically significant below 8 °C, while the warm flank turns negative above 22 °C with widening confidence intervals reflecting sparser data at the hot tail. The optimum closely matches the destination-side aggregate result ($T^* \approx 17$ °C from the Bartik-weighted quarterly panel), providing micro-level confirmation that the destination-side temperature-remittance relationship operates through migrant labor outcomes. Both exhibits are robust to the Jones et al. (2026) binning bias correction. Adding 21 per-day continuous temperatures (one for each day of the pre-period) alongside the bin counts leaves the bin coefficients essentially unchanged: the largest shift is 0.0016 in the hours model and 0.0008 in the earnings model. Although the 21 per-day controls are jointly significant in a Wald test ($p < 0.01$ for all three models), the economically relevant bin coefficients are unaffected. The quartic specification is immune by construction, as it uses a smooth polynomial rather than discrete bins.

4 Climate Change Projections

We now use our estimated temperature-remittance relationships to project the effect of climate change on remittance flows through mid-century and end-of-century. We combine the non-parametric bin coefficients from Section 3 with downscaled temperature projections from the CMCC-ESM2 Earth System Model (Cherchi et al., 2019) under two Shared Socioeconomic Pathway (SSP) scenarios: SSP2-4.5, a moderate mitigation pathway in which global warming reaches approximately 2.7°C by 2100, and SSP5-8.5, a high-emissions scenario reaching approximately 4.4°C (IPCC, 2021). This exercise quantifies how projected warming will reshape the temperature distributions faced by Mexican municipalities and US migrant-receiving counties,

Figure 8: Effect of US pre-period temperature on ln(weekly hours): quartic polynomial



Notes: Each panel shows the quartic polynomial response of log weekly hours to mean 21-day pre-period temperature, estimated from equation (9) with person fixed effects, county × calendar-month fixed effects, and state × year × month fixed effects. The curve is normalized to zero at its maximum. Shaded areas denote 95% confidence intervals from county-clustered standard errors. Rug plots show the distribution of mean pre-period temperatures in each sample. CPS Basic Monthly, 2005–2024, county-identified observations.

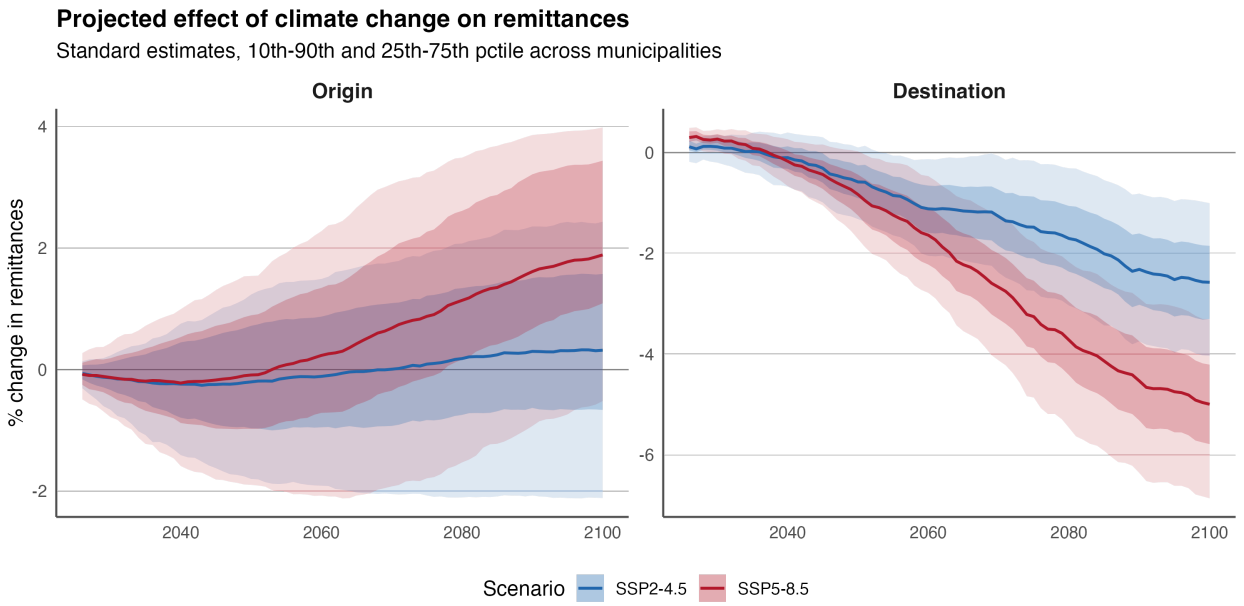
and translates those shifts into expected changes in remittance flows.

We employ the delta method, a standard approach in the climate impacts literature (Hsiang et al., 2017; Carleton et al., 2022; Burke et al., 2015a). For each administrative unit, calendar quarter, target year, and SSP scenario, we compute the difference between the CMCC-ESM2 projected mean temperature and the model's historical baseline (2000–2014), then shift the observed ERA5 daily temperature distribution by this delta and recount the days in each temperature bin. This preserves observed day-to-day variability while incorporating the projected mean shift. The projected change in remittances for each municipality is then computed by applying the estimated bin coefficients from Table B.2 to the change in bin days: on the origin side, directly to the projected change in Mexican temperature bins; on the destination side, to the change in Bartik-weighted US temperature exposure using the 2012 MCAS migrant shares. Following Borusyak et al. (2022), we apply the coefficients estimated on recentered exposure to the total change in exposure. We report projections using both the standard baseline estimates and the bias-corrected estimates from the Jones et al. (2026) correction. Appendix F provides the formal projection equations and full methodological details.

Figure 9 presents the projected time path of climate change effects on remittances from 2026 to 2100 under both SSP scenarios using the standard estimates. The left panel shows the origin side. The projected origin-side effects are modest in magnitude and near zero through mid-century, reflecting the fact that Mexican temperatures are currently concentrated near the thermal comfort zone, where the dose-response curve is relatively flat. Under the moderate SSP2-4.5 scenario, warming shifts the Mexican temperature distribution only slightly, producing a near-zero average effect by 2050 (−0.2%) that grows to +0.3% by 2100 as the rightward shift begins moving more days onto the steep warm flank of the U-shaped curve. Under the high-emissions SSP5-8.5 pathway, the effects are somewhat larger, reaching +1.9% by end-of-century, as

substantial warming pushes many municipality-quarter observations into the above-30°C range where each additional hot day raises remittance inflows.

Figure 9: Projected effect of climate change on remittances, 2026–2100



Notes: This figure shows the projected percentage change in remittances due to climate change under SSP2-4.5 (blue) and SSP5-8.5 (red), using the standard (baseline) non-parametric bin estimates. The left panel shows origin-side effects (Mexican temperature shifts); the right panel shows destination-side effects (US temperature shifts transmitted through migrant networks). Solid lines indicate the cross-municipality mean; dark shading denotes the 25th–75th percentile range and light shading the 10th–90th percentile range. Projections use the delta method applied to CMCC-ESM2 temperature output with ERA5 observed daily temperatures as the baseline (see Appendix F).

The origin-side response reflects a tension between two offsetting forces. As warming reduces the number of cold days in Mexico, the cold-bin contribution to remittance inflows decreases. Simultaneously, the increase in hot days raises remittance inflows through the warm bins. By mid-century, these two forces nearly cancel: the reduction in days below 15°C (which carried large positive coefficients) offsets the increase in days above 25°C. By 2100 under SSP5-8.5, the hot-bin contribution dominates, with the above-30°C bin alone contributing approximately 2 percentage points to the projected increase, as Mexico's temperature distribution shifts substantially into the extreme heat range (Figure F.20a in the appendix decomposes these contributions by bin).

The bias-corrected estimates from the Jones et al. (2026) correction amplify the origin-side projections. After correction, the above-30°C coefficient nearly doubles (from 0.048% to 0.093% per day), producing larger projected increases: +0.5% by 2050 and +4.3% by 2100 under SSP5-8.5 (compared to −0.1% and +1.9% with the standard estimates). The divergence grows over time because the bias-corrected coefficients place more weight on extreme heat, the bin whose day counts increase most dramatically under sustained warming (Figure F.20c,d in the appendix overlay the standard and bias-corrected time paths).

The destination-side effects are larger in magnitude and unambiguously negative (Figure 9, right panel). Warming in the United States shifts the temperature distribution in migrant-receiving counties further onto the steep hot flank of the inverted-U dose-response curve, reducing migrants' earnings capacity and hence

remittance outflows. Under SSP2-4.5, the average municipality experiences a destination-side decline of −0.6% by 2050, growing to −2.6% by 2100. Under SSP5-8.5, the decline reaches −0.8% at mid-century and −5.0% by end-of-century. These losses exceed the origin-side gains by a factor of two to three across all scenario-horizon combinations, implying that the net effect of climate change on remittances is negative: climate change simultaneously increases the need for transfers and reduces migrants' capacity to provide them, but the destination-side constraint dominates. The cross-municipality dispersion of destination-side effects (shaded ribbons in Figure 9) is narrower than on the origin side, reflecting the diversification inherent in migrant networks, where the Bartik-weighted temperature exposure averages across diverse US climatic conditions.

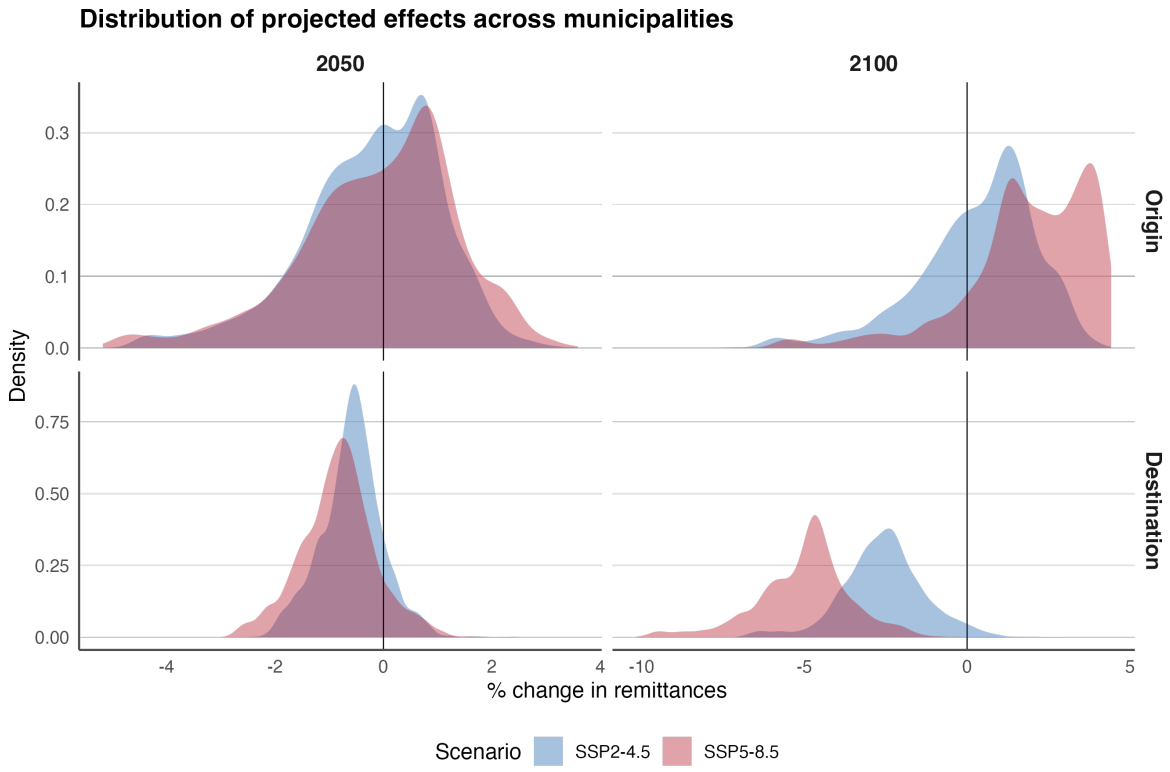
The bias-corrected estimates amplify the destination-side projections as well (Figure F.20c,d in the appendix). After the Jones et al. (2026) correction, the extreme cold coefficient (<0°C) strengthens substantially (from −0.094% to −0.238% per day), while the warm-bin coefficients remain similar. The net result is a projected destination-side decline of −1.7% by 2050 and −6.8% by 2100 under SSP5-8.5. Across all specifications, the bias-corrected estimates are systematically larger, suggesting that the standard estimates may understate the projected impacts (Figure F.20b in the appendix summarizes the full set of projections).

In dollar terms, using the 2024 baseline of $64 billion in annual remittances to Mexico (Banco de México, 2024), the projected destination-side losses range from $0.4 billion per year (SSP2-4.5, 2050, standard) to $4.4 billion per year (SSP5-8.5, 2100, bias-corrected). To put these magnitudes in perspective, the end-of-century destination-side loss under SSP5-8.5 is equivalent to approximately 7% of the budget of Mexico's current social protection system, or roughly the entire annual allocation of several major cash transfer programs. The origin-side effects partially offset these losses, but even under the most favorable interpretation (bias-corrected, SSP5-8.5, 2100), the origin-side increase of $2.8 billion offsets only about 60% of the destination-side decline. The distribution of effects across municipalities is heterogeneous (Figure 10). Under SSP5-8.5 by 2100, the 10th–90th percentile range of destination-side effects spans approximately −7% to −3%, with municipalities whose migrants are concentrated in already-hot US counties (Texas, Arizona) facing the largest losses.

These projections rest on several assumptions. First, the delta method assumes that the *shape* of the daily temperature distribution within each quarter is preserved, with only the mean shifting according to the climate model's projection. In practice, climate change may also alter the variance and higher moments of the temperature distribution, potentially increasing the frequency of extreme events beyond what a simple mean shift implies (IPCC, 2021). To the extent that this occurs, our projections would understate the true effect on the tails of the temperature distribution, particularly the extreme heat bin that drives the largest coefficients.

Second, the projections assume time-invariant coefficients, meaning that the estimated temperature-remittance relationships from our 2013–2024 sample period will continue to hold through 2100. This rules out behavioral adaptation on both sides of the corridor. Mexican households may develop alternative coping mechanisms (e.g., formal insurance, social protection programs, improved infrastructure), reducing their reliance on remittances during temperature shocks. Evidence that remittances already fund adaptive investments such as air conditioning (Randazzo et al., 2023) suggests that this channel is active and could evolve as warming progresses. US employers may adopt heat-mitigation technologies (e.g., workplace cooling, schedule adjustments) that attenuate the labor productivity losses from warming. If such adaptations occur,

Figure 10: Distribution of projected effects across municipalities



Notes: This figure shows the density distribution of projected municipality-level effects for 2050 and 2100, separately by channel (origin/destination) and SSP scenario. The origin-side distributions shift rightward over time as warming increases hot days. The destination-side distributions shift leftward, with the SSP5-8.5 distribution at 2100 centered near −5% to −7%.

our projections overstate the actual effects. Conversely, if warming pushes temperatures into ranges rarely observed in our sample period, the extrapolation beyond the support of the historical data may understate the effects, as the true dose-response curve could steepen at temperatures outside the observed range.

Third, we use a single climate model (CMCC-ESM2) rather than a multi-model ensemble. While CMCC-ESM2 is a well-validated Earth System Model that participates in the CMIP6 intercomparison project (Cherchi et al., 2019), relying on a single model does not capture the full range of climate projection uncertainty. Different models produce different spatial patterns of warming, and a multi-model ensemble would provide a more robust characterization of projection uncertainty, particularly at the regional scale relevant for this analysis.

Fourth, the projections hold the migrant network structure fixed at its 2012 configuration. In reality, climate change may itself alter migration patterns, as households in the most affected regions may send more (or fewer) migrants, and migrants may relocate to different US counties in response to local climate conditions (Cattaneo and Peri, 2016b). The destination-side projections also assume that the sectoral composition of migrant employment remains constant. If migrants shift toward less climate-exposed occupations over time, the destination-side effects would be smaller than projected.

