Federal Reserve Bank of Dallas

# Subnational Bilateral Remittance Flows in the U.S.-Mexico Corridor: A Multi-Source Imputation and Validation

Miguel Flores and Jesus Cañas

**Working Paper 2621**

**August 2026**

Research Department

https://doi.org/10.24149/wp2621

# Subnational Bilateral Remittance Flows in the U.S.-Mexico Corridor: A Multi-Source Imputation and Validation*

Miguel Flores† and Jesus Cañas‡

July 10, 2026

**Abstract**

More than $65 billion annually flows through Mexico's remittance corridor, the largest such pathway in the Western Hemisphere. Yet, the distribution of these flows across regions and the implications for financial inclusion—or alternatively, perpetuating inequality—remain poorly understood. This paper exploits bilateral remittance flows between U.S. and Mexican states during 2013–25 to construct and validate a multi-source origin–destination framework linking Banco de México's U.S.-state outflows with Mexican state-level receipts. We develop allocation matrices from consular registrations, household surveys, and transaction records and validate them against official state-level inflows. Validation against official Mexican state-level receipts yields a clear hierarchy, suggesting that matrices based on consular records explain more than four-fifths of cross-sectional variance in official receipts, outperforming survey-based and transaction-based alternatives, as the matrices preserve the concentrated geography of the settled diaspora. Using the preferred specification, we recover the bilateral structure of the U.S.–Mexico remittance system and analyze its evolution during a period that includes the COVID-19 shock. At the national level, growth in remittance inflows is largely explained by scale effects. However, at the regional level, the South-Southeast is the only region where remittance growth exceeds scale predictions. Shifts in migrant composition drive this divergence, particularly an increasing share of migrants from Southern-origin states in the United States, as well as the rising importance of specific origin states such as Oaxaca. These results highlight the importance of accounting for bilateral flow structures when studying remittances. By preserving spatial concentration, our framework provides new insights into how external shocks transmit unevenly across regions and how remittances shape subnational development trajectories in Mexico.

**Keywords:** remittances, bilateral imputation, subnational flows, Mexico, migration corridors, Matrículas Consulares, ENADID, matrix entropy, COVID-19

**JEL Codes:** F24, J61, O15, O17, R23

---

*The views expressed in this paper do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Miguel Flores, Departamento de Economía, Tecnológico de Monterrey, Campus Monterrey, miguelflores@tec.mx.

‡Jesus Cañas, Federal Reserve Bank of Dallas.

## 1. Introduction

Mexico received more than $65 billion (USD) in remittances in 2024, surpassing foreign direct investment and oil revenues as the country's principal source of foreign exchange (Banco de México, 2025; World Bank, 2022). For millions of households, concentrated in rural and semi-urban communities across the country's diverse regions, these flows finance consumption, education, housing, and small-scale investment, while providing insurance against adverse income shocks (Adams and Page, 2005; Acosta et al., 2008; Yang, 2011). Yet the central development question raised by this scale is not how much Mexico receives in aggregate, but who receives it, from where it flows, and how those origins shape the economic circumstances under which remittance income arrives. Recent data suggest a turning point in 2025. Remittance inflows fell to $61.8 billion, a 4.6 percent 12-month contraction that interrupted 11 consecutive years of expansion (BBVA, 2026). Understanding which Mexican regions are most exposed to such shifts and why they occur requires knowledge of the bilateral structure of the corridor that aggregate totals cannot provide.

The bilateral structure matters for development because Mexican states differ not only in the level of remittance receipts received, but also in the U.S. labor markets to which those receipts are tied. Exposure to California's agricultural cycle, Texas's construction sector, or emerging employment hubs in the Midwest and Southeast implies distinct transmission channels from U.S. economic conditions to Mexican regional outcomes. At the same time, migrants in these corridors face heterogeneous state-level policy environments in the United States, including differences in immigration enforcement, access to driver's licenses and local identification, and the willingness of state and local governments to cooperate with U.S. federal authorities. These regulatory and enforcement regimes shape migrants' labor market opportunities, earnings volatility, and risk of displacement, and thus condition how shocks in destination states translate into remittance flows and, ultimately, into income trajectories across Mexican regions

It also matters because the welfare impact of remittance income varies systematically with local economic conditions. South-Southeast Mexico, the country's poorest region, exhibits lower banking penetration, weaker access to formal credit, and higher rates of rural poverty than the traditional western-central sending region (CONEVAL, 2022; Demirgüç-Kunt et al., 2018). Consumption, investment, and financial inclusion effects from addition remittance income flowing into these state is therefore likely to differ from the impacts of comparable inflows to more-developed receiving regions (Giuliano and Ruiz-Arranz, 2009). Tracking the bilateral geography of remittances is thus not merely a measurement exercise but a necessary step toward understanding the uneven spatial incidence of migrant income and the distributional consequences of remittance growth.

This paper addresses that gap. No official dataset links U.S.-state remittance outflows to Mexican-state receipts at the bilateral level. Banco de México reports each side separately but does not identify the origin-destination connections between them. Private transaction records contain both endpoints but are proprietary and capture only formal transfers through specific institutions. As a result, most empirical work on the development impact of remittances in Mexico relies on unilateral measures, observing where money is received but not where it originates, which limits both the precision of causal estimates and the ability to trace how external shocks propagate across the corridor network.

We overcome this limitation by constructing and validating five bilateral allocation matrices that distribute official U.S.-state remittance outflows across Mexico's 32 states using migrant-origin shares as proxies for remittance shares. These matrices are derived from three source types: household emigration surveys (ENADID 2018 and 2023), administrative consular records (Matrículas Consulares 2018 and 2022), and transaction-based wire-transfer data (2021–2022). We validate all five against more than a decade of official Banco de México state-level receipts, providing a systematic empirical test of the corridor-proportionality approach — the assumption that, within a given U.S. state, the share of remittances sent to each Mexican destination is proportional to the origin composition of the resident Mexican-born population (Hanson, 2010; Docquier, Rapoport and Salomone, 2012; Clemens and McKenzie, 2018).

The validation results establish a clear hierarchy with direct implications for measurement practices. Matrices based on consular records systematically outperform survey-based and transaction-based alternatives, explaining more than four-fifths of the variance in official receipts over more than a decade. The key mechanism is geographic concentration, as matrices that preserve the concentrated settlement structure of the Mexican diaspora recover remittance geography more accurately than matrices that spread flows too broadly. Crucially, transaction data often presumed to carry the most direct information about remittance flows, perform substantially worse, because they reflect channel-specific remitting behavior shaped by financial inclusion and provider network availability rather than the underlying geography of the migrant population (Ambrosius and Cuecuecha, 2016; Clemens and McKenzie, 2018). This finding carries practical implications for statistical agencies and central banks seeking to estimate bilateral remittance flows in other corridors: Publicly available population-based data may be more informative than proprietary transaction records for identifying geographic structure.

Using the preferred matrix, we recover the full bilateral structure of the U.S.–Mexico remittance system and trace its evolution across a period that includes the COVID-19 shock. Three findings emerge. First, high-volume corridors linking California and Texas to traditional western-central Mexican states persistently anchor the system, consistent with the durable role of migration networks in shaping remittance geography. Second, a distinct growth frontier has emerged in non-traditional U.S. sending states, particularly in the Midwest, Mid-Atlantic, and Southeast, whose Mexican-born populations are disproportionately connected to South-Southeast Mexico. Third, and most consequential for regional development, a corridor-level Oaxaca–Blinder decomposition shows that while the COVID-era remittance surge is entirely scale-driven at the national level, it was not distributionally neutral. The South-Southeast Mexico is the only region to record a positive structural contribution, capturing a disproportionate share of growth through the reweighting of diaspora geography rather than through aggregate expansion alone. The West and North receive systematically less than their pre-pandemic network position would predict. These results provide direct evidence of an ongoing southward redistribution of remittance income toward Mexico's poorest region, a redistribution with development implications that cannot be observed in aggregate or in unilateral data.

This paper contributes to the literature in three ways. First, it provides a validated framework for estimating bilateral remittance flows in the absence of direct corridor-level data, showing that population-based matrices can depict subnational remittance geography with substantial accuracy. Second, it contributes to the literature on migration networks by documenting the coexistence of persistence and structural change. The remittance corridor retains a stable historical backbone, yet meaningful redistribution occurs at the margin through the rise of new origin-destination links, consistent with path-dependent network formation and gradual corridor consolidation (Massey, Durand and Malone, 2002; Munshi, 2003; Garip, 2012; Beine, Docquier and Özden, 2015). Third, it shows that aggregate remittance statistics can mask important spatial dynamics, with direct implications for how U.S. regional shocks are transmitted across Mexican states and how the gains from remittance growth are distributed within migrant-sending countries.

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature on remittances, migration corridors, bilateral flow estimation, and COVID-era migration restructuring. Section 3 describes the data sources, the construction of the allocation matrices, and the validation framework. Section 4 presents the validation results. Section 5 analyzes the structure and dynamics of the U.S.–Mexico remittance corridor, including a decomposition of COVID-era changes into scale and structural components. Section 6 concludes.

## 2. Literature Review

### 2.1 Remittances, Migration Corridors, and Development

A large literature documents the economic importance of remittances for migrant-sending households and regions. At the household level, remittances finance consumption, schooling, housing, and small-scale investment, while also providing insurance against income shocks and macroeconomic instability (Yang, 2008, 2011). At the aggregate level, remittance inflows have been linked to lower poverty and improved human capital outcomes, particularly where households face binding liquidity constraints and formal financial systems are shallow (Adams and Page, 2005; Acosta et al., 2008; Giuliano and Ruiz-Arranz, 2009). For Mexico, existing evidence shows that remittances can support microenterprise formation and generate local multiplier effects in high-migration communities (Durand, Parrado and Massey, 1996; Woodruff and Zenteno, 2007).

Yet most of this literature relies on unilateral or destination-only measures of remittances. Mexican states differ not only in the level of remittance receipts received, but also in the U.S. labor markets to which those receipts are tied. Exposure to California's agricultural cycle, Texas's construction sector, or New York's service economy implies distinct transmission channels from U.S. shocks to Mexican local outcomes. Without bilateral corridor information, empirical work must treat remittance receipts as if a common external process generated them, even though the origin structure may differ sharply across recipient states. This matters because the development impact of remittances depends not only on how much is received, but also from where those flows come and how sensitive those origin areas are to labor-market, policy, or macroeconomic shocks.

This bilateral dimension is especially important in a country such as Mexico, where migration networks are highly region-specific and where the effects of remittance inflows are likely to vary with local poverty and financial development. The evidence that remittances generate stronger developmental effects in financially constrained settings suggests that a redistribution of remittance growth toward poorer Southern states may have different welfare implications than comparable growth in the traditional western-central sending region (Giuliano and Ruiz-Arranz, 2009; Demirgüç-Kunt et al., 2018). Identifying the bilateral geography of remittances is therefore not merely a measurement exercise. It is also a necessary step toward understanding the uneven spatial transmission of migration-based income.

The theoretical foundation for this paper lies in the literature on migration networks. A large body of work shows that migrant flows are strongly path-dependent: Once a community establishes a foothold in a particular destination, social and informational ties reduce migration costs and sustain subsequent movement along the same corridor (Massey, 1987; Massey, Durand and Malone, 2002; Munshi, 2003; Garip, 2012). These networks also shape remittance patterns.

Migrants typically remit to family members and communities of origin, so remittance behavior is closely tied to the same corridor structure that governs migration itself (Rapoport and Docquier, 2006; Yang, 2011). This empirical regularity, that migrants remit primarily to where they are from, not merely from where they live, underpins the corridor-proportionality assumption in this paper.

The relevance of migration networks is not confined to the U.S.–Mexico case. Using global bilateral data, Beine, Docquier and Özden (2015) show that network externalities in migration are highly corridor-specific and strongly self-reinforcing. This evidence supports the view that the strongest remittance corridors should coincide with the densest and most persistent migrant networks. The contribution of the present paper is to bring this network perspective to the measurement of subnational remittance flows and to test directly whether alternative data sources differ in their ability to recover the concentrated geography implied by migration networks.

### 2.2 Bilateral Flow Estimation: Methods and Challenges

The challenge of estimating bilateral flows in the absence of direct measurement arises across several economic fields. In trade, gravity models infer bilateral flows from distance, common language, and other frictions (Anderson and van Wincoop, 2003). In international finance, benchmark surveys recover bilateral asset positions through direct reporting. Remittances pose a more difficult problem. Direct bilateral measurement is rare, especially below the national level, and formal transaction records are often proprietary, institution-specific, or incomplete with respect to informal channels. As a result, the dominant strategy in remittance measurement has been to infer bilateral flows from migrant stocks.

The core identifying idea is straightforward. If migrants from different Mexican states living in the same U.S. state remit similar amounts on average, then the origin composition of the migrant stock in that U.S. state can proxy for the destination composition of its remittance outflows. Hanson (2010) describes this logic in the context of developing-country remittances and notes that the assumption is plausible when migration and remitting behavior are both embedded in stable origin-based networks. At the international level, Docquier, Rapoport and Salomone (2012) show that bilateral migrant stocks are strong predictors of bilateral remittance flows, lending broader support to stock-based imputation methods. Related work for the U.S.–Mexico corridor shows that state-to-state remittance estimates based on migrant-origin shares explain a substantial fraction of the variation in Banco de México's state-level receipts and respond systematically to U.S. labor-market conditions (Flores and Gutiérrez Bernal, 2025; Flores and Torre, 2024).

Despite this progress, three issues remain unresolved. First, stock-based imputation has rarely been validated systematically across different types of source data. Household surveys, administrative registries, and transaction records each capture different segments of the migrant population and may therefore imply different bilateral structures. Second, even when transaction records appear to offer direct evidence regarding remittance flows, they may reflect only formal transfers processed through a particular institution rather than the true geography of aggregate remittance volumes. Third, most existing exercises operate either at the national bilateral level or for a single cross-section, leaving open the question of whether a fixed bilateral structure can explain receiving-side outcomes over extended time periods.

These concerns are particularly salient because formal and informal channel use varies systematically across migrant groups, destinations, and levels of financial inclusion. Ambrosius and Cuecuecha (2016) show that remittance behavior depends on access to financial services and channel-specific characteristics, implying that transaction data may capture the choice of how to remit rather than the structural geography of who remits to whom. Clemens and McKenzie (2018) identify bilateral remittance accounting as one of the major unresolved measurement problems in development economics and warn against treating partial transaction data as if they were a complete representation of remittance geography. The present paper contributes to this literature by offering a formal multi-source validation exercise and by showing that the source closest to actual transactions is not necessarily the one that best recovers aggregate geographic structure.

For the U.S.–Mexico corridor specifically, the literature remains constrained by the absence of direct bilateral subnational data. Existing work on remittances in Mexico typically studies state-level receipts, household welfare, or local labor-market responses without observing the U.S. origins of those flows (Amuedo-Dorantes and Pozo, 2004; Hanson, 2010; Mora-Rivera, García-Amador and Van Gameren, 2024). On the U.S. side, descriptive work documents the state-level distribution of the Mexican-born population but does not connect those distributions to remittance volumes (Passel and Cohn, 2009). To current knowledge, no previous study has systematically compared multiple bilateral allocation matrices for the U.S.–Mexico corridor and validated them against an independent state-level remittance panel spanning more than a decade. This paper addresses that empirical gap.

### 2.3 COVID-19, Migration Restructuring, and the Southward Shift

The bilateral framework developed in this paper is motivated not only by a measurement problem but also by a historical episode that exposed the limitations of aggregate remittance data. The period studied here, 2013–25, includes the COVID-19 pandemic, the most important shock over this interval to affect labor markets, migration dynamics, and remittance behavior along the U.S.–Mexico corridor. Whether the pandemic-era remittance boom simply scaled up an existing corridor structure or also altered the geography of remittance allocation is a central question of the present study.

The pandemic's effect on remittances was both large and, initially, counterintuitive. Despite a severe contraction in the U.S. economy, remittance flows from the United States to Mexico rose sharply between 2018 and 2021. Several mechanisms have been proposed to explain this pattern. One posits that U.S. fiscal transfers and labor-market support disproportionately benefited migrant workers who remained employed in essential sectors, sustaining their capacity to remit (World Bank, 2020; Genoni, Muñoz Boudet and Sahay, 2021). Another is that remittance intensity rose in response to the sharp deterioration experienced by recipient households in Mexico, consistent with the insurance motive emphasized in the remittances literature (Yang, 2008, 2011; Rapoport and Docquier, 2006). A third is compositional selection: Migrants with weaker labor-market attachment may have been more likely to return, leaving behind a more stably employed remitting population (Banco de México, 2022; World Bank, 2022).

At the same time, the pandemic period coincided with broader changes in migration geography. Recent ENADID rounds and Banco de México data confirm that this southward reweighting has been strong and sustained over the past decade, rather than a transient COVID-specific anomaly (INEGI, 2019; INEGI, 2024; Banco de México, 2025). This pattern is consistent with the logic of network formation. Once migrants from Southern states establish themselves in newer U.S. destinations, especially in the Southeast and parts of the Mid-Atlantic, those corridors become self-reinforcing through the same mechanisms that historically sustained the classic western-central routes to California and Texas (Massey, Durand and Malone, 2002; Munshi, 2003; Garip, 2012). What changes is not the existence of network persistence, but the geographic composition of the networks that persist.[1]

[1] Comparing the two quinquennial reference windows, ENADID 2018 recorded 645,458 recent emigrants to the United States over August 2013–October 2018, whereas ENADID 2023 recorded 832,627 emigrants over August 2018–October 2023 (INEGI, 2019; INEGI, 2024). This implies an increase of about 29 percent in recent emigration. Banco de México (2025) reports that emigration from Southern Mexican states more than tripled between the two quinquennia (approximately 230 percent growth), compared with 103 percent in the central region, 59 percent in the northern interior, and 21 percent in the northern border states.