Finally, the projections treat the origin and destination sides independently, whereas in practice they interact.

If destination-side constraints reduce remittances at the same time that origin-side shocks increase the need for transfers, the welfare consequences for recipient households may be larger than the sum of the two channels would suggest, as the insurance function of remittances fails precisely when it is most needed.

Despite these caveats, the projection exercise serves its primary purpose: translating the estimated causal relationships into economically interpretable magnitudes under plausible warming scenarios, and highlighting the asymmetry between the origin and destination channels that will shape the future of this important financial flow.

5 Conclusion

This paper shows that temperature shocks affect remittances through two distinct and opposing channels. On the origin side, temperature extremes in Mexico increase remittance inflows, consistent with an insurance motive. On the destination side, temperature extremes in US destination counties reduce remittances by depressing migrant earnings, particularly in outdoor-intensive sectors. The destination-side losses dominate: warming constrains migrants' capacity to remit more than it stimulates the need for transfers, implying that the insurance function of remittances is itself vulnerable to climate conditions at the destination.

Our findings highlight a previously unrecognized climate risk for remittance-dependent households: the financial flows on which millions of families rely for consumption smoothing are sensitive to warming at both ends of the migration corridor. Climate projections suggest that this vulnerability will deepen, as destination-side losses grow faster than origin-side gains. For policymakers in origin countries, social protection systems should account for the possibility that remittances may decline precisely when they are needed most. Moreover, the dependence of migrant communities on remittances, together with the fact that conditions in destination communities affect the most the flow of remittances, imply that origin communities should prepare and adapt to a warming climate where remittances cannot be as reliable. For destination countries, the results indicate that migrant labor, for which certain industries rely on, is particularly sensitive to temperature, and they may need to reconsider workplace heat standards and occupational safety enforcement in order to keep productivity levels. More broadly, our framework illustrates how bilateral financial flows can transmit climate shocks across borders, linking the welfare of non-migrant households in developing countries to weather realizations in distant labor markets.

References

Adao R., Kolesár M., and Morales E. (2019) Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134.

Afridi F., Mahajan K., and Sangwan N. (2022) The gendered effects of droughts: Production shocks and labor response in agriculture.. *Labour Economics*.

Agarwal R. and Horowitz A. W. (2002) Are International Remittances Altruism or Insurance? Evidence from Guyana Using Multiple-Migrant Households. *World Development*, 30(11): 2033–2044.

Aguilar-Gomez S., Gutierrez H.-D., E., Jaume D., and Tobal M. (2014) Thermal stress and financial distress: Extreme temperatures and firms' loan defaults in Mexico.. *Journal of Development Economics*, 168.

Ambrosius C., Quigua J., and Velasquez A. (2025) Externalities of U.S. Immigration Enforcement Policies in El Salvador. *Working Paper*.

Baldwin J. W., Benmarhnia T., Ebi K. L., Jay O., Lutsko N. J., and Vanos J. K. (2023) Humidity's Role in Heat-Related Health Outcomes: A Heated Debate. *Environmental Health Perspectives*, 131(5): 055001.

Banco de México (2024) Remesas. Sistema de Información Económica.

Barreca A., Clay K., Deschênes O., Greenstone M., and Shapiro J. (2016) Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the 20th Century. *Journal of Political Economy*, 124(1): 105–159, Includes state-by-month and month-by-year fixed effects and interacts temperature exposure with air conditioning penetration.

Baylis P., Bharadwaj P., Mullins J. T., and Obradovich N. (2025) Climate and migration in the United States.. *Journal of Public Economics*, 249.

Borusyak K. and Hull P. (2023) Nonrandom Exposure to Exogenous Shocks. *Econometrica*, 91(6): 2155–2185, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA19367.

Borusyak K., Hull P., and Jaravel X. (2022) Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89.

Bressler R. D., Papp A., Sarmiento L., Shrader J. G., and Wilson A. J. (2026) Occupation and temperature-related mortality in Mexico. *Journal of Human Resources*.

Bucheli J. R. and Fontenla M. (2025) The impact of return migration on economic development.. *Review of Economics and Statistics*: 393–407.

Burke M., Hsiang S., and Miguel E. (2015b) Higher temperatures, human capital, and economic growth in a changing climate. *Nature*, 527: 235–239.

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

——— (2023) The international transmission of local economic shocks through migrant networks.. *Journal of international economics*, 145.

Caballero M. E., Cadena B. C., and Kovak B. K. (2018) Measuring geographic migration patterns using matrículas consulares. *Demography*, 55: 1119–1145.

Carleton T., Duflo E., Jack K., and Zappalà G. (2025) Adaptation to Climate Change. Working Paper 33264, National Bureau of Economic Research.

Carleton T., Jina A., Delgado M., Greenstone M., Houser T., Hsiang S., Hultgren A., Kopp R. E., McCusker K. E., Nath I., Rising J., Rode A., Seo H. K., Viaene A., Yuan J., and Zhang A. T. (2022) Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits. *Quarterly Journal of Economics*, 137(4): 2037–2105.

Cattaneo C. and Peri G. (2016a) Migration, climate change, and shift-share designs. *Journal of Development Economics*, 122: 127–146.

——— (2016b) The migration response to increasing temperatures. *Journal of development economics*, 122: 127–146.

Cervantes J. (2015) Female migration and remittance flows to Mexico.Technical report, Center for Latin American Monetary Studies (cemla).

Cherchi A., Fogli P. G., Lovato T., Peano D., Iovino D., Gualdi S., Masina S., Scoccimarro E., Materia S., Bellucci A., and Navarra A. (2019) Global mean climate and main patterns of variability in the CMCC-CM2 models. *Journal of Advances in Modeling Earth Systems*, 11(12): 4497–4517.

Cohen F. and Dechezleprêtre A. (2022) Mortality, Temperature, and Public Health Provision: Evidence from Mexico. *American Economic Journal: Economic Policy*, XX(YY): ZZ–ZZ, Uses municipality-by-day and municipality-by-year fixed effects, and interacts the rollout of Seguro Popular with temperature exposure.

Congreso del Estado de Tlaxcala (2024) Presupuesto de Egresos del Estado de Tlaxcala para el Ejercicio Fiscal 2024. Decreto 317.

Conover E., Khamis M., and Pearlman S. (2023) Declining Outmigration and Local Labor Markets. *Journal of Human Resources*, 58: 1948–1973.

Dell M., Jones B., and Olken B. (2012) Temperature shocks and economic growth: Evidence from the developing world. *American Economic Journal: Macroeconomics*, 4(3): 66–95.

Deschênes O. and Greenstone M. (2011) Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US.. *American Economic Journal: Applied Economics*, 3.

Dinarte L., Jaume D., Medina-Cortina E., and Winkler H. (2021) Not by Land nor by Sea: The Rise of Formal Remittances during COVID-19. *Working Paper*.

Estrada F., Mendoza-Ponce A., Calderón-Bustamante O., and Botzen W. (2022) Impacts and economic costs of climate change on Mexican agriculture.. *Regional Environmental Change*, 22.

Faini R. (1994) Workers Remittances and the Real Exchange Rate: A Quantitative Framework. *Journal of Population Economics*, 7(2): 235–245, Publisher: Springer.

Frankel J. (2011) Are bilateral remittances countercyclical? *Open Economies Review*, 22: 1–16.

Goldsmith-Pinkham P., Sorkin I., and Swift H. (2020) Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110.

Graff Zivin J. and Neidell M. (2014) Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1): 1–26.

Hsiang S., Kopp R., Jina A., Rising J., Delgado M., Mohan S., Rasmussen D. J., Muir-Wood R., Wilson P., Oppenheimer M. et al. (2017) Estimating economic damage from climate change in the United States. *Science*, 356(6345): 1362–1369.

Ibáñez A. M., Quigua J., Romero M. J., and Velásquez A. (2025) Responses to Extreme Temperatures: Migrant Networks and International Migration from El Salvador.. *American Economic Journal: Economic Policy.*

IPCC (2021) *Climate Change 2021: The Physical Science Basis. Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*: Cambridge University Press.

Jaeger D. A., Ruist J., and Stuhler J. (2018) Shift-Share Instruments and the Impact of Immigration. *National Bureau of Economic Research No. w24285*.

Jones B. F., Moscona J., Olken B. A., and von Dessauer C. (2026) With or Without U? Binning Bias and the Causal Effects of Temperature Extremes. Working Paper 34671, National Bureau of Economic Research.

Joseph T., Nyarko Y., and Wang S. Y. (2018) Asymmetric information and remittances: Evidence from matched administrative data.. *American Economic Journal: Applied Economics*, 10.

Justino P., Monteiro G. F., Oliveira R. C., and Severnini E. R. (2025) Weathering challenges: Distributional impacts of climatic shocks on household consumption in Mozambique. *WIDER Working Paper*.

Leeper R. D., Harrington T., Palecki M. A., DePolt K., Scott E., Runkle J., and Diamond H. J. (2025) The Influence of Drought on Heat Wave Intensity, Duration, and Exposure. *Journal of Applied Meteorology and Climatology*.

Lucas R. E. B. and Stark O. (1985) Motivations to Remit: Evidence from Botswana. *Journal of Political Economy*, Publisher: The University of Chicago Press.

Malhi G. S., Kaur M., and Kaushik P. (2021) Impact of climate change on agriculture and its mitigation strategies: A review.. *Sustainability*, 13.

Massey D. S., Rugh J. S., and Pren K. A. (2010) The geography of undocumented Mexican migration. *Mexican studies/Estudios mexicanos*, 26(1): 129–152.

Michel-Cuello C. and Aguilar-Rivera N. (2022) Climate change effects on agricultural production systems in México.. *Handbook of Climate Change Across the Food Supply Chai*: 335–353.

Mohapatra S., Joseph G., and Ratha D. (2012) Remittances and natural disasters: ex-post response and contribution to ex-ante preparedness.. *Environment, Development and Sustainability*, 14: 365–387.

Monras J. (2020) Immigration and wage dynamics: Evidence from the mexican peso crisis.. *Journal of Political Economy*, 128.

Monterroso O. and Vilán D. (2025) Global Remittances Cycle. FEDS Notes, Board of Governors of the Federal Reserve System.

Ormandy D. and Ezratty V. (2012) Health and thermal comfort: From WHO guidance to housing strategies. *Energy Policy*, 49: 116–121.

Osuna-Gomez D. and Medina-Cortina E. (2025) The Effect of Immigration Enforcement Abroad on Immigrants' Home-Country Firms. *Journal of Political Economy: Microeconomics*.

Ares de Parga-Regalado G. and Prato M. (2026) The Effects of Declining U.S. Immigration on Local Labor Markets in Mexico. Upjohn Institute Working Paper 26-424, W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Parker S. W. and Todd P. E. (2017) Conditional Cash Transfers: The Case of Progresa/Oportunidades. *Journal of Economic Literature*, 55(3): 866–915.

Pearson T. (2023) US Immigration Enforcement and Mexican Labor Markets.. *Working Paper*.

Randazzo T., Pavanello F., and De Cian E. (2023) Adaptation to climate change: Air-conditioning and the role of remittances. *Journal of Environmental Economics and Management*, 120: 102818.

Riosmena F. and Massey D. S. (2012) Pathways to El Norte: Origins, destinations, and characteristics of Mexican migrants to the United States.. *International migration review*, 46: 3–36.

Santos Silva J. and Tenreyro S. (2006) The log of gravity. *Review of Economics and Statistics*, 88(4): 641–658.

Sarmiento L., Colelli F. P., and Pavanello F. (2025) Emergency department visits and temperature: Evidence from mexico.Technical report, Working paper, RFF.

Sarmiento L., Gilli M., Pavanello F., and Shayegh S. (2024) Income Shocks, Adaptation, and Temperature-Related Mortality: Evidence from the Mexican Labor Market.

Smith D. J. (2025) Climate Change, Violence, and Remittance Flows in Mexico. *International Migration Review*.

Somanathan E., Somanathan R., Sudarshan A., and Tewari M. (2021) The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6): 1797–1827.

Tymvios N., Behm J.-A. Y., M., and Johnson K. (2016) Heat Stress in the US Construction Industry.. *Environmental, Opportunities, and, challenges, Constructing, commitment, and, acknowledging, human, experiences*, 2.

Van Hook J., Morse A., Capps R., and Gelatt J. (2021) Uncertainty about the size of the unauthorized foreign-born population in the United States. *Demography*.

Wilson A. J., Bressler R. D., Ivanovich C., Tuholske C., Raymond C., Horton R. M., Sobel A., Kinney P., Cavazos T., and Shrader J. G. (2024) Heat disproportionately kills young people: Evidence from wet-bulb temperature in Mexico. *Science Advances*, 10(49): eadq3367.

World Bank (2023) Migration and Development Brief 39: Remittances Remain Resilient but Are Slowing.Technical report, World Bank.

Yang D. and Choi H. (2007) Remittances as insurance: Evidence from the Philippines. *World Bank Economic Review*, 21(2): 219–248.

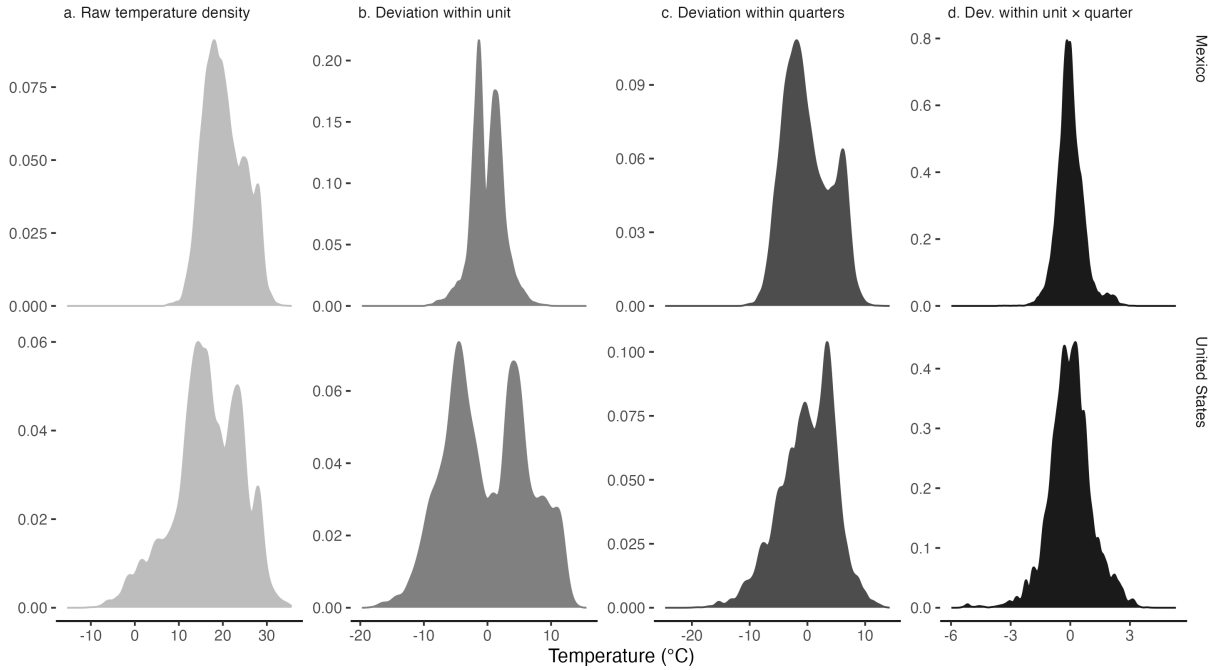
Zaiour R. (2023) Violence in Mexico, return intentions, and the integration of Mexican migrants in the US.. *Working Paper*.

Zhu J. L., Chau N., Rodewald A. D., and Garip F. (2024) Weather deviations linked to undocumented migration and return between Mexico and the United States.. *Proceedings of the National Academy of Sciences*, 121.

Online Appendix
The transnational effect of temperature on remittances

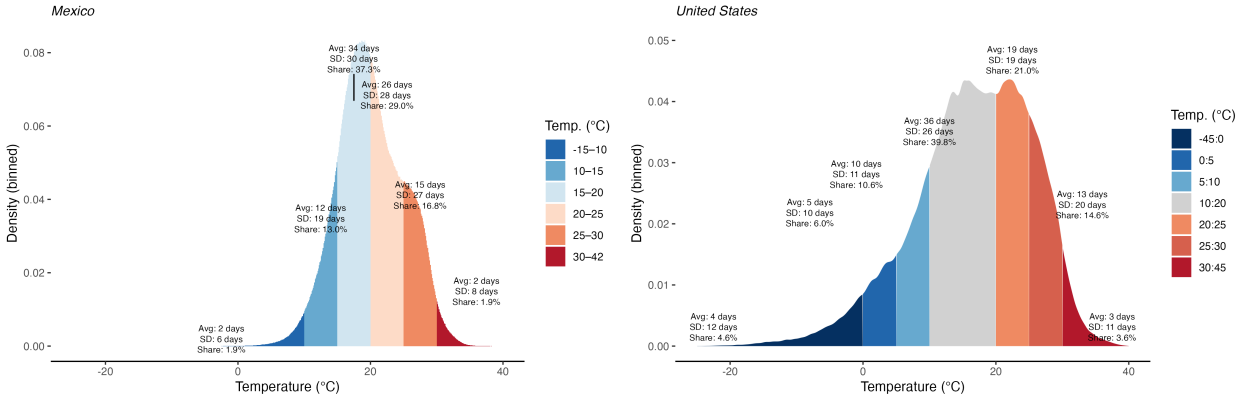
A Data

Figure A.1: Quarterly mean temperature distributions



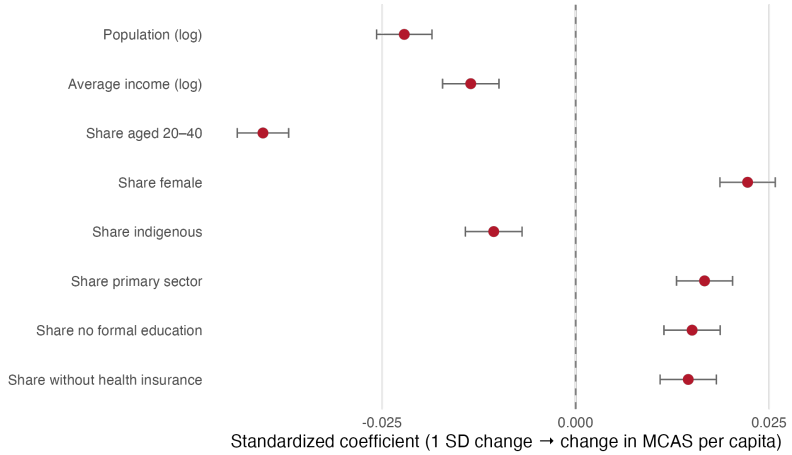
Notes: Density distribution of quarterly mean temperatures for Mexican municipalities (top row) and US counties hosting Mexican migrants per the 2012 MCAS (bottom row). Panel a shows the raw distribution; b shows the within-unit deviation (demeaned by municipality or county mean); c shows the within-quarter deviation (demeaned by quarter mean); d shows the residual variation after demeaning by both unit and quarter. This residual variation identifies the temperature effects in the parametric specifications.

Figure A.2: Daily temperature bin distributions



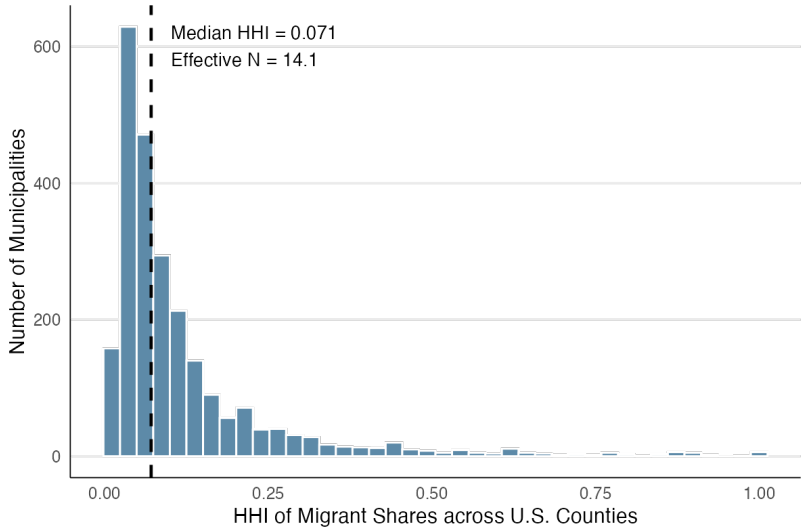
Notes: Distribution of daily temperatures across the bins used in the non-parametric specifications. Left panel: Mexican municipalities with six bins (<10, 10–15, 15–20, 20–25, 25–30, >30 °C); the reference category (15–25 °C) accounts for approximately 60% of all days. Right panel: US counties with seven bins (<0, 0–5, 5–10, 10–20, 20–25, 25–30, >30 °C); the reference category (10–20 °C) accounts for 30% of all days. Labels report the average number of days per quarter, the standard deviation, and the share of total days in each bin.

Figure A.3: Balance test: baseline municipality characteristics and migrant network intensity



Notes: Each point reports the coefficient from a bivariate OLS regression of MCAS per capita (2012 MCAS registrations divided by 2010 Census population) on a standardized (zero mean, unit variance) municipality characteristic from the 2010 Population Census. Horizontal bars denote 95% confidence intervals based on HC1 robust standard errors. Red points indicate statistical significance at the 5% level. The sample includes 2,436 municipalities with at least one MCAS registration. Municipalities with stronger migrant networks are systematically smaller, poorer, more female, more rural, and less educated, consistent with selective emigration of young working-age males. A joint regression of MCAS per capita on all eight characteristics yields an R^2 of 0.32, confirming substantial cross-sectional differences. These differences are absorbed by the municipality-by-year and municipality-by-quarter fixed effects used in all specifications.

Figure A.4: Geographic concentration of migrant networks



Notes: This figure shows the distribution of the Herfindahl-Hirschman Index (HHI) of migrant shares across Mexican municipalities, computed as $HHI_m = \sum_c s_{mc}^2$ where s_{mc} is the share of municipality m's registered migrants residing in US county c (2012 MCAS). Lower values indicate more geographically dispersed networks. The median HHI is 0.071, corresponding to an effective number of destination counties of approximately $1/0.071 \approx 14$. The median municipality sends migrants to 100 US counties, with the single largest destination hosting roughly 19% of its migrants. Only 3.4% of municipalities have $HHI > 0.5$ (highly concentrated networks). The dashed line marks the median.

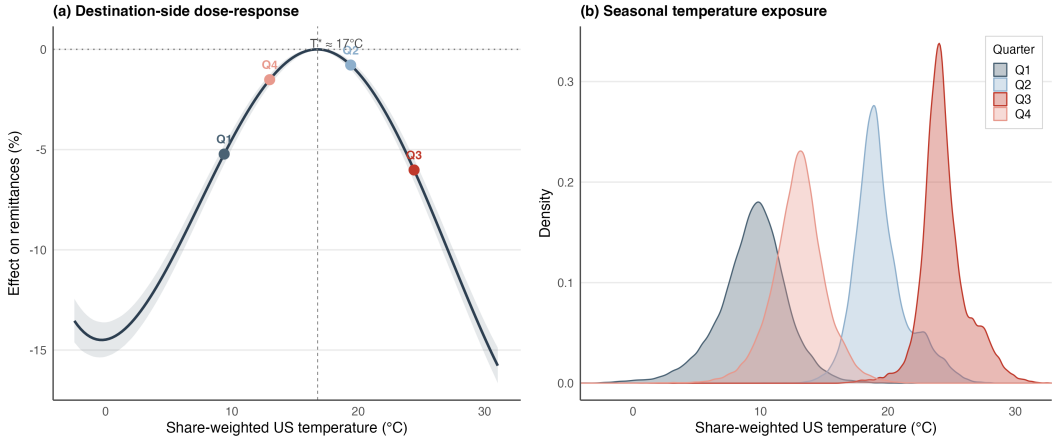
B Results

Table B.1: The effect of temperature on quarterly remittances: parametric specification

	(1) Origin Local MX temperature	(2) Destination Transnational US temperature
T / T^{US}	0.0441 (0.0705)	0.00123 (0.00369)
T^2 / $(T^{US})^2$	−0.00616 (0.00492)	0.00200*** (0.000684)
T^3 / $(T^{US})^3$	0.000262* (0.000149)	−0.000119*** (0.0000435)
T^4 / $(T^{US})^4$	−3.54e-06** (1.67e-06)	1.71e-06** (7.58e-07)
T^* (interior optimum)	20.4 °C (min)	≈ 17 °C (max)
Model information		
N.Obs	94,444	93,789
Avg population	49,756	—
Recentering FE	—	Mun × year-quarter
Weather controls	Yes	Yes
MX weather controls	—	Yes
FE: mun × year	Yes	Yes
FE: mun × quarter	Yes	Yes

Notes: This table reports quartic polynomial coefficient estimates from PPML specifications of remittances on temperature, with log population as an offset. Column (1) shows the origin-side (local) specification: the effect of Mexican municipality quarterly mean temperature on remittances received; T^* denotes the temperature that minimizes remittances. Column (2) shows the destination-side (transnational) specification: each polynomial power of US county-level quarterly mean temperature is recentered separately with municipality-by-time fixed effects before collapsing with pre-determined 2012 MCAS migrant shares; T^* denotes the temperature that maximizes remittances. Weather controls include precipitation and environmental variables; column (2) additionally includes Mexican weather controls. Standard errors are clustered at the municipality level. Significance codes: *** < 0.01,** < 0.05,* < 0.1.

Figure B.5: Seasonal temperature exposure of migrant networks



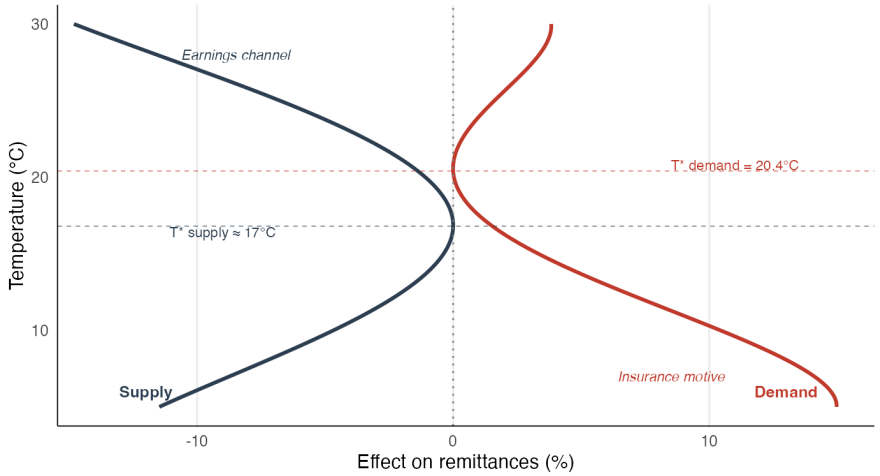
Notes: Panel (a) shows the quartic destination-side dose-response: the effect of share-weighted US temperature on remittances, normalized to zero at the optimum ($T^* \approx 17$ °C). Colored points mark each quarter's average network temperature on the curve. Panel (b) shows the distribution of share-weighted US temperature $\bar{T}_m = \sum_c s_{mc} T_c$ across municipalities, separately by quarter. In Q1 (winter), networks are centered around 8.7 °C on the steep cold flank of the dose-response, losing approximately 6% relative to the peak; in Q3 (summer), networks are at 24.9 °C on the hot flank, losing about 5%. In Q2 (spring) and Q4 (fall), networks sit near the flat peak where temperature shocks have minimal effects.

Table B.2: The effect of temperature on quarterly remittances: non-parametric specification

Panel A: Origin (local MX temperature, ref: 15–25 °C)						
	(−15:10]	(10:15]		(25:30]	(30:42]	
λ_b	0.00246***	0.00050*		0.00045***	0.00048*	
	(0.00070)	(0.00030)		(0.00012)	(0.00026)	
Days/quarter (SD)	5.22	18.24		27.37	8.37	
Days/quarter (Cell-SD)	0.83	3.07		3.72	1.23	
Panel B: Destination (transnational US temperature, ref: 10–20 °C)						
	(−45:0]	(0:5]	(5:10]	(20:25]	(25:30]	(30:45]
β_b	−0.00094***	−0.00258***	−0.00139***	−0.00091***	−0.00208***	−0.00214***
	(0.00015)	(0.00019)	(0.00018)	(0.00008)	(0.00007)	(0.00010)
Avg. days/quarter	6.15	7.98	11.47	18.57	14.89	2.11
Model information						
	Panel A			Panel B		
N.Obs	94,444			93,789		
Avg population	49,756			—		
Recentering FE	—			Mun × year-quarter		
Weather controls	Yes			Yes		
MX weather controls	—			Yes		
FE: mun × year	Yes			Yes		
FE: mun × quarter	Yes			Yes		

Notes: Non-parametric bin estimates from PPML specifications with log population offset. Panel A (origin): each λ_b is the effect of one additional day per quarter in each Mexican temperature bin (ref: 15–25 °C). Panel B (destination): each β_b is the effect of one additional share-weighted day in each US bin (ref: 10–20 °C), recentered with municipality-by-time FE and collapsed with 2012 MCAS shares. Weather controls as in Table B.1. SEs clustered at municipality level. $^{***} < 0.01, ^{**} < 0.05, ^{*} < 0.1$.

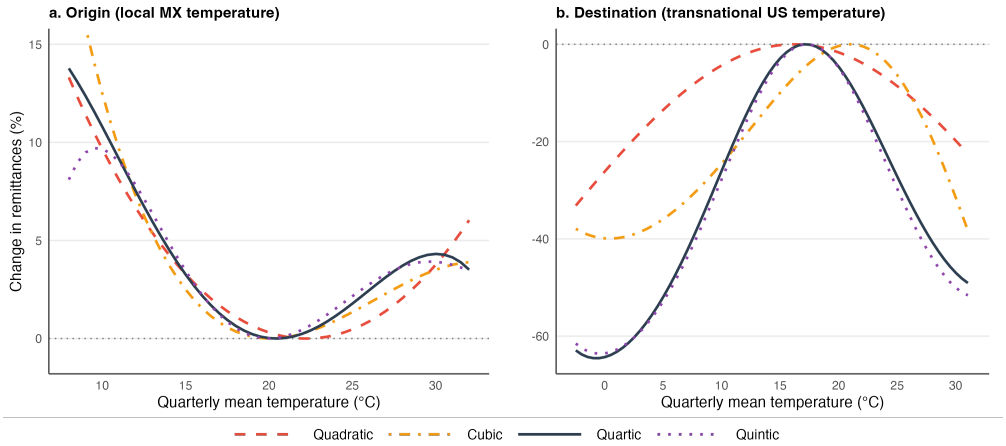
Figure B.6: Origin-side vs. destination-side dose-response curves



Notes: Panel (a) shows the origin-side quartic exposure-response curve: the effect of local (Mexican) temperature on remittances received, normalized to zero at the interior minimum ($T^* = 20.4$ °C). Panel (b) shows the destination-side quartic curve: the effect of US county temperature on remittances sent, normalized to zero at the maximum ($T^* \approx 17$ °C). The two curves are mirror images. On the origin side, temperature extremes in Mexico *increase* remittances, consistent with an insurance motive: households facing adverse weather conditions receive more transfers from abroad. On the destination side, temperature extremes in the US *decrease* remittances, consistent with an earnings channel: migrants exposed to extreme temperatures earn less and send less home.