Southern Mexico differs systematically from the traditional sending region in both socioeconomic structure and financial inclusion. It exhibits higher poverty rates, weaker banking penetration, and greater dependence on informal and subsistence activities than the West or North (CONEVAL, 2022; Demirgüç-Kunt et al., 2018). As a result, a redistribution of remittance growth toward Southern destinations carries welfare implications that differ from those arising from a proportional expansion of pre-existing flows. Aggregate state-level remittance data can show that Southern states are receiving more, but they cannot identify whether this reflects faster growth in U.S. outflows from traditional origins, the consolidation of new migration corridors, or both. A bilateral framework is needed to distinguish between these possibilities.

This is the point at which the present paper enters the literature. By constructing and validating bilateral allocation matrices from multiple sources, the analysis recovers the origin-destination structure of the U.S.–Mexico remittance system and uses it to decompose post-COVID changes into scale and structural components. In doing so, it links the measurement problem of bilateral remittance estimation to the substantive question of whether the remittance boom of the pandemic period reflected the persistence of established migration networks or the reallocation of growth toward new ones. The analysis therefore speaks simultaneously to the literatures on remittance measurement, migration networks, and the regional incidence of migrant income.

## 3. Data and Methodology

### 3.1 Conceptual Framework

The empirical strategy rests on a corridor-proportionality assumption: Within a given U.S. state, the share of remittances sent to each Mexican destination is proportional to the origin composition of the Mexican-born population residing in that state. In substantive terms, migrants are assumed to remit primarily to their place of origin, so the geography of migrant stocks determines the geography of remittance flows. This assumption is consistent with a broad literature showing that migration networks are highly corridor-specific and path-dependent, and that remittance behavior is closely linked to origin-based family and community ties rather than to location alone (Massey, 1987; Munshi, 2003; Rapoport and Docquier, 2006; Yang, 2011; Garip and Asad, 2016). At the international level, bilateral migrant stocks have been shown to be strong predictors of bilateral remittance flows, lending external support to this approach (Docquier, Rapoport and Salomone, 2012).

Formally, let $\boldsymbol{R_j(t)}$denote total remittance outflows from U.S. state $\boldsymbol{j}$ at time $\boldsymbol{t}$, and let $\boldsymbol{Y_i(t)}$ denote remittance receipts in Mexican state $\boldsymbol{i}$. The objective is to recover bilateral flows from U.S. state $\boldsymbol{j}$ to Mexican state $\boldsymbol{i}$. These imputed bilateral remittance flows are defined as:

$$\widehat{\boldsymbol{R}}_{ji}(\boldsymbol{t}) = \boldsymbol{F}_{ji}(\boldsymbol{t}) = \boldsymbol{w}_{ji} \cdot \boldsymbol{R}_{j}(\boldsymbol{t}), \tag{1}$$

where $\boldsymbol{w}_{ji}$ represents the share of remittances originating in U.S. state $\boldsymbol{j}$ that are allocated to Mexican state $\boldsymbol{i}$. By construction, the allocation weights satisfy $\sum_{i} \boldsymbol{w}_{ji} = \mathbf{1} \forall \boldsymbol{j}$, ensuring that total outflows from each U.S. state are fully distributed across Mexican destinations. Aggregating across U.S. states yields total remittance receipts in Mexican states:

$$\boldsymbol{Y}_{i}(\boldsymbol{t}) = \sum_{j} \widehat{\boldsymbol{R}}_{ji}(\boldsymbol{t}) = \sum_{j} \boldsymbol{w}_{ji} \cdot \boldsymbol{R}_{j}(\boldsymbol{t}) \tag{2}$$

The key empirical task is therefore to construct the allocation weights $\boldsymbol{w}_{ji}$. These weights are derived from a bilateral migrant stock matrix $\boldsymbol{M}$, where $\boldsymbol{M}_{ij}$denotes the number of migrants born in Mexican state $\boldsymbol{i}$ and residing in U.S. state $\boldsymbol{j}$. From this matrix, two related but conceptually distinct objects can be defined.

First, the diaspora matrix $\boldsymbol{w}_{ij} = \frac{\boldsymbol{M}_{ij}}{\sum_{j'} \boldsymbol{M}_{ij'}}$ captures how migrants from each Mexican state are distributed across U.S. destinations, that is, where emigrants from a given origin state settle in the United States. Second, and crucial for remittance imputation, the remittance allocation matrix $\boldsymbol{w}_{ji} = \frac{\boldsymbol{M}_{ij}}{\sum_{i'} \boldsymbol{M}_{i'j}}$ describes the origin-state composition of migrants residing in each U.S. state—in other words, for a given U.S. state of residence, the Mexican states from which its Mexican-born population originates. Although both objects are normalizations of the same underlying bilateral migrant-stock matrix, they serve fundamentally different purposes. The diaspora matrix describes where migrants from a given Mexican state go, while the remittance allocation matrix describes who lives in a given U.S. state. For remittance imputation, the latter is the relevant construct, because total outflows from a U.S. state must be allocated across Mexican destinations according to the origin shares of that state's Mexican-born population.

Under the corridor-proportionality assumption, the expected share of remittances from U.S. state $\boldsymbol{j}$ sent to Mexican state $\boldsymbol{i}$, conditional on the observed migrant stock distribution $\mathbf{M}_{\cdot j}$ in that U.S. state, equals $\boldsymbol{w}_{ji}$:

$$\mathbb{E}\left[\frac{\hat{\boldsymbol{R}}_{ji}(\boldsymbol{t})}{\boldsymbol{R}_j(\boldsymbol{t})}\,\middle|\, \boldsymbol{M}_{\cdot j}\right] = \boldsymbol{w}_{ji}, \tag{3}$$

where equation (3) formalizes the paper's core empirical strategy, suggesting that conditional on where migrants live in the U.S., differences in bilateral remittance flows across Mexican destinations arise from differences in the composition of migrant origins, not from origin-specific remitting behavior *per-se* (e.g., migrants from Michoacán do not remit systematically more or less per person than migrants from Oaxaca within the same U.S. state). The assumption does not require all migrants to remit the same amount. Rather, it requires that within-state differences in average remittances across origin groups are not large enough to overturn the allocation implied by migrant-origin shares. Its empirical plausibility is evaluated indirectly through the validation exercise showing that if the restriction held poorly, no matrix constructed from migrant stocks would reproduce the observed distribution of state-level remittance receipts with meaningful accuracy.

### 3.2 Source Datasets

The analysis combines one official remittance source with three types of bilateral origin-destination data. The official remittance benchmark is Banco de México, which reports quarterly remittance outflows by U.S. state and remittance receipts by Mexican state. The outflow series provides the observed totals $R_j(t)$ that are allocated across Mexican destinations, while the receipt series provides the observed benchmark against which alternative imputation matrices are validated.

Three source types are used to construct the bilateral allocation matrices. The first is ENADID (Encuesta Nacional de la Dinámica Demográfica), a nationally representative Mexican household survey that includes migration modules identifying recent emigrants, their Mexican state of origin, and their U.S. state of destination. The 2018 and 2023 waves are used because they provide the bilateral state-to-state detail required for this exercise and bracket the COVID period. ENADID captures recent emigration flows rather than the settled migrant stock, which makes it informative about structural changes in corridor geography but also more vulnerable to sampling variability in low-volume corridors (Docquier et al., 2014; Flores and Gutiérrez Bernal, 2025).

The second source is the Matrícula Consular de Alta Seguridad, an administrative record generated by Mexican consulates in the United States. Applicants report their Mexican state of birth and current U.S. address, which allows registrations to be aggregated into a bilateral matrix by Mexican origin state and U.S. state of residence. The 2018 and 2022 editions are used to align with the ENADID windows and to capture pre- and post-COVID diaspora geography. Relative to household surveys, consular data provide broader coverage of the settled Mexican diaspora and are especially informative for undocumented migrants, who are a remittance-intensive segment of the migrant population and often rely on consular identification to access local services and financial institutions (Passel and Cohn, 2009; Massey and Gentsch, 2014).

The third source is transaction-based, wire-transfer records from a major Mexican financial institution (Banorte), capturing the volume-weighted origin–destination distribution of formal remittance transfers across U.S.–Mexico corridor pairs. In contrast to the population-based sources above, this dataset measures behavioral transaction frequency—how often migrants remit through formal channels—rather than where migrants reside. Transaction frequency is shaped by financial inclusion, provider network availability, transfer size preferences, and individual remitting habits, all of which are at best loosely correlated with underlying migrant population geography (Ambrosius & Cuecuecha, 2016; Clemens & McKenzie, 2018). This distinction between behavioral frequency and structural population distribution provides the theoretical basis for expecting transaction data to underperform in bilateral volume imputation—a prediction confirmed by the validation results in Section 4. Specifically, the transaction data used in this paper originate from the same Banorte wire-transfer records analyzed in a joint (Center for Latin American Monetary Studies) CEMLA–Banorte study on female Mexican emigration and remittances (CEMLA, 2024). That study focused on gender patterns in remitting behavior, financial inclusion, and women's aggregate contribution to Mexico's remittance income. The present paper uses the same underlying data, facilitated by CEMLA, for constructing a matrix to impute subnational origin–destination remittance flows. While the CEMLA–Banorte study aggregated transfers by sender and recipient gender, this paper leverages the geographic identifiers (U.S. state of origin and Mexican state of destination) embedded in aggregated data to estimate corridor-level flows and validate them against official Banco de México receipts.[2]

To ensure comparability across sources, the analysis is restricted to the 45 U.S. states that appear in ENADID 2018 and for which Banco de México reports quarterly remittance outflows. The excluded states account for a negligible share of the Mexican-born population in the United States and have no meaningful effect on aggregate imputation results.[3] The sample therefore balances

[2] The author is grateful to the Center for Latin American Monetary Studies (CEMLA) for providing the aggregated transaction data by remittance corridor used in this study. Special thanks are due to Jesús Cervantes for his generous support, guidance, and assistance in facilitating access to these data.