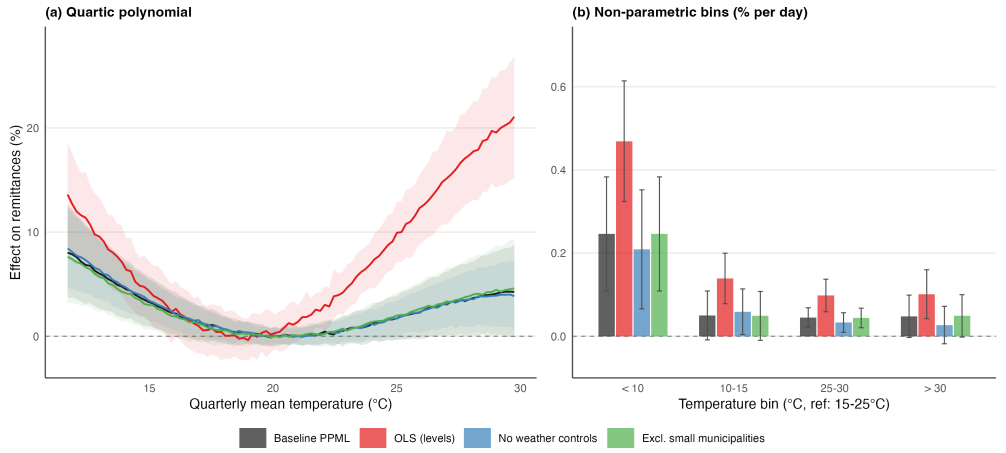
C Robustness

Figure C.7: Sensitivity of the dose-response curve to polynomial degree



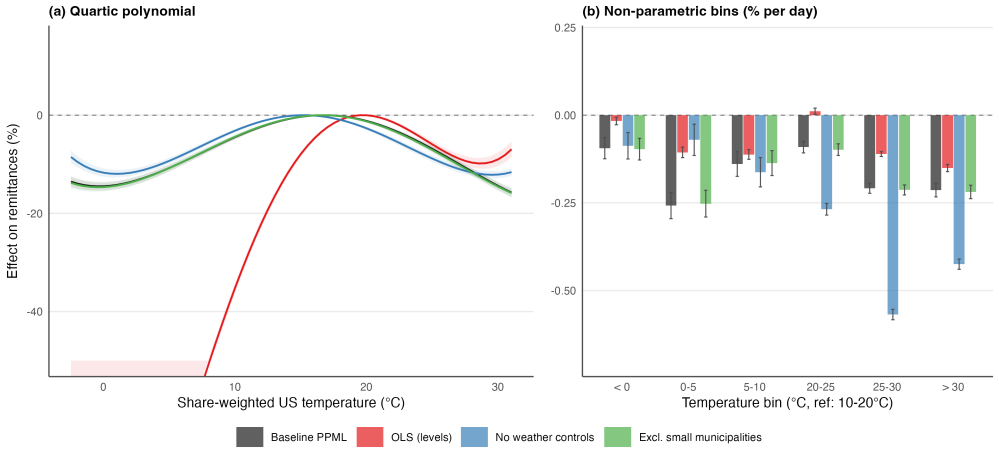
Notes: Each panel overlays the estimated dose-response curve for polynomial degrees 2 (quadratic), 3 (cubic), 4 (quartic, baseline), and 5 (quintic). Panel (a) shows the origin-side relationship between local Mexican temperature and remittances per capita (PPML), normalized so the minimum is zero; positive values indicate the percentage increase in remittances at temperature extremes relative to the optimum (≈20 °C). Panel (b) shows the destination-side relationship between transnational US temperature and remittances, using the share-weighted shift-share specification with full controls, normalized so the maximum is zero. The quartic (solid dark line) and quintic (dotted purple) are nearly indistinguishable on both sides, indicating that the quartic provides sufficient flexibility. All specifications include the full set of weather controls and municipality-by-year and municipality-by-quarter fixed effects. The origin-side polynomial is estimated directly on quarterly mean temperature via PPML; the destination-side polynomial uses the shift-share interaction specification evaluated at the migrant-share margin.

Figure C.8: Robustness to alternative specifications: origin side



Notes: Panel (a) overlays quartic polynomial dose-response curves from four variants of the origin-side specification. "Baseline PPML" is the main specification. "OLS (levels)" estimates remittances per capita in levels by OLS with population weights; the fitted curve is normalized by dividing through the sample mean of remittances per capita (≈80 pesos), converting the level effect to a percentage change comparable to the PPML semi-elasticity. "No weather controls" excludes precipitation, humidity, pressure, and wind speed from the PPML specification. "Drop low-pop (P25)" removes municipalities below the 25th percentile of population. All curves are normalized at their interior minimum, so values represent the percentage increase in remittances relative to the optimum. Shaded areas are pointwise 95% confidence bands clustered at the municipality level. Panel (b) reports the corresponding non-parametric bin estimates (reference: 15–25 °C). All specifications include municipality-by-year and municipality-by-quarter fixed effects.

Figure C.9: Robustness to alternative specifications: destination side



Notes: Panel (a) overlays quartic polynomial dose-response curves from four variants of the destination-side specification. "Baseline PPML" is the main specification. "OLS (levels)" estimates remittances in levels by OLS with log population as a regressor; the fitted curve is normalized by dividing through the sample mean of remittances, converting the level effect to a percentage change comparable to the PPML semi-elasticity. "No weather controls" excludes all share-weighted US weather and Mexican weather controls from the PPML specification. "Drop low-pop (P25)" removes municipalities below the 25th percentile of population. All curves are normalized at their interior maximum, so values represent the percentage change in remittances relative to the optimum temperature. Shaded areas are pointwise 90% confidence bands. Panel (b) reports the corresponding non-parametric bin estimates (reference: 10–20 °C) with 95% confidence intervals. All specifications include municipality-by-year and municipality-by-quarter fixed effects and use municipality-by-time recentered temperature variables (Eq. 4). Standard errors use the AKM variance estimator for PPML specifications and an analogous shift-share sandwich estimator for OLS.

Table C.3: Destination-side sensitivity to the COVID-19 pandemic

	(−45,0]	(0,5]	(5,10]	(20,25]	(25,30]	(30,45]	N
*Interaction model: temperature bins × **1**(2020–2021)*							
Base effect (non-COVID)	−0.00059***	−0.00336***	−0.00110***	−0.00108***	−0.00272***	−0.00278***	93,774
	(0.00016)	(0.00019)	(0.00019)	(0.00009)	(0.00008)	(0.00010)	
COVID differential (Δ)	−0.00258***	0.00349***	−0.00322***	0.00192***	0.00245***	0.00195***	
	(0.00056)	(0.00051)	(0.00044)	(0.00021)	(0.00020)	(0.00020)	
Total during COVID	−0.00318***	0.00013	−0.00432***	0.00084***	−0.00027	−0.00083***	
	(0.00053)	(0.00048)	(0.00039)	(0.00020)	(0.00018)	(0.00020)	
Subsample exercises							
Pre-COVID (2013–2020Q1)	−0.00258***	0.00085***	−0.00437***	0.00328***	−0.00057***	−0.00284***	50,311
	(0.00016)	(0.00018)	(0.00015)	(0.00009)	(0.00008)	(0.00013)	
Post-COVID (2022–2024)	−0.00579***	−0.00722***	−0.00027***	−0.00145***	−0.00006	−0.00075***	21,227
	(0.00016)	(0.00014)	(0.00011)	(0.00008)	(0.00007)	(0.00010)	

Notes: This table examines how the COVID-19 pandemic affected the destination-side temperature-remittance relationship. The top panel reports an interaction model estimated on the full sample, where each temperature bin is interacted with an indicator for 2020–2021. "Base effect" captures the temperature-remittance relationship outside that period; all six bins are negative and significant. "COVID differential" captures the additional effect during 2020–2021; the positive coefficients on the warm bins indicate that the pandemic *attenuated* the negative warm-temperature effect, consistent with disruption of outdoor labor markets, federal stimulus transfers, and the surge in remittance flows during 2020. "Total during COVID" is the linear combination of base and differential effects, with standard errors computed from the joint variance-covariance matrix. The bottom panel reports subsample exercises: the pre-COVID period (2013–2020Q1) confirms the inverted-U pattern existed before the pandemic, while the post-pandemic period (2022–2024) shows the negative relationship persisting. All specifications use the baseline controls and municipality-by-time fixed effects. Standard errors use the AKM variance estimator of Adao et al. (2019). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table C.4: Origin-side robustness to alternative inference strategies

	(−15,10]	(10,15]	(25,30]	(30,42]
Municipality (baseline)	0.00246***	0.00050*	0.00045***	0.00048*
	(0.00070)	(0.00030)	(0.00012)	(0.00026)
State	0.00246**	0.00050	0.00045**	0.00048
	(0.00111)	(0.00060)	(0.00020)	(0.00036)
Conley (25 km)	0.00246*	0.00050	0.00045***	0.00048
	(0.00126)	(0.00058)	(0.00017)	(0.00034)
Conley (50 km)	0.00246*	0.00050	0.00045**	0.00048
	(0.00150)	(0.00063)	(0.00020)	(0.00036)
Conley (100 km)	0.00246	0.00050	0.00045**	0.00048
	(0.00168)	(0.00066)	(0.00022)	(0.00041)
N	94,444	94,444	94,444	94,444

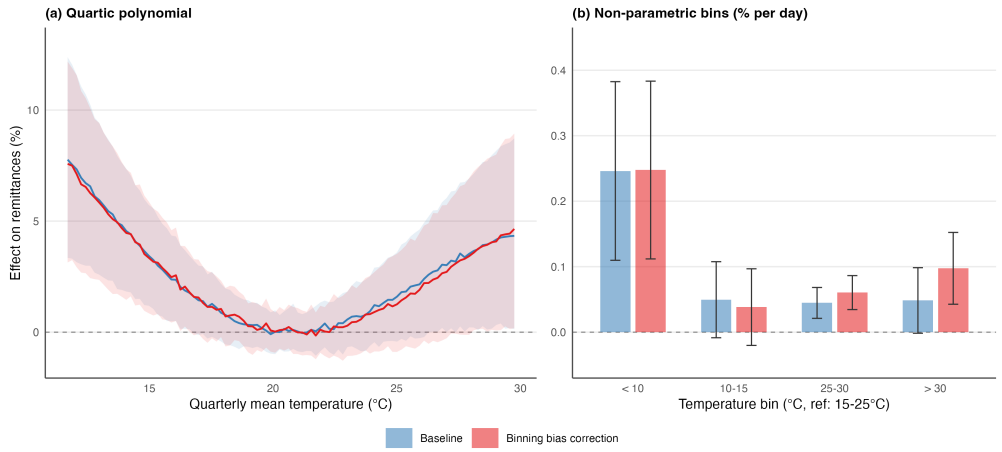
Notes: This table reports the origin-side non-parametric bin point estimates under alternative standard error computations. The reference category is 15–25 °C. Each row reports the same coefficients with a different variance estimator: municipality-level clustering (baseline), Mexican state-level clustering, and Conley spatial HAC standard errors at 25, 50, and 100 km distance cutoffs. The Conley estimator accounts for spatial correlation in weather shocks across nearby municipalities without imposing discrete cluster boundaries. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table C.5: Destination-side robustness to alternative inference strategies

	(−45,0]	(0,5]	(5,10]	(20,25]	(25,30]	(30,45]
AKM (baseline)	−0.00094***	−0.00258***	−0.00139***	−0.00091***	−0.00208***	−0.00214***
	(0.00015)	(0.00019)	(0.00018)	(0.00008)	(0.00007)	(0.00010)
Municipality	−0.00094	−0.00258*	−0.00139***	−0.00091	−0.00208**	−0.00214**
	(0.00105)	(0.00132)	(0.00052)	(0.00060)	(0.00082)	(0.00096)
State	−0.00094	−0.00258*	−0.00139*	−0.00091	−0.00208*	−0.00214*
	(0.00142)	(0.00153)	(0.00076)	(0.00109)	(0.00117)	(0.00111)
N	93,789	93,789	93,789	93,789	93,789	93,789

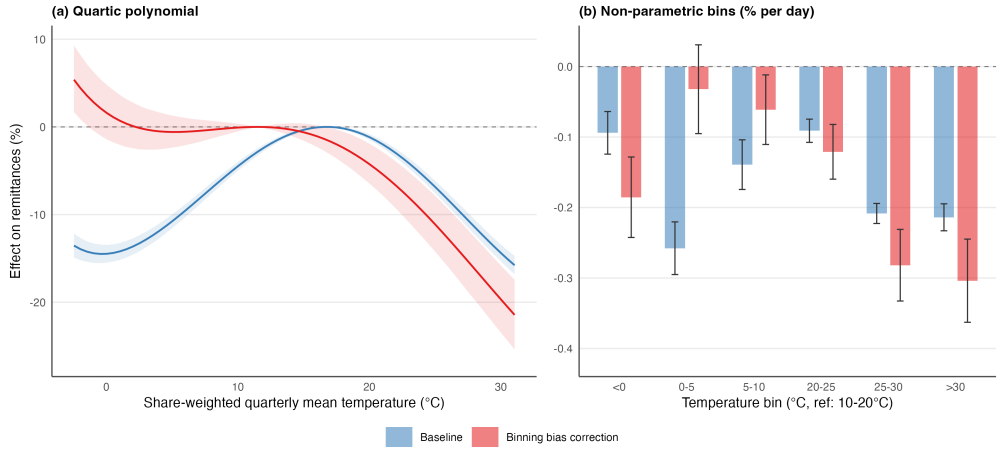
Notes: This table reports the destination-side non-parametric bin point estimates under alternative standard error computations. The reference category is 10–20 °C. Each row reports the same coefficients with a different variance estimator: the AKM estimator of Adao et al. (2019) clustering at the county × quarter level (baseline), municipality-level clustering, and Mexican state-level clustering. The AKM estimator yields smaller standard errors because it exploits the shift-share structure, concentrating inference on the county-level shock rather than the municipality level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure C.10: Origin-side robustness to binning bias correction (Jones et al., 2026)



Notes: This figure compares origin-side estimates with and without the binning bias correction of Jones et al. (2026). The correction constructs counterfactual expected temperature exposure by estimating municipality-by-month linear warming trends, shrinking them toward month-specific means via empirical Bayes, de-trending daily temperatures, and shifting the resulting ECDF forward by the estimated trend. Including the expected counts (or their polynomial projection) as additional controls absorbs the predictable component of temperature exposure. Panel (a) shows quartic polynomial dose-response curves with 95% simulation intervals (2,000 draws from the coefficient distribution); the correction enters as a fourth-degree polynomial of the counterfactual quarterly mean. Panel (b) shows non-parametric bin coefficients (days per quarter in each bin relative to 15–25 °C); the correction enters as five counterfactual bin counts. Both panels use the full specification with rain quintiles and relative humidity, atmospheric pressure, and wind speed quintile controls. The Wald F-statistic for joint significance of the counterfactual controls is 18.5 ($p < 0.001$). Standard errors clustered at the municipality level. $N = 94{,}429$.

Figure C.11: Destination-side robustness to binning bias correction (Jones et al., 2026)



Notes: This figure compares destination-side estimates with and without the binning bias correction of Jones et al. (2026). Panel (a) shows quartic polynomial dose-response curves with 95% simulation intervals, normalized to the interior local maximum (the optimal temperature). In the parametric specification, the correction enters as the counterfactual expected share-weighted quarterly mean temperature (linear). Panel (b) shows non-parametric bin coefficients (days per quarter in each bin relative to 10–20 °C); the correction enters as counterfactual expected bin counts, following the recommended functional-form matching of Jones et al. (2026). Both panels include the full set of controls (US precipitation, MX weather quartic). AKM standard errors (Adao et al., 2019). $N = 93{,}789$.