[3] The excluded U.S. states, Maine, New Hampshire, Vermont, Wyoming, and Connecticut, account for less than 1 percent of total remittances from the United States to Mexico in 2025, based on Banco de México data, with a combined share of 0.72–0.77 percent across quarters. Connecticut represents about 0.29 percent, while the remaining states each contribute 0.05 percent or less. These

geographic coverage with temporal and cross-source consistency. Throughout the analysis, Mexican states are also grouped into four regions following the standard CONAPO classification: North, West, Central, and South-Southeast. (CONAPO is Consejo Nacional de Población, the national population council.) This regional structure is used to organize the interpretation of bilateral dynamics and the decomposition of COVID-era changes.

### 3.3 Comparison of diaspora matrices and validation of imputed values

Before turning to the external validation against observed remittance receipts, examining the internal structure of the diaspora matrices is useful. This step clarifies how source-specific features of each matrix shape the imputed bilateral flows they generate. To compare the diaspora matrices, we use an additional set of metrics, three internal structural diagnostics computed from each matrix's row-conditional distribution of U.S-to-Mexican state allocation shares. These diagnostics summarize geographic concentration (the Herfindahl–Hirschman Index, HHI), data density (sparsity), and distributional uniformity (Shannon entropy). Taken together, they describe the underlying shape of each matrix before it is used for imputation and clarify why matrices derived from different sources systematically deliver different levels of predictive accuracy.

Geographic concentration (HHI) is measured as the Herfindahl–Hirschman Index, computed column-wise and averaged across U.S. states:

$$\mathrm{HHI} = \sum_{j=1}^{32} w_{ij}^{2} \tag{4}$$

where $w_{ij}$ is the allocation share assigned to Mexican state (i) by U.S. state (j). HHI takes values between $1/32 \approx 0.031$ (perfect uniformity) and 1.0 (complete concentration in a single destination). A higher HHI indicates that a matrix assigns most of a given U.S. state's remittance outflows to a small number of Mexican states, consistent with the empirically documented concentration of the Mexican diaspora in a handful of high-emigration states such as Michoacán, Jalisco, and Guanajuato (Massey et al., 2002). Under the corridor-proportionality assumption, a matrix that reproduces this concentration should outperform one that spreads flows too evenly, because the true geographic structure of remittance allocation is itself highly concentrated (Hanson, 2010).

figures are consistent with close to 1 percent share of the Mexican-born population in the United States (ACS). Their omission therefore has a negligible effect on aggregate flows but avoids sparsity and comparability issues in the validation exercises.

Sparsity is defined as the share of zero-valued cells in the allocation matrix, expressed as a percentage of all possible U.S.–Mexico corridor pairs. A sparse matrix contains many bilateral pairs with no observed migrant connection, either because the underlying data source lacks sufficient sample coverage to populate low-volume corridors or because no meaningful activity occurs along those routes, as in many international migration matrices (Beine, Docquier, & Özden, 2015). The relationship between sparsity and predictive accuracy is therefore ambiguous. Denser matrices can, in principle, allocate flows to a broader set of destinations, but density that arises from behavioral transaction data rather than genuine population geography can reduce accuracy by assigning positive weights to corridors with negligible actual migrant presence, a concern echoed in work comparing survey-based and transaction-based remittance measures (Clemens & McKenzie, 2018).

We also use Shannon entropy is used as a complementary measure of dispersion. It is computed row-wise for each U.S. state and then averaged across Mexican states:

$$\boldsymbol{H_i = -\sum_j w_{ji} \log_2(w_{ji})} \tag{6}$$

$$\boldsymbol{H = \frac{1}{J}\sum_j H_j} \tag{7}$$

where $\boldsymbol{w_{ji}}$is the remittance allocation share from U.S. state $\boldsymbol{j}$ to Mexican state $\boldsymbol{i}$. Entropy measures the effective number of destinations across which a state's outflows are distributed. A fully concentrated allocation yields $\boldsymbol{H = 0}$, while a uniform distribution across all 32 Mexican states yields $\boldsymbol{H = \log_2(32) = 5}$bits. The intuition for using entropy as a diagnostic is direct. Under the corridor-proportionality assumption, the most accurate matrix should be the one whose geographic distribution best mirrors the true concentration of the Mexican diaspora in each U.S. state, in line with evidence on path-dependent migration networks (Munshi, 2003; Garip, 2012). A matrix that is excessively diffuse, spreading flows thinly across many Mexican destinations, will tend to underestimate receipts for high concentration states such as Michoacán and Jalisco and overestimate receipts for peripheral destinations. Entropy therefore serves as a mechanism variable linking differences in source type population or administrative-based data versus transaction-based measures to differences in imputation accuracy.

### 3.4 Validation strategy

The alternative matrices are evaluated by comparing their implied remittance receipts with Banco de México's official state-level series. For each matrix, bilateral outflows are imputed at the quarterly level and then aggregated across U.S. states to obtain predicted receipts for each Mexican state. These predicted values are compared with observed receipts over the 2013–25 period.

Validation relies on multiple metrics because no single statistic captures all relevant dimensions of fit. The primary measure is [$R^2$], which indicates how much of the variation in official receipts is explained by the imputed values. Pearson's correlation coefficient is also reported because it is scale-invariant and therefore captures whether the matrix correctly ranks Mexican states by remittance receipts even when predicted levels are mechanically below official totals. This distinction matters because the bilateral matrices allocate only U.S.-origin remittances, whereas Banco de México's state-level receipts include inflows from all countries. Finally, mean absolute percentage error (MAPE) is used to summarize average proportional deviations between predicted and observed receipts.

The validation covers the full 2013–25 period and all 32 Mexican states, yielding a large state-by-quarter panel for each matrix. Because the matrices are constructed from discrete cross-sections rather than annual bilateral data, the exercise also provides a demanding test of temporal stability. If a single snapshot of diaspora geography continues to explain observed remittance receipts over more than a decade, the result would support the view that migration networks are highly persistent and that time-invariant bilateral weights are informative about the long-run structure of the corridor.

Once the preferred matrix is identified, it is used to recover bilateral remittance flows and to study the structure and dynamics of the U.S.–Mexico corridor. The analysis first examines the long-run architecture of the system under a fixed-matrix specification in which changes in U.S. state-level outflows drive all variation in corridor-level flows, while the bilateral structure remains constant. This provides a transparent benchmark for identifying the historical backbone of the remittance network and the relative importance of established versus emerging corridors.

The analysis then turns to structural change during the COVID period. To separate aggregate expansion from geographic reallocation, the analysis applies an Oaxaca-Blinder decomposition to changes in remittance receipts between the pre- and post-COVID periods. In this setting, changes in U.S. state-level outflows capture the scale effect, while changes in the allocation weights capture structural reallocation across Mexican destinations. This decomposition allows the analysis to determine whether the remittance boom of the pandemic period reflected a uniform expansion of existing corridors or a deeper redistribution of flows toward new receiving regions.

## 4. Results

### 4.1 Structural Properties, Predictive Performance and Validation

Table 1 summarizes the structural properties of the five allocation matrices together with their validation performance against Banco de México's annual state-level receipts during 2013–25. *MCAS* 2018 shows higher concentration and lower entropy, indicating that they allocate flows to a relatively small set of destinations, consistent with the known concentration of Mexican migration networks. ENADID matrices are even more concentrated but substantially sparser, reflecting smaller underlying migrant stocks and undercoverage of lower-volume corridors. The transaction-based matrix, by contrast, is extremely diffuse, with near-zero sparsity and the highest entropy, implying that flows are spread broadly across destinations.