Table C.6: Rotemberg weight diagnostics for the destination-side shift-share design

	(−45,0]	(0,5]	(5,10]	(20,25]	(25,30]	(30,45]
Panel A: Summary statistics by temperature bin						
Effective N of counties	48.8	26.6	17.0	6.7	27.9	9.7
% counties with $\hat{\alpha} < 0$	15.2	30.9	40.9	33.0	26.5	34.2
Negative weight share (%)	1.8	4.1	3.3	3.0	4.0	7.4
Top-5 county share (%)	27.1	29.7	40.2	53.6	37.1	61.3
Top-10 county share (%)	37.3	39.7	52.7	62.9	48.1	80.6

Panel B: Top 10 counties by Rotemberg weight, (30,45] bin

	$\hat{\alpha}_c$ (%)	$\hat{\beta}_c^{\text{jid}}$	N municipalities
Harris, TX (Houston)	25.8	0.0018	1,888
Dallas, TX	12.5	0.0005	1,762
Hidalgo, TX (McAllen)	12.0	−0.0041	1,216
Pima, AZ (Tucson)	8.4	0.0025	816
Bexar, TX (San Antonio)	7.8	−0.0031	1,448
Maricopa, AZ (Phoenix)	6.5	−0.0032	1,826
Travis, TX (Austin)	4.8	0.0037	1,424
Tarrant, TX (Fort Worth)	3.6	0.0009	1,435
Fresno, CA	3.6	−0.0037	1,542
Webb, TX (Laredo)	2.5	−0.0058	918

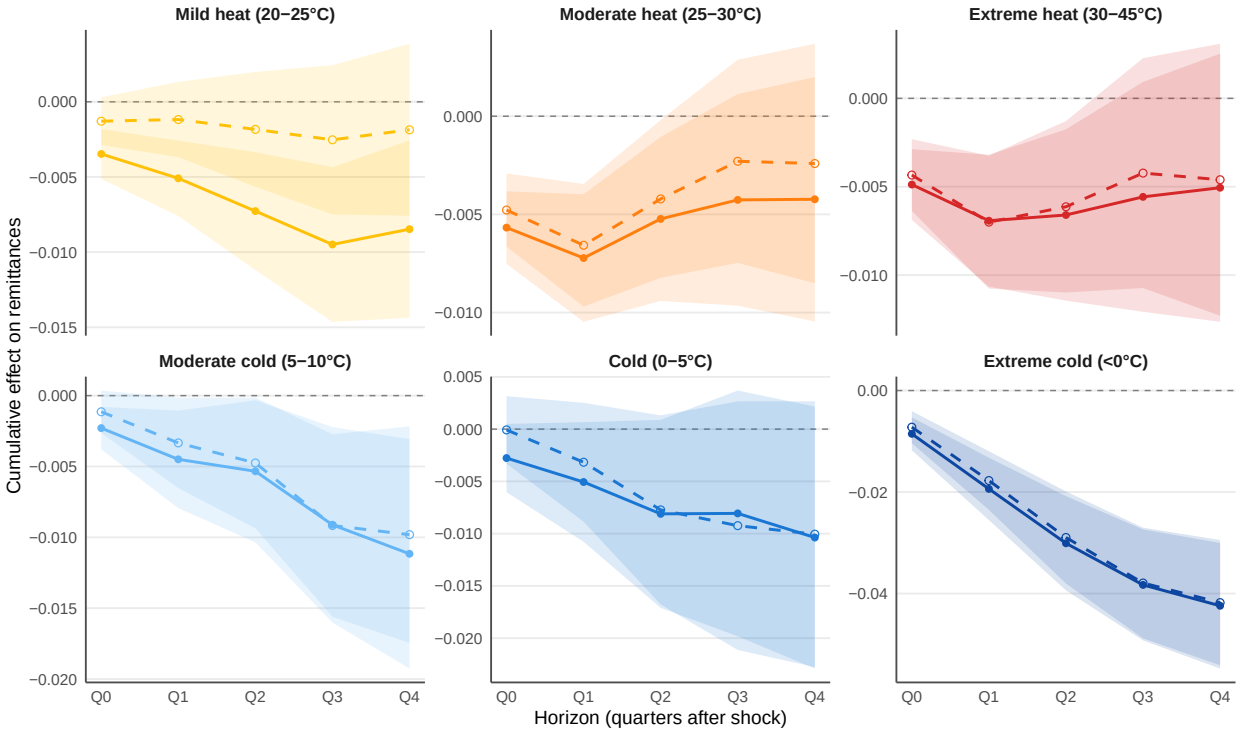
Notes: This table reports Rotemberg weight diagnostics (Goldsmith-Pinkham et al., 2020) for the destination-side shift-share (Bartik) design. The OLS coefficient on each temperature bin decomposes as $\hat{\beta}_b = \sum_c \hat{\alpha}_c^b \hat{\beta}_c^b$, where $\hat{\alpha}_c^b$ is the Rotemberg weight measuring county c's contribution to identification and $\hat{\beta}_c^b$ is the just-identified estimate using only county c's variation. Weights are computed via the Frisch-Waugh-Lovell residualization that partials out the other five bins, all weather controls, and municipality-by-year and municipality-by-quarter fixed effects. Panel A reports summary statistics: effective $N = 1/\text{HHI}$ of positive weights, the share of counties with negative weights, and the concentration of positive weight mass in the top-5 and top-10 counties. Panel B lists the ten most influential counties for the (30,45] °C bin, reporting each county's Rotemberg weight ($\hat{\alpha}_c$), just-identified OLS estimate ($\hat{\beta}_c^{\text{jid}}$), and the number of Mexican municipalities connected to it through the 2012 MCAS migrant network. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table C.7: Destination-side estimates excluding the ten most influential counties

	(−45,0]	(0,5]	(5,10]	(20,25]	(25,30]	(30,45]
Baseline	−0.094***	−0.258***	−0.139***	−0.091***	−0.209***	−0.214***
	(0.015)	(0.019)	(0.018)	(0.008)	(0.007)	(0.010)
Baseline + BB	−0.238***	0.071**	−0.013	−0.066***	−0.201***	−0.216***
	(0.029)	(0.030)	(0.025)	(0.022)	(0.028)	(0.032)
LOO-10	−0.199**	0.115	0.227*	−0.964***	−0.496***	−1.083***
	(0.090)	(0.100)	(0.135)	(0.083)	(0.070)	(0.100)
LOO-10 + BB	−0.236***	−0.025	0.369***	−0.789***	−0.544***	−0.465***
	(0.086)	(0.094)	(0.126)	(0.087)	(0.078)	(0.117)
N	93,789 / 93,745					

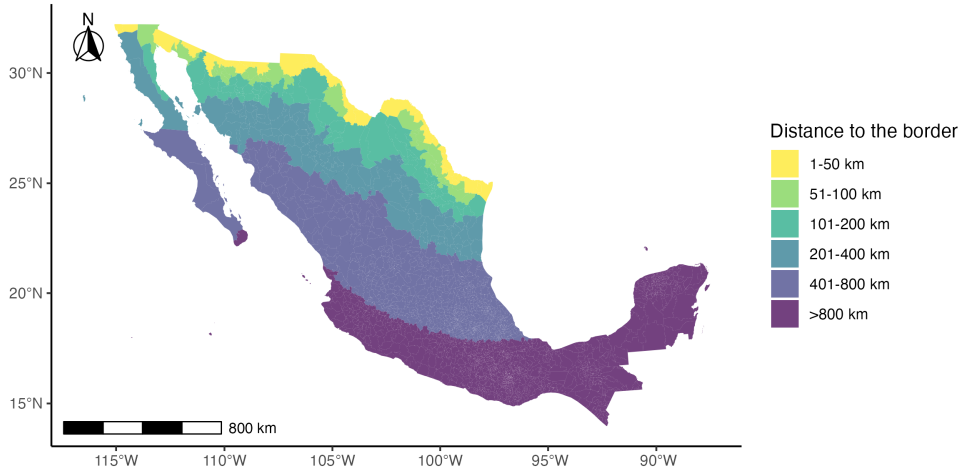
Notes: This table compares destination-side non-parametric bin estimates (days per quarter relative to 10–20 °C) using the full set of US counties with estimates that simultaneously exclude the ten counties with the largest Rotemberg weights for the (30,45] °C bin (see Table C.6, Panel B). After dropping these counties, migrant shares are renormalized and the shift-share instrument is reconstructed. "BB" denotes the binning bias correction of Jones et al. (2026). The warm bins (20–45 °C) remain robustly negative across all specifications. The moderate cold bins (0–5 and 5–10 °C) are fragile, changing sign in the LOO-10 sample, consistent with the binning bias diagnosis: secular warming has predictably reduced moderate cold days, confounding the baseline cold-bin estimates. Extreme cold (<0 °C) is stable across all specifications. AKM standard errors (Adao et al., 2019) for full-sample specifications; municipality-clustered standard errors for LOO-10 specifications. Coefficients are multiplied by 100 (percentage points per day). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure C.12: Distributed lag model: cumulative impulse response functions



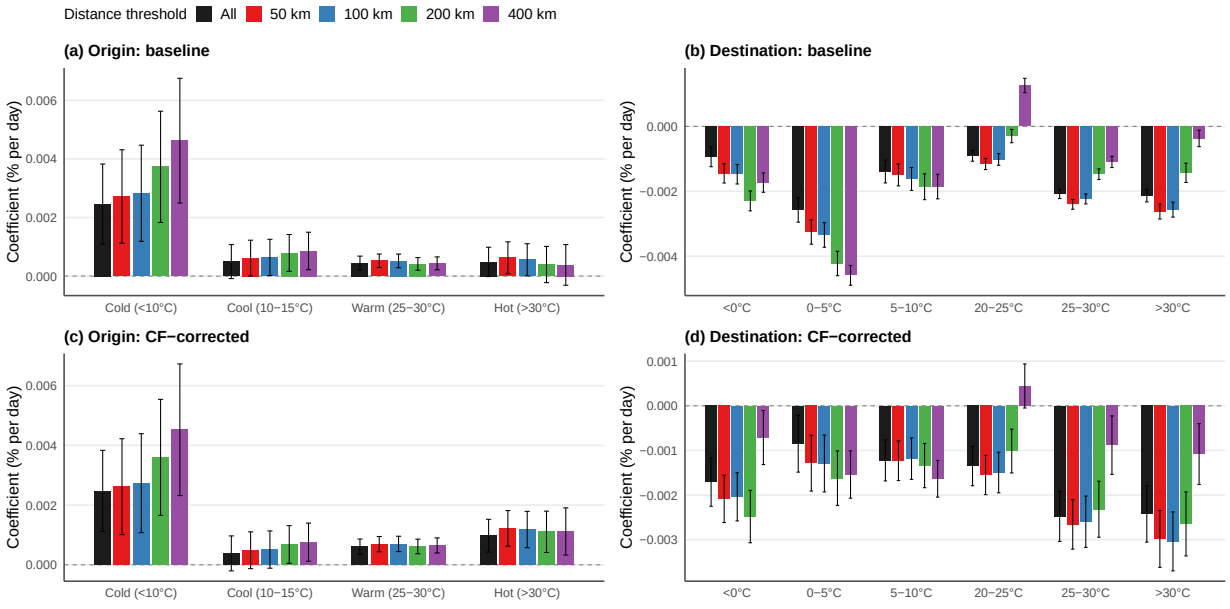
Notes: Each panel plots destination-side cumulative impulse response functions from a distributed lag model with four quarterly lags for one US temperature bin (via the shift-share instrument, reference: 10–20 °C). Solid lines show the baseline specification; dashed lines show the binning bias correction (Jones et al., 2026), which includes share-weighted counterfactual expected bin counts as additional controls (Wald $F = 25.4$, $p < 0.001$). The top row shows warm bins: all three show negative contemporaneous effects that are significant at impact, but confidence intervals widen at longer horizons, suggesting that destination-side heat disruptions are largely transitory. The BB correction has negligible effects on the warm bins. The bottom row shows cold bins: extreme cold (<0 °C) stands out with a large and persistent cumulative effect of approximately −4.2 percentage points by Q4, virtually unchanged after the BB correction (−4.18%). The moderate cold bins show modest attenuation at impact under the BB correction, but cumulative effects at Q4 are stable. Shaded areas represent 95% confidence intervals based on municipality-level clustering. The origin side is omitted because the strong seasonal autocorrelation of quarterly temperature ($r = 0.85$ between a given quarter and the same quarter one year prior) renders the four-lag structure unreliable after municipality-by-quarter fixed effects absorb average seasonal patterns.

Figure C.13: Distance of Mexican municipality centroids to the US–Mexico border



Notes: This figure maps the distance (in kilometers) from each Mexican municipality's population centroid to the nearest point on the US–Mexico border. Darker shading indicates greater distance. The border proximity robustness checks in Figure C.14 and Figure C.14 progressively exclude municipalities (origin side) or US counties (destination side) within 50, 100, 200, and 400 km of the border.

Figure C.14: Border distance robustness: baseline and CF-corrected estimates



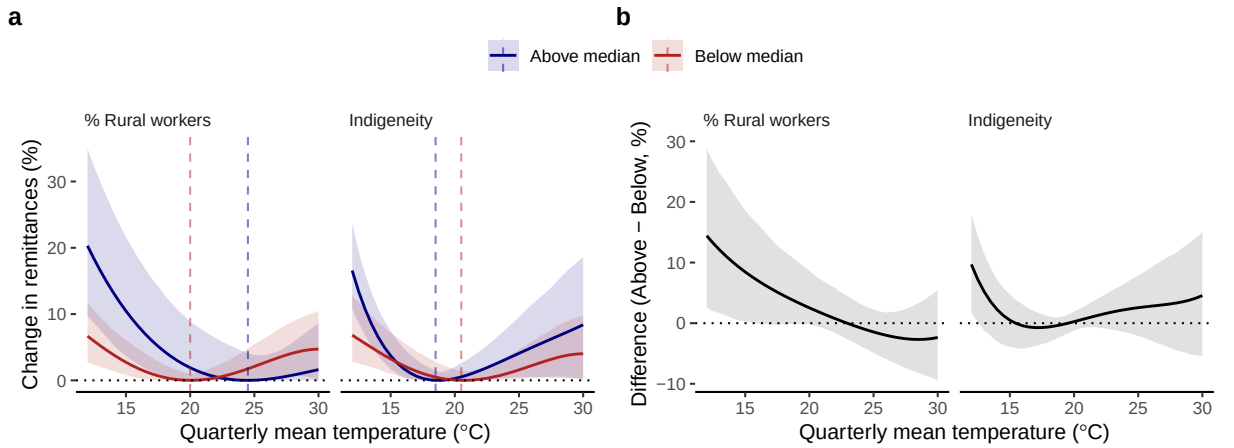
Notes: Each panel re-estimates the non-parametric bin specification after progressively excluding municipalities (origin side, panels a and c) or US counties (destination side, panels b and d) within 50, 100, 200, and 400 km of the US–Mexico border. The top row shows baseline estimates; the bottom row adds the Jones et al. (2026) binning bias correction (CF-corrected). Bars show point estimates; whiskers show 95% confidence intervals. On the origin side, the cold bin (<10 °C) strengthens as border municipalities are excluded, while the warm bins remain stable. On the destination side, cold bins (0–5 °C, 5–10 °C) strengthen as border counties are excluded; warm bins (25–30 °C, >30 °C) are stable through the 100 km threshold but attenuate beyond 200 km, consistent with major migrant-receiving counties in Texas and Arizona being located 200–500 km from the border. Municipality×year and municipality×quarter fixed effects; population weights (origin) or offset (destination); municipality-clustered standard errors (origin) and AKM standard errors (destination).

Heterogeneous effects by municipality characteristics

We examine whether the temperature-remittance relationship varies with observable characteristics of Mexican municipalities. We focus on two time-invariant characteristics that proxy for climate vulnerability and migrant occupational exposure: the share of the labor force in primary (agricultural) occupations and the indigenous population share. For each characteristic, we split municipalities at the median and estimate separate models for the above-median and below-median groups. Both variables are constructed from Mexican census data and averaged across available census years. We considered additional splits by municipality-level income and age structure but omit them for two reasons. First, income is highly correlated with both chosen variables (rural workforce share and indigeneity jointly capture economic marginalization), so it adds limited independent information. Second, the age composition of sending communities is endogenous to the migration process itself: selective emigration of young working-age adults mechanically reshapes the demographic profile of origin municipalities (Appendix Figure A.3), making it difficult to interpret heterogeneity along this dimension as reflecting pre-existing vulnerability rather than the consequences of migration.