Table 1 extends this validation framework by jointly presenting the structural properties of each allocation matrix and their corresponding fit against Banco de México's annual state-level receipts during 2013–25. Taken together with the preceding discussion, the table makes clear that predictive performance is tightly linked to how concentrated each matrix is. MCAS 2018 combines relatively high concentration and moderate sparsity with the strongest fit and it attains the highest explanatory power in the full state–year panel ($R^2 = 0.828$), and the largest Pearson correlation with official receipts ($r = 0.913$). ENADID-based matrices display somewhat lower $R^2$ and r, reflecting their more fragmented structure and the underrepresentation of low-volume corridors. By contrast, the transaction-based matrix is highly diffuse, with near-zero sparsity, the highest entropy, and much weaker validation metrics ($R^2 = 0.469$; $r = 0.685$; MAPE = 47.5 percent). This pattern anticipates the mechanism highlighted in Figure 1 that across the five sources, lower entropy (more concentrated allocations) is almost perfectly associated with higher $R^2$. Because all matrices share the same mechanical downward bias—allocating only U.S.-origin remittances while official receipts include all countries Pearson's r is the most informative summary of their usefulness for empirical work, and on this scale-invariant metric MCAS 2018 is clearly preferred.[4]

**Table 1. Table 1. Structural Properties and Validation Performance of Bilateral Allocation Matrices, 2013–2025**

[4] All matrices allocate only U.S.-origin remittances, while official receipts include also non-U.S. flows (Canada, Europe, etc.), which average 7.6–7.9% of national receipts. Because every matrix is row-stochastic, the national aggregate of any estimate exactly equals total U.S. outflows, making the value-weighted national gap purely mechanical and nearly identical across matrices (≈ −8%). Additionally, we calculated the equal-weighted mean percentage error (*Bias_EW*), which captures additional cross-state misallocation. For *MCAS* 2018, *Bias_EW* = −16.8%, about 9 points more negative than the mechanical benchmark, reflecting the concentration of non-U.S. remittances in northern and more urban states and residual error in low-volume destinations. ENADID matrices show smaller negative bias (−2.3% and −4.1%) not because they allocate more accurately, but due to their sparse structure (65–57% zero cells) as they avoid imputing for small states altogether, omitting those errors from the average. The transaction matrix exhibits slight positive bias (+8.0%), consistent with its overrepresentation of formal transfer channels.

| Source | HHI | Entropy (bits) | Sparsity (%) | Non-zero corridors | $R^2$ | Pearson r | MAPE (%) |
|---|---|---|---|---|---|---|---|
| *MCAS* 2018 | 0.094 | 3.97 | 9.9% | 1,299 | 0.828 | 0.913 | 28.3 |
| *MCAS* 2022 | 0.101 | 3.88 | 9.9% | 1,442 | 0.783 | 0.885 | 31.9 |
| ENADID 2018 | 0.285 | 2.54 | 65.0% | 493 | 0.730 | 0.854 | 34.1 |
| ENADID 2023 | 0.193 | 2.86 | 56.6% | 597 | 0.677 | 0.823 | 33.3 |
| Transacciones 2021 | 0.080 | 4.19 | 0.5% | 1,401 | 0.469 | 0.685 | 47.5 |

Notes: Annual $R^2$ by matrix reported in Appendix Table A1. Source: Authors' calculations.

**Figure 1. Entropy and Predictive Accuracy Across Allocation Matrices, 2013–2025**

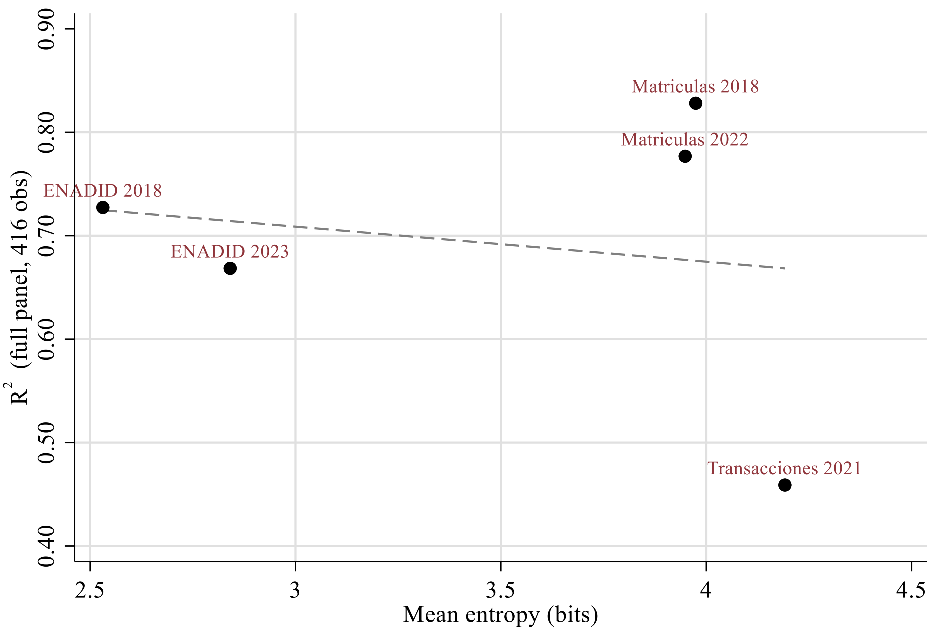


Notes: Each point represents one allocation matrix. Source: Authors' calculations.

Figure 2 evaluates the fit of the preferred matrix (*MCAS* 2018) by plotting predicted against observed remittance receipts for two benchmark years, the beginning (2013) and the end (2025) of the sample period. In both panels, the distribution of points lies close to a positively sloped line, with no evidence of structural breaks or systematic misallocation across states. Two features are particularly noteworthy. First, the stability of the relationship over more than a decade confirms that a fixed diaspora distribution can track substantial variation in aggregate remittance flows. This provides strong evidence of the persistence of migration networks and validates the use of a time-invariant allocation matrix for structural analysis. Second, the largest recipient states—Michoacán, Jalisco, and Guanajuato—are accurately captured, indicating that the model performs best precisely where the economic magnitude of remittances is greatest.

All matrices exhibit a systematic negative bias, which is expected as they allocate only U.S.-origin remittances, while the official series includes inflows from all countries. This bias does not affect the relative allocation across states and is therefore not indicative of model misspecification. The transaction-based matrix is the only exception, displaying a slight positive bias, consistent with its tendency to overrepresent formal transfer channels rather than underlying population geography.

**Figure 2. Predicted and Observed Remittance Receipts, 2013 and 2025**

**a. 2013**

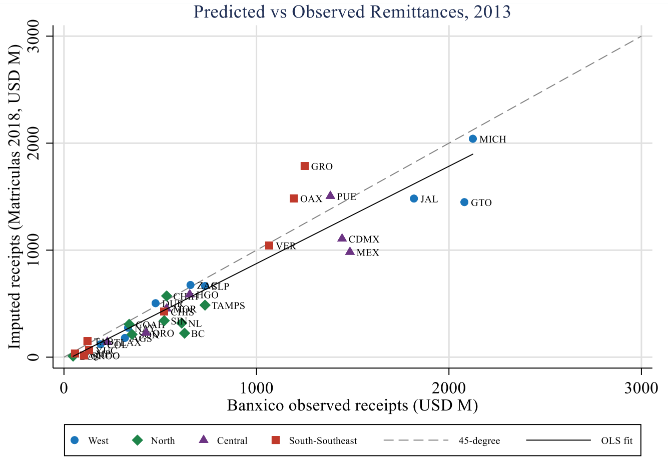


**b. 2025**

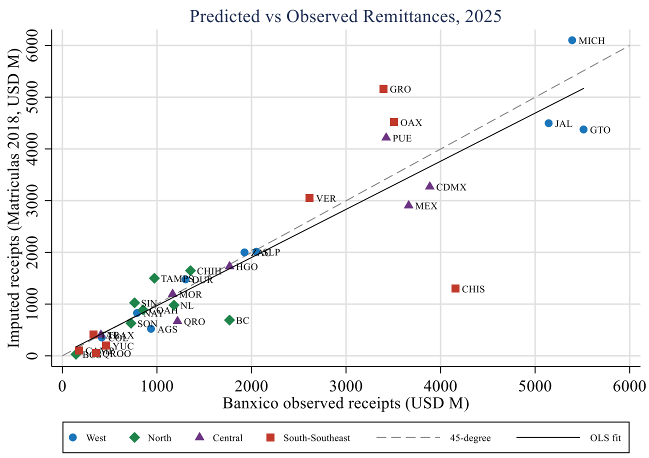


Notes: n = 32 Mexican states per panel. Dashed = 45° perfect prediction. Solid = OLS fit. Point colour = INEGI region.
Source: Authors' calculations.

This evidence establishes a clear hierarchy of data quality, where consular-based matrices best recover the geographic structure of remittance flows, followed by survey-based matrices, while transaction data perform poorly for volume imputation. More importantly, the results identify the mechanism behind this hierarchy: Predictive accuracy depends on capturing the concentration of the migrant stock, not the frequency of remitting behavior.

## 5. Bilateral Corridor Structure and Dynamics

The validation results establish *MCAS* 2018 as the preferred allocation matrix and provide a foundation for recovering the bilateral structure of the U.S.–Mexico remittance system. This section uses that matrix to analyze both the long-run architecture of the corridor and the structural changes associated with the COVID period. The analysis proceeds in two steps. Section 5.1 characterizes the system under a fixed-diaspora counterfactual, isolating the role of U.S. state-level outflow growth. Section 5.2 introduces time variation in allocation weights to identify structural reallocation.