Origin side. Figure C.15 reports the results for the origin side. Panel (a) shows quartic polynomial dose-response curves estimated separately for each subgroup. Each curve is normalized at its interior minimum (the temperature that minimizes remittances), so all values represent percentage increases relative to that optimum. Municipalities with higher shares of rural workers exhibit a steeper U-shaped response, with remittances rising to roughly 20% above the optimum at the coldest observed temperatures, compared to only 5–7% for less rural municipalities. A similar pattern appears for indigeneity: more indigenous municipalities show greater sensitivity to temperature deviations. Both patterns are consistent with the agricultural channel documented in prior work on temperature and economic activity in Mexico (e.g., Cohen and Dechezleprêtre, 2022; Sarmiento et al., 2024): municipalities more dependent on agriculture face larger local income shocks from temperature extremes, generating stronger remittance inflows as insurance. Panel (b) formalizes these comparisons by plotting the difference (above minus below median) in the normalized dose-response curves, with simulation-based 95% confidence bands. The rural workforce panel shows a positive and significant difference at cool temperatures that narrows toward zero at higher temperatures. The indigeneity panel displays a similar shape, though the confidence bands are wider.

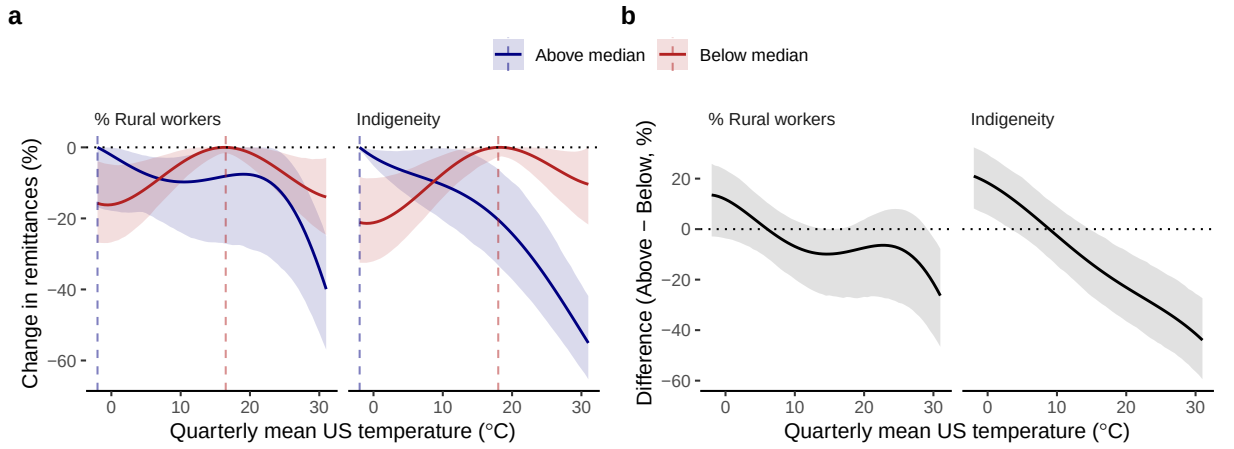
Figure C.15: Origin-side heterogeneity by rural workforce share and indigeneity



Notes: Panel (a) splits municipalities at the median of the rural workforce share (left) and indigenous population share (right), estimating separate quartic polynomial dose-response curves for each group. The outcome is remittances per capita (PPML with population weights); all specifications include municipality×year and municipality×quarter fixed effects, a quadratic in precipitation, and average relative humidity, wind speed, and atmospheric pressure. Curves are normalized at their respective interior minima (all effects ≥ 0) and converted to percentages via $(\exp(\hat{f}) - 1) \times 100$. Panel (b) shows the difference between above-median and below-median curves; positive values indicate that above-median municipalities experience a larger increase in remittances at that temperature. Shaded areas are pointwise 95% confidence bands from 2,000 multivariate normal draws. Municipality-clustered standard errors.

Destination side. Figure C.16 presents the analogous exercise for the destination side. Panel (a) shows quartic dose-response curves in the Borusyak-recentered shift-share US temperature variables, estimated separately for each subgroup. Each power of county-level temperature (T, T^2, T^3, T^4) is residualized against municipality×quarter fixed effects before collapsing with migrant shares, following the same approach as the main specification. The temperature grid spans the 1st to 99th percentiles of the share-weighted US temperature distribution, so the support reflects actual migrant exposure. Curves are normalized at their respective maxima, so all values represent percentage declines relative to the optimal US temperature. Both splits reveal that above-median municipalities (more rural, more indigenous) experience steeper destination-side declines at high US temperatures, reaching roughly -45% to -55% at 30°C compared to -10% to -20% for below-median municipalities. This is consistent with migrants from more rural and indigenous communities being disproportionately employed in outdoor, heat-exposed occupations in the United States (agriculture, construction, landscaping). Panel (b) plots the difference (above minus below median) in the normalized dose-response curves. For both splits, the difference is positive at cold US temperatures (above-median municipalities fare relatively better) and turns negative at high temperatures (above-median municipalities fare worse), with the indigeneity split showing a monotonically declining difference across the temperature range.

Figure C.16: Destination-side heterogeneity by rural workforce share and indigeneity



Notes: Panel (a) splits municipalities at the median of the rural workforce share (left) and indigenous population share (right), estimating separate quartic dose-response curves using Borusyak-recentered shift-share US temperature variables for each group. The outcome is total remittances (PPML with population offset); all specifications include share-weighted US weather controls, a quartic in Mexican temperature, and municipality×year and municipality×quarter fixed effects. The temperature grid spans the 1st to 99th percentiles of the MCAS share-weighted US county temperature distribution. The dose-response is evaluated as $\hat{f}(T) = \sum_{k=1}^{4} \hat{\beta}_k (T^k - \bar{T}^k)$, normalized at its maximum (all effects ≤ 0), and converted to percentages via $(\exp(\hat{f}) - 1) \times 100$. Panel (b) shows the difference between above-median and below-median curves; positive values indicate that above-median municipalities experience a smaller decline in remittances at that US temperature. Shaded areas are pointwise 95% confidence bands from 2,000 multivariate normal draws. Municipality-clustered standard errors.

D Validity of the Migrant Network Data

Our shift-share design relies on the geographic distribution of Mexican migrants across US counties, which we construct from the *Matrícula Consular de Alta Seguridad* (MCAS) registry. This appendix describes the institutional background of the MCAS and presents validation exercises that demonstrate the reliability of these data as a measure of the Mexican migrant population in the United States.

D.1 Institutional background on the MCAS

The MCAS is an identification card issued by Mexican consulates to Mexican nationals living abroad. The card serves as an official form of identification accepted by many US banks, landlords, and employers, making it a practical necessity for migrants regardless of their legal status. To obtain a card, applicants must schedule an appointment at their nearest Mexican consulate and present a birth certificate or valid passport along with proof of US residence. The card costs 39 USD, is valid for five years, and is typically issued within one business day.

For research purposes, the MCAS registry is valuable because each application records both the migrant's municipality of origin in Mexico and the US county of residence. This allows for the construction of a bilateral migration matrix linking Mexican municipalities to US counties. The registry is maintained by Mexico's *Secretaría de Relaciones Exteriores* (SRE) and, as of our data extract, covers the stock of valid cards as of 2012.

Prior research supports the representativeness of the MCAS as a measure of the Mexican migrant population. Caballero et al. (2018) estimate that approximately 38 percent of Mexican-born individuals residing in the United States hold an MCAS card. Massey et al. (2010) report even higher coverage among undocumented migrants, with 75 to 80 percent holding one, likely because these individuals have fewer alternative forms of identification. Cervantes (2015) provides additional evidence on the demographic composition of MCAS holders.

One limitation is that recently arrived migrants may not yet have obtained an MCAS card, so the registry could undercount new arrivals. However, new migrants overwhelmingly settle in locations where existing migrant networks are already established, so the geographic distribution of MCAS holders remains a reliable guide to migrant stocks even when levels are understated for recent cohorts.

D.2 Geographic consistency of migrant stocks and remittances

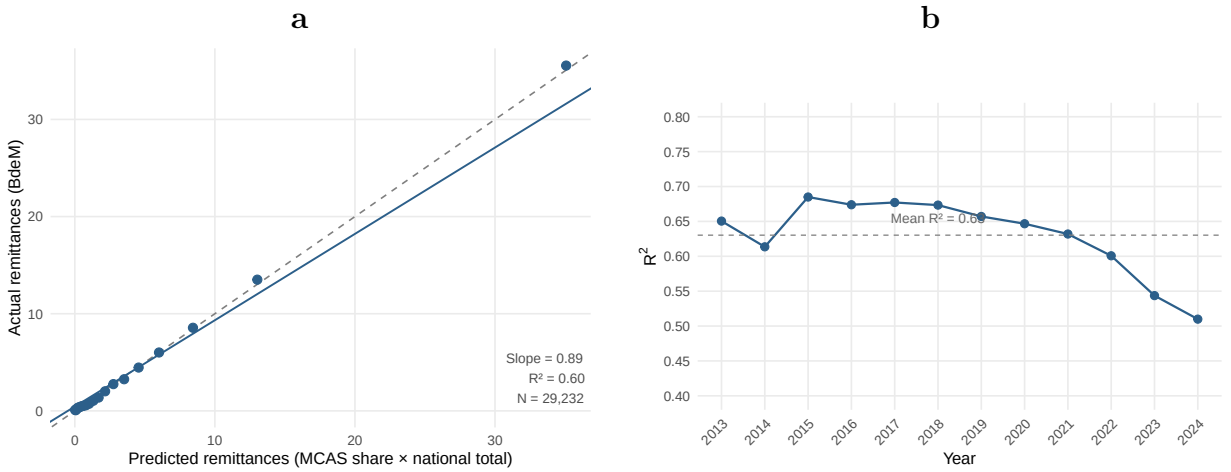
We validate the MCAS-based migrant shares by testing whether they predict the geographic distribution of remittances observed in independent administrative data from the Banco de México (BdeM).

Our primary validation exercise distributes total national remittance inflows to each Mexican municipality in proportion to its MCAS-based migrant share and compares these predicted remittances with actual municipality-level remittance inflows reported by BdeM. If the MCAS shares accurately capture the geographic distribution of remittance-sending migrants, predicted and actual remittances should be tightly correlated. Figure D.17a presents a binscatter of actual versus predicted remittances pooled across 2,436 municipalities and 12 years (2013–2024). The fitted slope is 0.89, close to the 45-degree line, and the R^2 is 0.60, indicating that MCAS shares alone explain roughly two-thirds of the cross-municipal variation in remittance inflows. Table D.8 reports formal regression results with and without a constant and with year fixed effects; in all specifications the slope on predicted remittances is close to and statistically indistinguishable from one.

An important concern is whether the 2012 MCAS shares remain predictive of remittance geography throughout our 2013–2024 sample period. Figure D.17b repeats the municipality-level validation for each year separately. The cross-sectional R^2 is stable between 0.61 and 0.69 through 2020, with a gradual decline to 0.51 by 2024 (mean $R^2 = 0.63$). This pattern is consistent with slow geographic redistribution of migrant networks

over time: the 2012 MCAS stock remains a strong predictor of remittance geography a decade later, although predictive power erodes in the most recent years as new migration destinations emerge. Importantly, the slope coefficients follow a similar trajectory, starting near unity in 2013 ($\hat{\beta} = 1.01$) and declining to 0.84 by 2024, suggesting that municipalities with historically large migrant networks grow more slowly than newer destinations. The municipality fixed effects in our main specifications absorb these level differences, so the gradual shift in the cross-sectional distribution does not threaten identification.

Figure D.17: MCAS validation: predicted vs. actual remittances



Notes: Panel (a): each dot represents the mean of actual and predicted remittances within one of 20 equal-sized bins of predicted remittances, pooled across 2,436 municipalities and 12 years (2013–2024). Predicted remittances allocate national BdeM totals to municipalities in proportion to their 2012 MCAS migrant shares. The solid line is the OLS fit; the dashed line is the 45-degree line. Both variables are measured as quarterly means in millions of USD. Panel (b): each point reports the R^2 from a cross-sectional regression of actual on predicted remittances estimated separately for each year. The dashed line marks the mean R^2 across all years (0.63). Sample: 2,436 municipalities per year.

Table D.8: Regression of actual on predicted municipality-level remittances

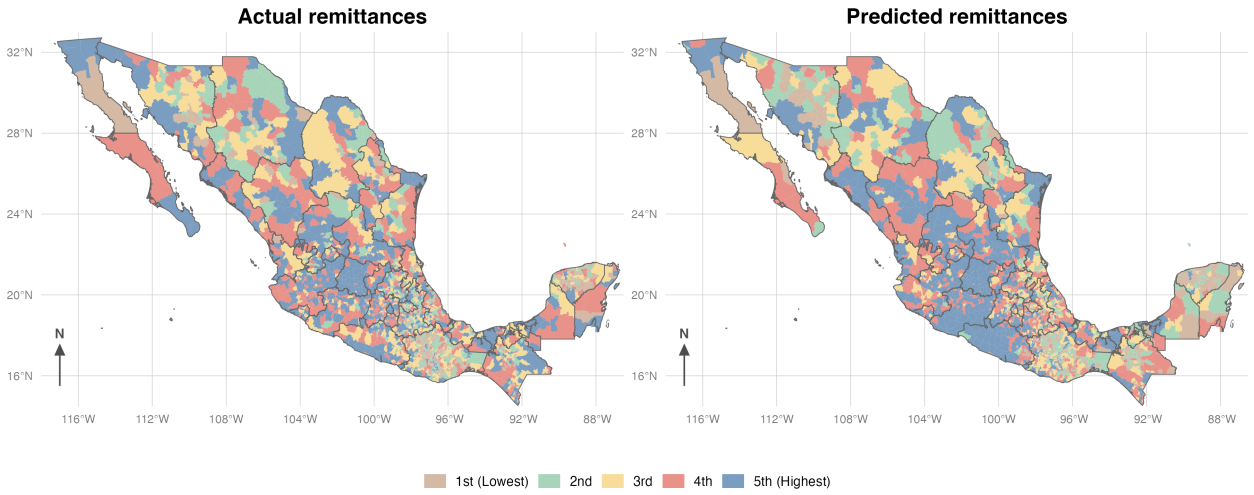
	With constant	No constant	Year FE
Predicted remittances	0.889***	0.904***	0.886***
	(0.004)	(0.004)	(0.004)
Constant	0.443***		
	(0.046)		
Observations	29,232	29,232	29,232
R^2	0.603	0.647	0.603

Notes: Dependent variable is actual BdeM municipality-level remittance inflows (quarterly mean, millions USD). Predicted remittances distribute national BdeM totals to municipalities in proportion to their 2012 MCAS migrant shares. Panel of 2,436 municipalities observed annually over 2013–2024. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

As a complementary check, Figure D.18 compares quintile maps of actual and predicted remittance distributions across municipalities, averaged over the full 2013–2024 period. The two maps display strong visual agreement: 50 percent of municipalities fall in the same quintile under both measures, and the broad spatial patterns (high remittances in western and central Mexico, low remittances in the south and southeast) are reproduced faithfully.

BdeM also publishes remittance outflows by US state of origin. This allows a state-level validation that is independent of the municipality-level exercise. Preliminary analysis of predicted vs. actual remittance outflows at the US state level yields $R^2 \approx 0.97$, providing strong additional evidence that the MCAS registry accurately captures the geographic distribution of remittance-sending migrants. A formal presentation of this exercise is left for a future version.

Figure D.18: Quintile maps of actual and predicted remittance distributions



Notes: Municipalities are classified into quintiles based on average remittance inflows over 2013–2024. Left panel: actual BdeM remittances. Right panel: predicted remittances (national total × 2012 MCAS share). Municipalities without MCAS registrations are assigned zero predicted remittances. State boundaries overlaid for reference.