### 5.1 Long-Run Corridor Structure: 2013-2025

Using *MCAS* 2018 matrix as a fixed allocation device, bilateral flows are estimated for 2013–25. Under this specification, all variation in corridor-level flows is driven exclusively by changes in U.S. state-level outflows, while the structure of bilateral linkages remains constant. Table 2 lists the largest U.S.–Mexico bilateral corridors, ordered by their estimated flow levels in 2025. It shows that the system is highly concentrated within a small set of persistent links. California alone accounts for most of the top positions, including the single largest corridors to Michoacán, Jalisco, Oaxaca, and Guerrero, while Texas appears as the second structural pillar, with its main corridors directed toward West and North Mexican states. This pattern reflects the durability of long-standing migration networks, classic corridors such as New York–Puebla that remain structurally important but lower in the ranking relative to other corridors originating in California and Texas. Appendix Table A2 lists the fastest-growing bilateral corridors over 2013–25. These are dominated by Ohio, Iowa, and New Jersey as sending states, with South-Southeast Mexican destinations, particularly Oaxaca, Veracruz, and Chiapas, accounting for the majority of the top-ranked corridors. This pattern confirms that the growth frontier of the remittance system is geographically distinct from its absolute backbone

**Table 2. Top 15 Bilateral Remittance Corridors by Estimated Flow in 2025**

| # | U.S. State | Mexican State | Region | $w_{ij}$ 2018 (%) | Flow 2025 (USD M) |
|---|---|---|---|---|---|
| 1 | California | Michoacán | West | 15.9% | $3,011M |
| 2 | California | Jalisco | West | 13.3% | $2,521M |
| 3 | California | Oaxaca | South-Southeast | 10.3% | $1,938M |
| 4 | California | Guerrero | South-Southeast | 9.7% | $1,825M |
| 5 | California | Puebla | Central | 6.8% | $1,279M |
| 6 | Texas | Guanajuato | West | 12.8% | $1,130M |
| 7 | California | CDMX | Central | 6.0% | $1,129M |
| 8 | California | Guanajuato | West | 5.8% | $1,091M |
| 9 | Texas | San Luis Potosí | West | 12.1% | $1,068M |
| 10 | Texas | Tamaulipas | North | 11.6% | $1,024M |
| 11 | New York | Puebla | Central | 44.7% | $809M |
| 12 | California | Estado de México | Central | 3.9% | $734M |
| 13 | Texas | Nuevo León | North | 8.1% | $715M |
| 14 | California | Sinaloa | North | 3.3% | $621M |
| 15 | California | Zacatecas | West | 3.2% | $602M |

Source: Authors' calculations.

Figure 3 complements the corridor-level analysis by aggregating bilateral flows into four origin-state groups: California; Texas; a Southeastern cluster of Florida, Georgia, North Carolina, and South Carolina (FL–GA–NC–SC); and a Mid-Atlantic/Midwestern cluster of New York, New Jersey, Illinois, and Ohio (NY–NJ–IL–OH). This grouping captures both the historical backbone of the corridor—anchored in California and Texas—and the rise of newer sending hubs in the Southeast and Mid-Atlantic/Midwest as Mexican migration has diversified geographically (Massey, Durand, and Malone, 2002; Massey et al., 2010). As shown, each Mexican region exhibits a distinct composition of U.S.-origin groups, where California overwhelmingly dominates in the West, Texas plays a relatively larger role in the North, while the emerging Eastern clusters are increasingly influential in South-Southeast Mexico. Thus, regional exposure to U.S. shocks is mediated not only by remittance volume but also by the underlying mix of origin states. Additionally, two patterns stand out. First, the system displays strong persistence at its core: California and Texas remain the dominant sources of remittances to all Mexican regions, especially the West and Central. Second, marginal growth is increasingly driven by the non-traditional clusters: Contributions from FL–GA–NC–SC and NY–NJ–IL–OH rise steadily, with the sharpest gains accruing to South-Southeast Mexico. This evolution reveals a growth frontier that is distinct from, yet nested within, the historical backbone. Although each emerging sending state is relatively small on its own, together they account for a disproportionate share of incremental flows and are disproportionately connected to Southern destinations, so that the coexistence of a stable core and an expanding frontier becomes the defining structural feature of the remittance system.

**Figure 3. Remittances Inflows to Mexican Regions by U.S.-Origin, 2013-2025**

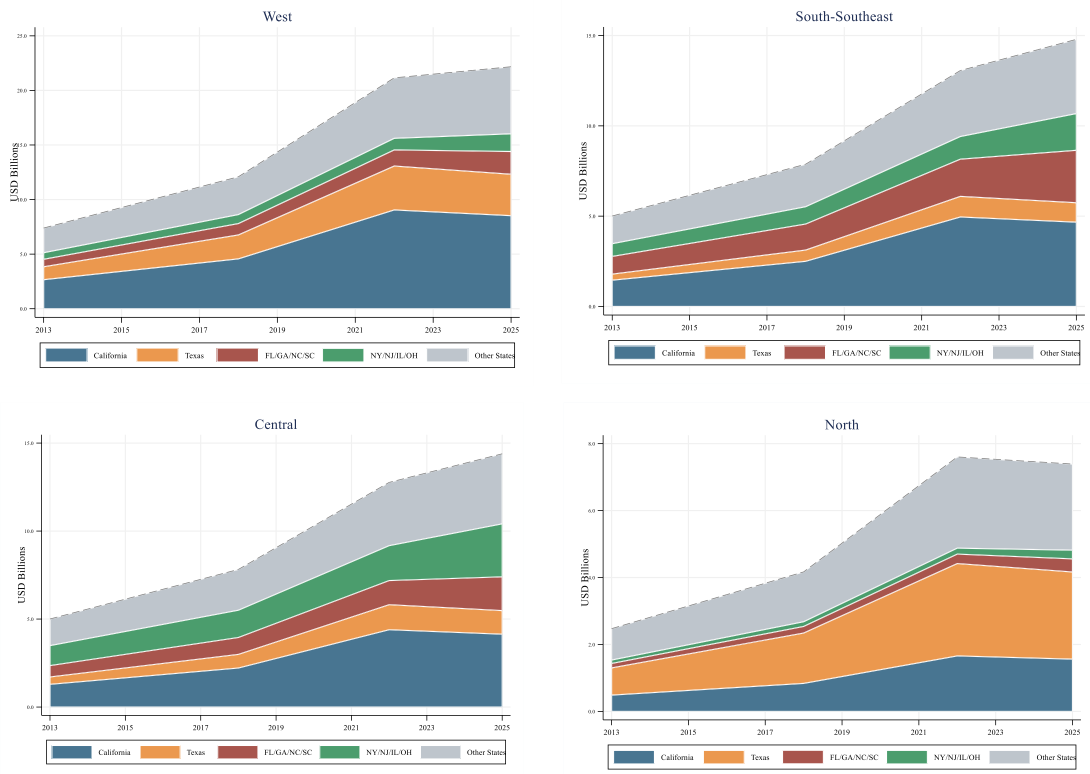


Notes: Groups colors in blue: California, orange: Texas, red: FL+GA+NC+SC, green: NY+NJ+IL+OH, grey:Other States.
Source: Authors' calculations.

## 5.2 Structural Change: Scale, Reallocation, and the Southward Shift

The analysis now moves from the long-run structure of the remittance corridor to a closer look at how that structure changed during the COVID period. A natural question arises: Did the unprecedented rise in remittances simply reflect a *scale effect*—that is, larger outflows from the United States applied to a largely unchanged corridor map—or was there also a meaningful *reallocation* of flows across Mexican states, consistent with the southward shift that recent research has pointed to?

The fixed-matrix results from the previous section offer a useful baseline. Under that counterfactual, bilateral remittances evolve only as U.S. state-level outflows change, while the pattern of bilateral links stays constant. By construction, the allocation weights satisfy $\sum_{i} w_{ji} = 1,$ for every U.S. state $j$, so total outflows from each state are fully distributed across Mexican destinations. Aggregating over U.S. states then gives estimated remittance receipts for each Mexican state as $Y_i(t) = \sum_{j} \hat{R}_{ji}\,(t)$. To quantify structural change over the COVID period, we require a method that separates *how much* of the observed change in remittance receipts is due to expansion in aggregate U.S. outflows (*scale*) versus how much is due to changes in the geographic allocation of those outflows across Mexican destinations (*structural reallocation*). The appropriate tool for this purpose is an Oaxaca–Blinder decomposition, adapted from labor economics (Blinder, 1973; Oaxaca, 1973) and widely used to decompose group differences in outcomes into endowment effects and coefficient effects. Here, the "groups" are the pre-period (2018) and post-period (2022), the outcome is remittance receipts $Y_i$, endowments are U.S. state-level outflows $R_j$, and coefficients are the allocation weights $w_{ji}$.

This framework helps address the question of the extent to which the COVID-era remittance surge reflect aggregate growth versus a redistribution of flows across Mexican regions. Formally, let $t_0 = 2018$ (pre-period) and $t_1 = 2022$ (post-period). The change in receipts for Mexican state i is:

$$\boldsymbol{\Delta Y_i = Y_i(t_1) - Y_i(t_0) = \sum_{j} \left[w_{ji}(t_1)R_j(t_1) - w_{ji}(t_0)R_j(t_0)\right]} \qquad (7)$$

Adding and subtracting $\sum_{j} w_{ji}\,(t_0)R_j(t_1)$ yields the standard Oaxaca–Blinder decomposition into three components:

$$\boldsymbol{\Delta Y_i = \sum_{j} \Delta\, w_{ji} * R_j(t_0) + \sum_{j} w_{ji}\,(t_0) * \Delta R_j + \sum_{j} \Delta\, w_{ji}(t_0) * \Delta R_j} \qquad (8)$$

where
$\Delta w_{ji} = w_{ji}(t_1) - w_{ji}(t_0)$ and
$\Delta R_j = R_j(t_1) - R_j(t_0)$.

The structural effect holds U.S. outflows fixed at their pre-period level, $R_j(t_0)$, and captures redistribution due to shifts in migrant-origin geography. Specifically, it considers the southward shift that this paper hypothesizes by measuring how much receipts would change if only the diaspora allocation weights changed. For the scale effect, it holds the bilateral allocation structure fixed at its pre-period level, $w_{ji}(t_0)$, and captures the pure aggregate expansion channel by measuring how much remittance receipts would change if only U.S. state-level outflows grew. Finally, the interaction term reflects the covariance between changes in allocation weights and changes in outflows; a positive interaction indicates that corridors experiencing increased diaspora shares also experienced above-average outflow growth, thereby amplifying the structural effect.

Figure 4 summarizes the pattern by plotting changes in bilateral allocation weights between the pre- and post-COVID periods. Despite differences in magnitude across sources, a common picture emerges: Allocation shares shift systematically toward South-Southeast Mexico and away from the West and North, with geographically widespread changes that span multiple U.S. origin states and are especially marked in corridors linking emerging sending states to Southern destinations. For robustness, the decomposition is carried out using bilateral allocation weights from ENADID 2018–23, which reflect more recent emigration flows and yield an upper-bound estimate. The corresponding visualization for the ENADID-based matrices is reported in Appendix Figure A3.

**Figure 4. Changes in Bilateral Allocation Weights Between the Pre- and Post-COVID Periods: Matrículas Consulares, 2018–2022**

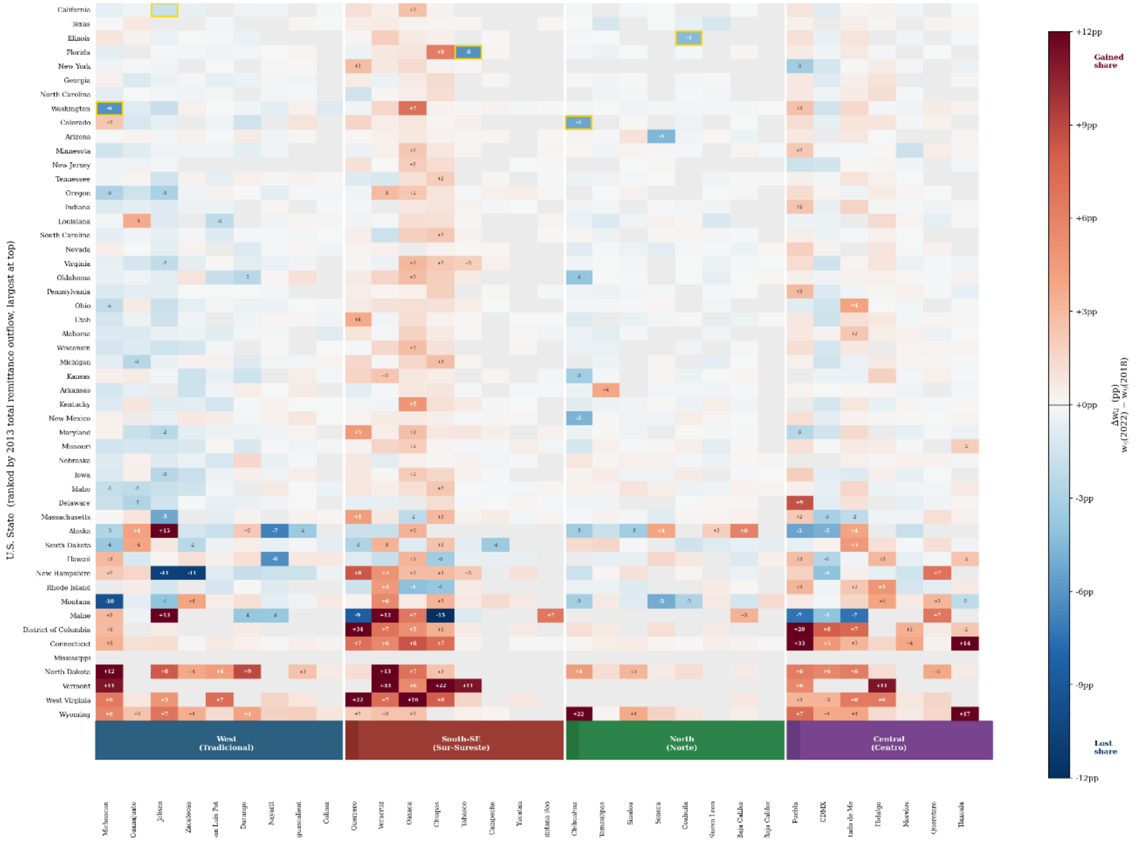


Notes: Red = corridor gained diaspora share | Blue = corridor lost share. Source: Authors' calculations.

While the previous figures document the direction of structural change, they do not quantify its economic significance. Figure 5 fills this gap by applying an Oaxaca–Blinder decomposition to the change in remittance receipts between the pre-COVID (2018) and post-COVID (2022) periods, for both MCAS (panel a) and ENADID (panel b). The decomposition separates two main channels: a scale effect, which measures how much receipts would have changed if only aggregate U.S. outflows had grown while the bilateral corridor structure remained fixed, and a structural effect, which measures how much receipts changed because migrants from different Mexican origin states came to represent larger or smaller shares of the diaspora in each U.S. state, independently of aggregate outflow growth. A third component, the interaction effect, captures the covariance between these two forces.

In both panels, the dominant pattern at the national level is that the boom is overwhelmingly driven by the scale component where in every region, the single largest contribution to growth comes from the expansion of aggregate U.S. outflows rather than from changes in corridor geography. At the same time, the figure makes clear that this aggregate expansion is redistributed unevenly across Mexican regions. South-Southeast Mexico is the only region with a positive structural effect in both specifications, with roughly one-sixth to one-third of its growth attributable to changes in allocation weights rather than to scale alone. By contrast, the West and North exhibit sizable negative structural bars, indicating that, once pure scale is netted out, they capture less of the boom than their pre-COVID network position would predict. Meanwhile, the Central region remains broadly neutral.

The ENADID panel (Figure 5b), which reflects the most recent migration patterns from the household survey, is particularly striking. For South-Southeast Mexico, the scale and structural bars are almost identical in height, each accounting for roughly half of the region's total remittance increase, so diaspora reallocation contributes nearly as much as aggregate U.S. outflow growth to its COVID-era gains. This is a revealing result, which implies that the southward shift is not a minor side effect of a scale-driven boom but a central driver of regional divergence.

In both panels, the interaction bars are small and of mixed sign across regions, confirming that the joint covariance between changes in allocation weights and changes in U.S. outflows is second-order relative to the scale and structural components; adding them to the structural term would only slightly raise the southward contribution without changing the conclusion that the remittance surge was nationally scale-driven yet regionally redistributive, with a particularly pronounced tilt toward South-Southeast Mexico.

Considering the decomposition in Figure 5, these findings help reconcile the apparent paradox of the COVID-era remittance boom. Nationally, growth is largely a scale phenomenon, but regionally it reflects a non-trivial redistribution of diaspora geography. The pattern is fully consistent with the emergence of new migration corridors documented in Section 5.1 and underscores the rising role of non-traditional U.S. sending states in channeling remittances toward Southern destinations. Overall, the evidence points to a dual structure. The remittance system remains remarkably stable at its core, anchored in long-standing migration networks, yet its marginal growth is increasingly driven by new nodes with distinct destination profiles. The COVID shock did not overturn this architecture but amplified it, accelerating a pre-existing southward shift in the geography of Mexican migration and remittance flows.[5]

[5] This interpretation is also consistent with the broader migration context of the period. The growing presence of Central American migrants in Southern border states such as Chiapas and Tabasco may have contributed to the regional scale component by increasing local remittance inflows, while affecting the structural component to a lesser extent, since the corridor mapping focuses on U.S.–Mexico origin–destination links. As a result, the main conclusions about southward reallocation and the role of Mexican diaspora geography remain valid. A more detailed examination of how Central American migration interacts with subnational remittance patterns in Southern Mexico is an important topic for future research, but lies beyond the scope of this paper.