E Mechanism Evidence: Temperature Effects on Migrants' Outcomes

This appendix includes additional analysis relative to mechanisms that was not included in the main paper.

E.1 Temperature and mortality in Mexico

Temperature extremes affect household welfare not only through labor markets but also through health. Mortality is a direct channel through which temperature shocks generate demand for remittances: funeral costs, loss of household income, and medical expenses preceding death all create financial needs that migrant transfers can help cover. Mortality is also a proxy for the broader burden of temperature-related morbidity. For every death attributed to temperature extremes, many more individuals experience non-fatal health shocks (hospitalizations, emergency visits, lost workdays due to illness) that impose costs on households (Sarmiento et al., 2025). The mortality estimates presented here thus represent a lower bound on the health-related demand for remittances. We use monthly data from 1998 to 2022, with mortality counts from INEGI death certificates and population weights from interpolated census data.

We estimate the effect of temperature on municipality-level mortality rates using the same panel of 2,453 Mexican municipalities and ERA5 temperature data as in our main analysis. Following Sarmiento et al. (2024), we estimate Poisson pseudo-maximum likelihood regressions of the form:

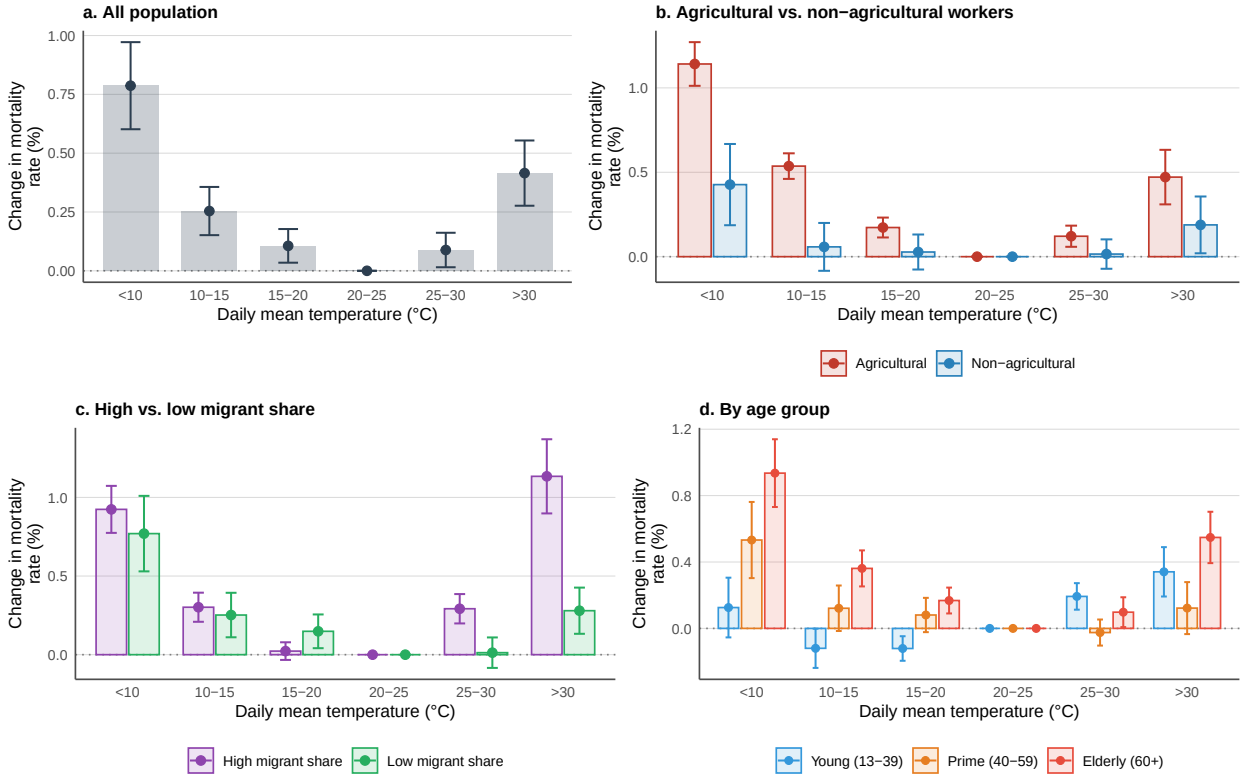
$$\text{MRate}_{mt} = \exp\left[\sum_k \beta_k \,\text{BinDays}_{k,mt} + \gamma_1 \text{Rain}_{mt} + \gamma_2 \text{Rain}_{mt}^2 + \alpha_{m\times\text{month}} + \delta_{m\times y}\right] \times \varepsilon_{mt} \tag{10}$$

where MRate_{mt} is the mortality rate per 10,000 population in municipality m and month t; $\text{BinDays}_{k,mt}$ counts days in temperature bin k (reference: 20–25°C); $\alpha_{m\times\text{month}}$ and $\delta_{m\times y}$ are municipality-by-calendar-month and municipality-by-year fixed effects. Standard errors are clustered at the municipality level.

Figure E.19 presents results for four samples. Panel (a) shows all-population mortality: temperature extremes increase mortality at both tails, with cold days (<10°C) raising mortality by 0.78% per day and hot

days (>30°C) by 0.42%. The U-shaped relationship mirrors the origin-side remittance response, consistent with health shocks generating demand for transfers. Panel (b) splits by occupation: agricultural workers exhibit a steeper response, particularly to cold, while non-agricultural worker mortality is flatter. This aligns with evidence that outdoor labor exposure drives temperature-related mortality (Sarmiento et al., 2024). Panel (c) splits municipalities at the median of the MCAS-to-population ratio. Both groups show the U-shaped pattern, but the heat effect (>30°C) is notably larger in high-migrant-share municipalities. This is consistent with the observable characteristics of migrant-sending communities: they tend to be more rural, poorer, and more marginalized, with limited access to healthcare infrastructure and cooling technologies that could buffer the health effects of extreme heat (Cattaneo and Peri, 2016b). The result reinforces the demand-side mechanism: the communities most dependent on remittances are also the most vulnerable to temperature-related health shocks. Panel (d) disaggregates by age group. The elderly (60+) exhibit the largest mortality response to both cold and heat extremes, followed by prime-age adults (40–59) and young adults (13–39). This age gradient is relevant for the remittance channel because migrants typically send transfers to support aging parents and other dependents. Temperature-induced mortality among the elderly and prime-age household members can trigger transfers to cover medical and funeral expenses, or to compensate for the loss of a household earner.

Figure E.19: Effect of temperature on mortality rates in Mexican municipalities



Notes: Each panel shows non-parametric bin estimates from equation (10). The dependent variable is the mortality rate per 10,000 population (or subgroup population for panels b and d). Reference bin: 20–25°C. Panel (a): all-cause mortality, full sample. Panel (b): agricultural workers (primary sector) vs. non-agricultural workers, with occupation-specific population denominators. Panel (c): municipalities split at the median MCAS 2012 stock per capita. Panel (d): mortality by age group (young: 13–39, prime: 40–59, elderly: 60+), with age-specific population denominators. All specifications include municipality-by-month, municipality-by-year fixed effects, and rainfall controls. Error bars: 90% confidence intervals from municipality-clustered standard errors. Monthly data, 1998–2022.

F Climate Projection Methodology: The Delta Method

This appendix describes in detail the procedure used to produce the climate change projections reported in Section 4. The pipeline consists of four steps: (i) subsetting the global climate model output, (ii) constructing a population-weighted crosswalk between the climate model grid and administrative units, (iii) computing temperature deltas at the administrative-unit level, and (iv) applying those deltas to the observed temperature record to produce projected temperature bin counts.

Climate model. We use daily mean near-surface temperature (`tas`) from CMCC-ESM2, an Earth System Model developed at the Centro Euro-Mediterraneo sui Cambiamenti Climatici (Cherchi et al., 2019). CMCC-ESM2 operates on a nominal $1.25° \times 0.94°$ longitude-latitude grid (288×192 grid cells globally) with a 365-day calendar (no leap years). We use the historical simulation (1850–2014, variant `r1i1p1f1`) and two Shared Socioeconomic Pathway scenarios from the CMIP6 ScenarioMIP: SSP2-4.5 (moderate mitigation, 2015–2100) and SSP5-8.5 (high emissions, 2015–2100). We subset each global NetCDF file to a North America bounding box (130°W–60°W, 10°N–55°N) using CDO `sellonlatbox`, reducing file sizes from approximately 1.1 GB to 40 MB each.

Population-weighted crosswalk. Climate model grid cells are coarse relative to the administrative units used in the analysis (2,453 Mexican municipalities and 3,108 US counties). A single grid cell may span multiple administrative units, and a single unit may be covered by multiple grid cells. We construct a population-weighted spatial crosswalk that maps each grid cell to overlapping administrative units, weighting by the population residing in the intersection.

Specifically, we (i) build polygons for each CMCC-ESM2 grid cell centered on the cell's latitude-longitude coordinates; (ii) compute the geometric intersection of these polygons with municipality and county shapefiles; (iii) for each intersection polygon, extract the total population using the UN-adjusted 1-km resolution population raster (GPWv4, 2020 vintage) via `exactextractr`; and (iv) compute population weights as $w_{gc} = \text{pop}(g \cap c) / \sum_{g'} \text{pop}(g' \cap c)$, where g indexes grid cells and c indexes administrative units. For the small number of units with zero population in all intersection polygons (e.g., uninhabited desert areas), we fall back to area-based weights. The resulting crosswalk ensures that temperature changes are weighted by where people live, so that a grid cell covering a dense urban area receives more weight than one covering uninhabited terrain. Weights sum to one within each administrative unit by construction.

Computing deltas. For each administrative unit, calendar quarter (Q1–Q4), target year $y \in \{2026, \ldots, 2100\}$, and SSP scenario, we compute:

$$\Delta_{c,q,y,s} = \bar{T}^{\text{future}}_{c,q,y,s} - \bar{T}^{\text{hist}}_{c,q}, \tag{11}$$

where $\bar{T}^{\text{hist}}_{c,q}$ is the population-weighted mean daily temperature from the historical CMCC-ESM2 simulation over 2000–2014, and $\bar{T}^{\text{future}}_{c,q,y,s}$ is the corresponding mean over a centered 20-year window $[y-10, y+9]$ (clamped to $[2015, 2100]$) from the SSP simulation. The population-weighted aggregation from grid cells to administrative units uses the crosswalk weights described above. By construction, the delta captures only the *change* in mean temperature projected by the climate model, not the model's absolute temperature level, which may be biased relative to observations.

Projecting temperature bin counts. For each (administrative unit, quarter, year, SSP) combination, we apply the delta to the full pool of observed ERA5 daily temperatures for that unit-quarter cell. Specifically, we:

1. Pool all observed daily mean temperatures for unit c in quarter q across all years in the ERA5 record (approximately 12 years $\times$ 90 days = 1,080 daily observations per unit-quarter cell).

2. Shift each daily temperature by the delta: $T_d^{\text{proj}} = T_d^{\text{obs}} + \Delta_{c,q,y,s}$.
3. Count the fraction of shifted days falling in each temperature bin using the same bin boundaries as the main analysis (left-closed intervals).
4. Scale the fractions to the number of days in the quarter (90, 91, or 92) to obtain projected bin day counts.

This procedure preserves the observed day-to-day variability within each quarter while shifting the entire distribution by the amount projected by the climate model. It is equivalent to assuming that climate change affects the mean of the temperature distribution but not its higher moments (variance, skewness, kurtosis). The projected bin counts are then differenced from the observed baseline bin counts (the average across all years of ERA5 data) to obtain $\Delta D^{MX}_{b,c,q,y,s}$ for Mexican municipalities and $\Delta D^{US}_{b,c,q,y,s}$ for US counties.

Projection equations. The projected percentage change in remittances for municipality m in year y under scenario s is computed separately for each channel. On the origin side:

$$\%\Delta Rem^{D}_{mys} = \exp\left(\sum_b \hat{\lambda}_b \, \Delta D^{MX}_{bmys}\right) - 1, \tag{12}$$

where $\hat{\lambda}_b$ are the estimated origin-side bin coefficients from Table B.2 and ΔD^{MX}_{bmys} is the change in the number of days per quarter in Mexican temperature bin b relative to the observed baseline. On the destination side, we first compute the change in Bartik-weighted US temperature exposure:

$$\Delta Z^{(b)}_{mys} = \sum_c s_{mc} \, \Delta D^{US}_{bcys}, \tag{13}$$

where s_{mc} are the 2012 MCAS migrant shares and ΔD^{US}_{bcys} is the change in days in US bin b for county c. The projected destination-side effect is then:

$$\%\Delta Rem^{S}_{mys} = \exp\left(\sum_b \hat{\beta}_b \, \Delta Z^{(b)}_{mys}\right) - 1, \tag{14}$$

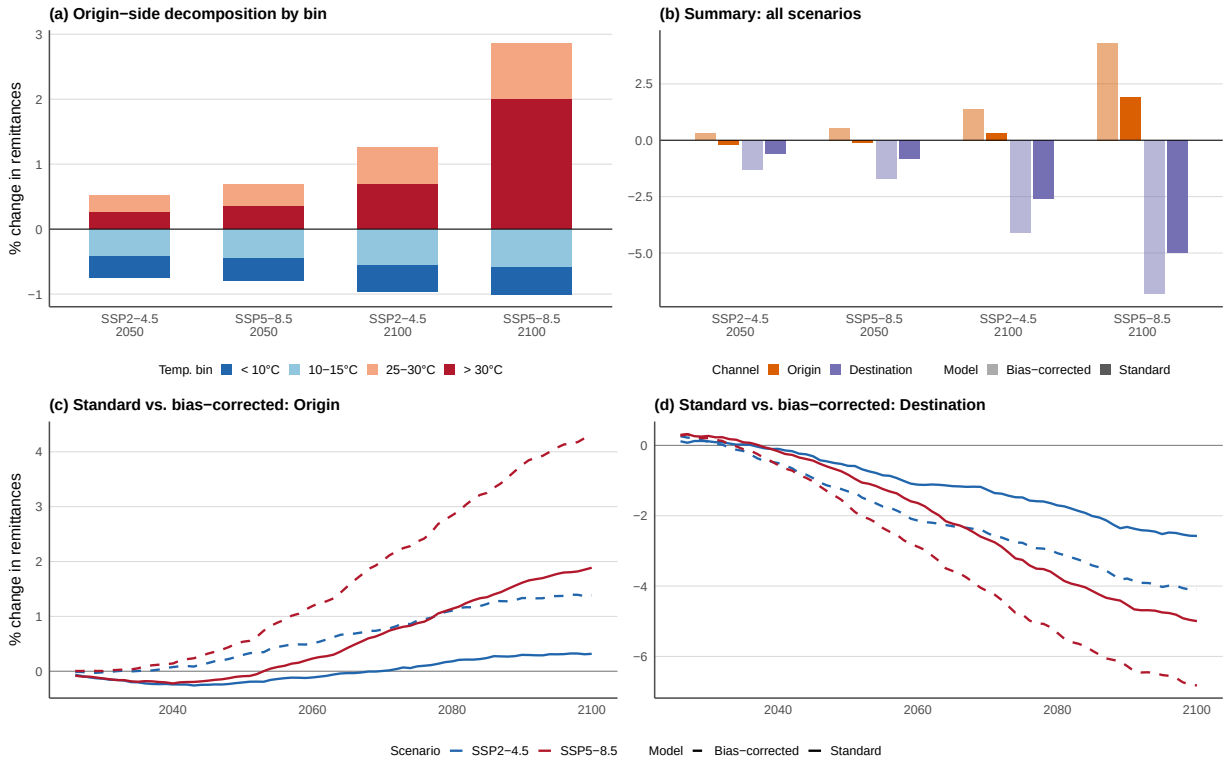
where $\hat{\beta}_b$ are the estimated destination-side bin coefficients. Following Borusyak et al. (2022), we apply the coefficients estimated on recentered exposure ($\hat{Z}$) to the total change in exposure (ΔZ), since the recentering identifies the causal effect of the underlying shocks.

Verification. We verify the pipeline through several internal consistency checks. First, the projected bin day counts sum to the expected number of days per quarter (90–92) for every unit-quarter-year-scenario combination. Second, under both SSP scenarios, hot-temperature bins (e.g., >30°C for Mexico, >30°C for the US) increase monotonically over time while cold bins decrease, consistent with warming. Third, the SSP5-8.5 deltas are uniformly larger than SSP2-4.5 deltas, as expected given the higher forcing pathway. Fourth, mid-century deltas are in the range of 1–3°C and end-of-century deltas reach 3–5°C under SSP5-8.5, consistent with IPCC AR6 projections for North America (IPCC, 2021).

G Network Diversification and Insurance

The destination-side effect of temperature on remittances operates through US counties where Mexican migrants work. Each municipality's exposure is a weighted average of county-level shocks, with weights given by 2012 MCAS migrant shares. A natural question is whether the structure of these migrant networks

Figure F.20: Additional climate projection results



Notes: Panel (a) decomposes the projected origin-side effect into contributions from each non-reference temperature bin (standard estimates): cold bins (<10°C, 10–15°C) contribute negatively as warming reduces cold days, while warm bins (25–30°C, >30°C) contribute positively. Panel (b) summarizes the projected percentage change in remittances for mid-century (2050) and end-of-century (2100) under SSP2-4.5 and SSP5-8.5, separately for the origin side (orange) and destination side (purple); darker bars use standard estimates, lighter bars use bias-corrected estimates. Panels (c) and (d) overlay the projected effects using the standard baseline estimates (solid lines) and the bias-corrected estimates from the Jones et al. (2026) correction (dashed lines) for the origin and destination sides, respectively.

matters: do municipalities with migrants spread across climatically uncorrelated destinations experience smoother remittance effects?

Using the non-parametric bin coefficients $\hat{\beta}_b$ from the destination-side baseline (Table B.2), we compute the predicted county-level effect for each county c in quarter t as:

$$f_{ct} = \sum_b \hat{\beta}_b \cdot \text{days}_{b,ct}, \tag{15}$$

where $\text{days}_{b,ct}$ counts the number of days county c spent in temperature bin b during quarter t. Because the Bartik instrument decomposes linearly across bins, the municipality-level predicted effect is a share-weighted average of county-level effects:

$$\hat{y}_{mt} = \sum_c s_{mc} \cdot f_{ct}, \tag{16}$$

where s_{mc} is the share of municipality m's migrants residing in county c. This decomposition allows us to trace municipality-level variation back to the temperature realizations at individual destinations.

The relevant diversification concept is not geographic spread or climate-mean heterogeneity, but whether destination counties experience correlated temperature shocks. Houston and Dallas have different mean climates, but both sit in Texas and tend to experience similar heatwaves. Houston and Chicago, by contrast,

may have less different mean temperatures but substantially less correlated quarterly shocks.

We measure this by computing the pairwise correlation of county-level predicted effects f_{ct} across quarters. Let R denote the correlation matrix, where $R_{cc'} = \text{Cor}(f_{ct}, f_{c't})$ computed over all 48 quarters (2013–2024). For each municipality, we construct a share-weighted average pairwise shock correlation:

$$\rho_m = \frac{\mathbf{s}_m' R \, \mathbf{s}_m - \text{HHI}_m}{1 - \text{HHI}_m}, \tag{17}$$

where $\mathbf{s}_m$ is the vector of migrant shares and $\text{HHI}_m = \sum_c s_{mc}^2$ is the Herfindahl index. The numerator $\mathbf{s}_m' R \, \mathbf{s}_m$ includes both diagonal terms ($R_{cc} = 1$, weighted by s_c^2) and off-diagonal pairwise correlations. Subtracting HHI_m removes the diagonal contribution, and dividing by $1 - \text{HHI}_m$ normalizes the measure to $[-1, 1]$.

When $\rho_m \approx 1$, all destinations experience similar shocks simultaneously, offering no insurance. When $\rho_m \approx 0$, shocks are uncorrelated and the share-weighted average is smoother. Values below zero indicate partially offsetting shocks. In the data, the median ρ_m is 0.32, with a range from -0.76 to 0.98.

Figure G.21 presents four panels. To separate the effect of network correlation from the effect of average destination climate, panels (b) and (c) residualize the outcome variables on a quartic polynomial of the municipality's mean Bartik temperature $\bar{Z}_m = \sum_c s_{mc} \bar{T}_c$ (Frisch-Waugh decomposition).

Panel (a) plots the density of time-averaged predicted effects $\bar{y}_m$ separately for terciles of ρ_m. The three distributions are similar in location (means of -7.6%, -7.5%, and -8.1%), but the high-correlation group exhibits a wider, more left-skewed distribution, consistent with less diversification of shocks.

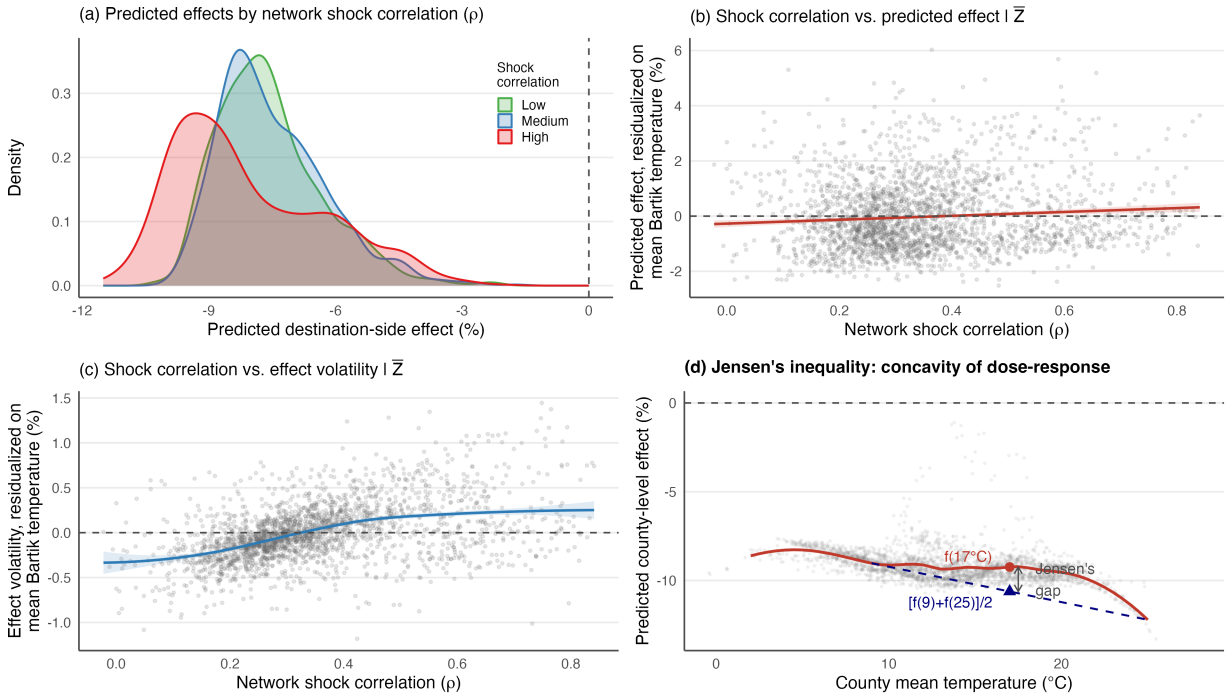
Panel (b) shows that, after conditioning on mean Bartik temperature, shock correlation has no meaningful relationship with the level of predicted effects. A linear fit (shown in the figure) is nearly flat, with a residualized correlation of 0.09. This is the expected result: ρ_m measures how shocks co-move across destinations, not the average level of temperature exposure. Municipalities with correlated and uncorrelated destinations lose similar amounts on average, holding $\bar{Z}$ constant. The contrast between the flat relationship in panel (b) and the steep gradient in panel (c) confirms that the insurance channel operates through volatility, not levels.

Panel (c) presents the main insurance result. There is a strong positive relationship between ρ_m and effect volatility (the time-series standard deviation of $\hat{y}_{mt}$), with a residualized correlation of 0.44. By tercile, volatility increases monotonically from 1.9 percentage points (low ρ) to 2.2 pp (medium) to 3.1 pp (high). Municipalities with uncorrelated destinations experience approximately 40% less quarter-to-quarter variation in their predicted destination-side effects.

Panel (d) illustrates Jensen's inequality on the county-level dose-response. Because the inverted-U relationship between temperature and remittances is concave near its peak ($T^* \approx 17$°C), spreading exposure across counties at different temperatures yields a worse average outcome than concentrating at the mean: $\frac{1}{2}[f(T-d) + f(T+d)] < f(T)$ when f is concave. At $T = 17$°C with a spread of ± 8°C, the Jensen's gap is -1.4 pp. On the convex portion of the curve (far from T^*), diversification provides a small benefit ($+0.2$ pp at $T = 10$°C). This channel operates on the level of predicted effects, complementing the volatility channel captured by ρ_m.

These results characterize the temperature-driven component of remittance variation. Temperature, however, accounts for a small share of overall remittance volatility: the standard deviation of predicted effects is approximately 2 percentage points, compared with 73 pp for residual log remittances after absorbing municipality and year-quarter fixed effects. Regressing actual remittance volatility on ρ_m confirms that the insurance channel does not detectably affect total remittance variation.

Figure G.21: Network diversification and destination-side temperature insurance



Notes: Panel (a) shows the density of time-averaged predicted destination-side effects by tercile of network shock correlation ρ_m (Equation 17). Panel (b) plots residualized predicted effects against ρ_m with a linear fit, after partialling out a quartic in mean Bartik temperature. Panel (c) plots residualized effect volatility (time-series SD) against ρ_m with a loess smoother. Panel (d) illustrates Jensen's inequality on the county-level dose-response: the red curve is a loess fit of county mean predicted effects against county mean temperature; the dashed chord connects $f(9°\text{C})$ and $f(25°\text{C})$, showing that the diversified average (blue triangle) falls below the concentrated value (red circle) due to the concavity of the dose-response near its peak. Predicted effects use the non-parametric bin coefficients from the destination-side baseline specification. Panels (b) and (c) trim ρ_m to the 1st–99th percentile range. Shaded bands represent 95% confidence intervals. The sample includes 2,430 municipalities with at least two destination counties in the 2012 MCAS.

Pre-residualized county temperature shocks

A natural concern with the shift-share design is that the identifying variation may partly reflect predictable seasonal components of county-level temperature rather than genuinely idiosyncratic weather shocks. To address this, we implement a stricter variant of the destination-side specification that isolates "unexpected" temperature shocks. Before the standard Borusyak recentering (Eq. 4), we first residualize each county-level temperature variable against county-by-year and county-by-quarter fixed effects:

$$g_{ct} = \gamma_{c,y} + \delta_{c,q} + g^*_{ct},$$

where $\gamma_{c,y}$ and $\delta_{c,q}$ absorb county-specific annual and seasonal patterns, so that g^*_{ct} captures only the deviation of county c's temperature in quarter t from its own historical norm for that time of year. We then apply the standard recentering ($g^*_{ct} = \alpha_{mt} + \tilde{g}^*_{ct}$), collapse with migrant shares, and estimate the same PPML specification as in Eq. 6.

Figure G.22 reports the results. Point estimates are numerically identical to the baseline across both the quartic polynomial (panel a) and non-parametric bin (panel b) specifications. This equivalence is not a coincidence: it follows from an algebraic absorption argument. The pre-residualization shifts the collapsed

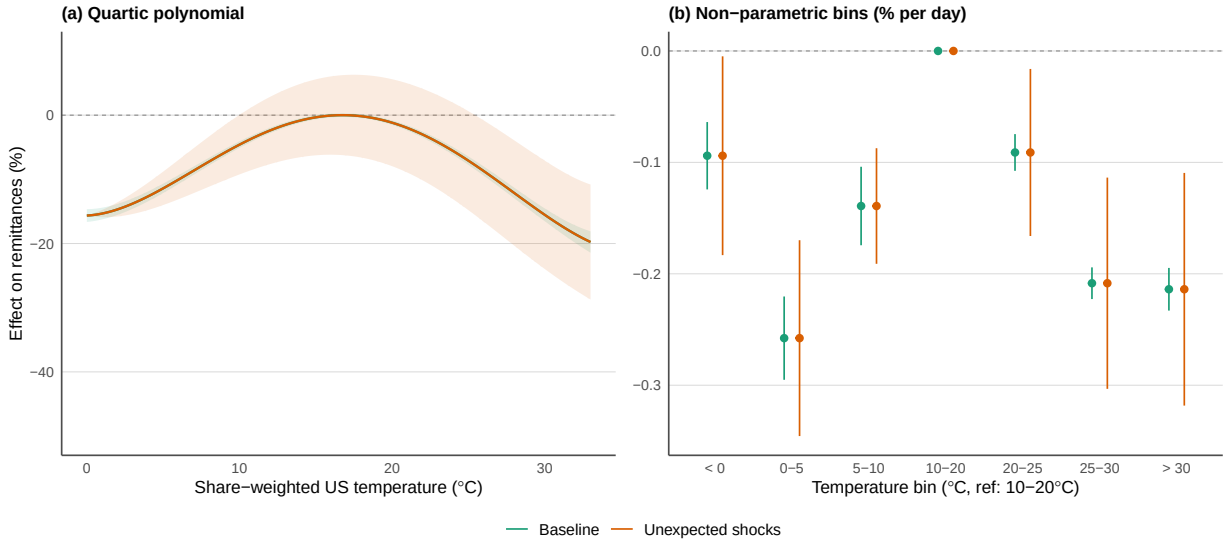
instrument by

$$\tilde{Z}^{*}_{mt} = \tilde{Z}_{mt} - \underbrace{\left[\sum_{c} s_{mc}\,\gamma_{c,y} - \bar{\gamma}_{m,y}\right]}_{h_{m,y}} - \underbrace{\left[\sum_{c} s_{mc}\,\delta_{c,q} - \bar{\delta}_{m,q}\right]}_{k_{m,q}},$$

where the bars denote unweighted network means. The correction terms $h_{m,y}$ and $k_{m,q}$ vary only at the municipality-by-year and municipality-by-quarter level, and are therefore perfectly absorbed by the fixed effects $\delta_{m\times y}$ and $\delta_{m\times q}$ already included in the PPML specification. Consequently, the pre-residualization is numerically redundant for point estimation: the baseline recentering combined with the regression fixed effects already purges the predictable county-level seasonal and annual components.

The only difference is in the standard errors. The AKM variance estimator clusters score contributions at the county-by-quarter level. Pre-residualizing the shocks reduces the within-cluster variation available for inference, widening the estimated standard errors by a factor of approximately two to four. Despite this precision loss, all six non-reference temperature bins in the non-parametric specification remain statistically significant (four at the 1% level, two at the 5% level), and the quartic coefficients $\hat{\beta}_2$ through $\hat{\beta}_4$ retain significance at conventional levels. The inverted-U dose-response relationship is fully robust to restricting attention to unexpected county-level temperature shocks.

Figure G.22: Destination-side robustness: pre-residualized county temperature shocks



Notes: This figure compares the baseline destination-side estimates (green) with a variant that pre-residualizes each county-level temperature variable by county×year and county×quarter fixed effects before applying the Borusyak recentering (orange). Panel (a) shows quartic polynomial dose-response curves normalized at their interior maximum; shaded areas denote pointwise 95% confidence bands. Panel (b) shows non-parametric bin estimates (reference: 10–20°C); whiskers denote 95% confidence intervals. Point estimates are numerically identical across the two approaches (the curves overlap perfectly in panel a and the dots align in panel b); only the standard errors differ, widening by a factor of 2–4 due to reduced within-cluster variation in the AKM variance estimator. All specifications include the full set of weather controls, Mexican temperature controls, and municipality×year and municipality×quarter fixed effects. PPML with population offset; AKM standard errors.