**Figure 5. Oaxaca-Blinder Decomposition COVID-Era Remittance Growth by Region, 2018-2022**

**a. *MCAS***

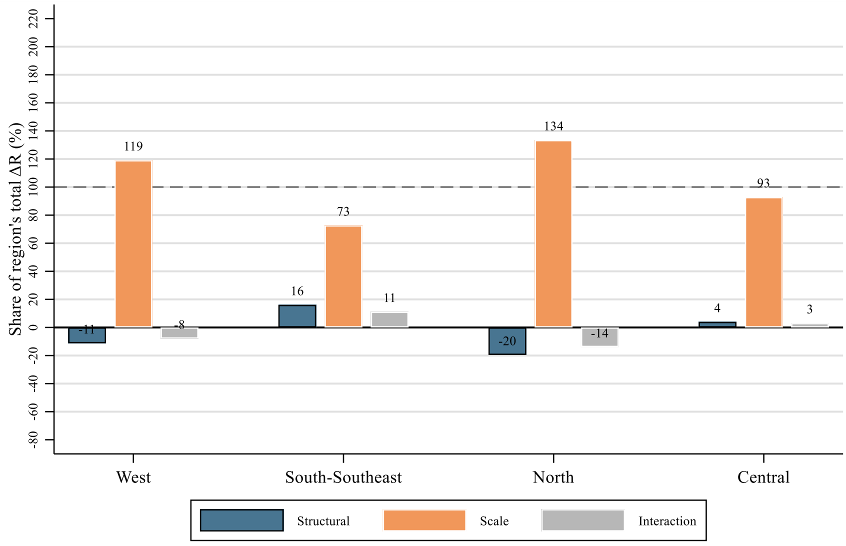


**b. ENADID**

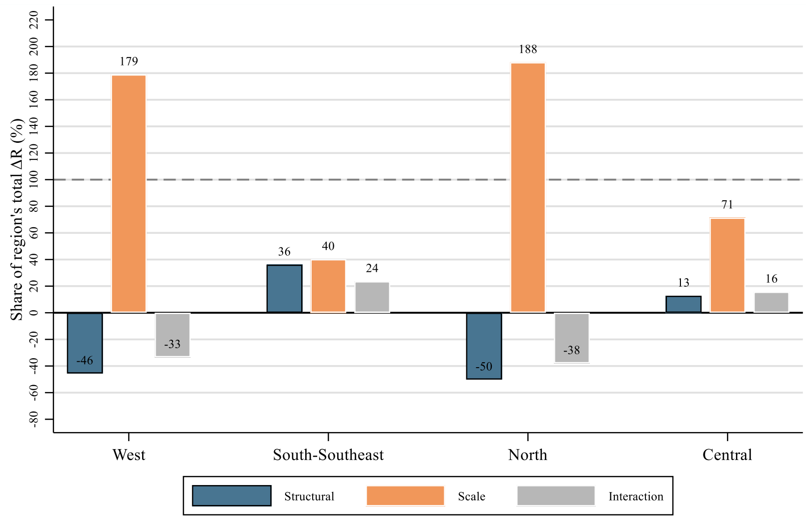


Notes: *MCAS* 2018/2022; ENADID 2018/2023; Banco de México.
Source: Authors' calculations.

### 5.3 Robustness

Appendix B evaluates whether the positive structural contribution to South-Southeast Mexico could be generated mechanically by violations of the corridor-proportionality assumption. Two complementary exercises are reported there. First, the preferred *MCAS* 2018 matrix is compared with a transaction-based matrix derived from Banorte wire transfers. The comparison shows that proportionality holds most closely in the backbone corridors that account for the bulk of remittance volume, while the main deviations arise in weaker or institutionally distorted corridors and do not imply that Southern destinations are spuriously overweighted.

Second, Appendix B implements targeted perturbation tests that intentionally reallocate post-COVID weights away from South-Southeast Mexico in the weakest-fitting corridors. In every scenario, the structural effect for South-Southeast remains positive. Taken together, these exercises indicate that the southward shift identified in Section 5.2 is not an artifact of the imputation procedure, but a robust feature of the bilateral remittance system.

Third, the corridor assumption holds most closely in the backbone corridors that account for the majority of remittance volume, while the deviations that appear in weaker corridors are economically interpretable and quantitatively insufficient to overturn the results.

## 6. Conclusion

This paper develops and validates a framework to estimate subnational bilateral remittance flows in the absence of direct corridor-level data, using the U.S.–Mexico corridor as a test case. By combining official U.S. state-level remittance outflows with migrant-origin distributions from multiple sources and validating the resulting matrices against more than a decade of Banco de México state-level receipts, the analysis provides a systematic empirical assessment of the corridor-proportionality approach.

Three findings stand out. First, there is a clear hierarchy in data quality for bilateral imputation. Matrices based on consular records deliver the strongest predictive performance, survey-based matrices perform moderately well, and transaction-based data perform substantially worse. The central mechanism is geographic concentration: Matrices that preserve the concentrated structure of the Mexican diaspora provide a more accurate accounting of the spatial distribution of remittances than matrices that diffuse flows too broadly across destinations. This result suggests that, for the purpose of recovering aggregate remittance geography, population-based data may be more informative than transaction records that reflect channel-specific remitting behavior.

Second, the bilateral remittance system exhibits a highly persistent long-run structure. A small set of corridors linking California and Texas to traditional western-central Mexican states continues to dominate the system, consistent with the durable role of migration networks. At the same time, the analysis identifies a distinct growth frontier in non-traditional U.S. sending states, especially in the Midwest, Mid-Atlantic, and Southeast, whose migrant populations are disproportionately connected to South-Southeast Mexico. The remittance corridor is therefore best understood as a stable core combined with an expanding frontier, rather than as either a fixed system or a fully reconfigured one.

Third, the COVID-era remittance surge was scale-driven at the national level but not distributionally neutral. The expansion of aggregate U.S. outflows accounts for essentially all of the increase in total remittances, yet bilateral decomposition shows that structural reallocation shifted part of this growth toward South-Southeast Mexico. This region is the only one to display a positive structural effect across specifications, while the West and North receive less than their pre-pandemic network position would predict. The results therefore provide direct evidence of an ongoing southward shift in the geography of Mexican migration and remittances.

These findings carry implications for measurement, theory, and policy. From a measurement perspective, they show that bilateral remittance estimation is feasible with publicly available population-based data, but that transaction data drawn from a single financial institution should be used cautiously when the goal is to recover geographic structure. From a theoretical perspective, the results reinforce the view that migration networks are highly persistent but not static. Aggregate structure changes slowly, yet meaningful redistribution can occur through the gradual reweighting of corridors as new sending and receiving nodes emerge. From a policy perspective, this reallocation matters because the marginal impact of remittance growth is increasingly concentrated in poorer and less financially tied regions of Southern Mexico, where additional inflows are likely to have different welfare and development effects than in traditional remittance-receiving states.

The framework developed here is portable to other major remittance corridors provided two inputs are available: subnational remittance outflows in the origin country and migrant-stock data convertible into origin–destination allocation matrices. Both inputs exist, in varying degrees of quality, for several of the world's largest corridors. The World Bank's KNOMAD bilateral remittance matrix uses migrant stock data to estimate corridor-level flows globally, though at the national rather than subnational scale (Ratha et al., 2022). For the Philippines–U.S. corridor, De Arcangelis et al. (2023) show that sender surveys capture more than 90 percent of true remittance amounts, though accuracy varies with migrant characteristics such as legal status and remittance frequency, a pattern consistent with this paper's finding that behavioral transaction data underperform population-based sources for geographic imputation. For the India–Gulf corridor, where migration is predominantly temporary and circular and informal hawala channels are widespread, formal transaction data likely undercount total flows, and the proportionality assumption may be weaker because per-migrant remittance amounts vary systematically by occupation, contract type, and financial access (Chhibber, 2025). Applying the same validation logic employed here, structural diagnostics, multiple fit metrics, and a scale-versus-structure decomposition would allow central banks and national statistical offices in these and other corridors to test, rather than assume, whether diaspora-based imputation can depict bilateral remittance geography with sufficient accuracy for policy use.

The development stakes of this kind of measurement investment are substantial. In the U.S.–Mexico corridor alone, the southward redistribution of remittance growth documented here implies that the marginal dollar of remittance income is increasingly flowing to states where households face binding liquidity constraints, where formal financial systems are shallowest, and where the consumption and investment multipliers of additional income are likely highest. At the same time, the 2025 contraction in aggregate remittance inflows, the first in more than a decade, will not be absorbed uniformly. Corridors anchored in emerging, lower-volume U.S. sending states, which are disproportionately linked to South–Southeast Mexican destinations, may prove more fragile than the historical backbone precisely because their network ties are newer, less institutionalized, and more exposed to shifts in U.S. immigration and labor-market policy, such as changes in enforcement intensity, work authorization, or access to state-level protections. In this context, identifying which Mexican regions face the greatest exposure to a reversal in remittance growth, and how that exposure interacts with evolving policy regimes in destination states, is a question that bilateral corridor data can answer and aggregate statistics cannot.

**References**

Acosta, P., Calderón, C., Fajnzylber, P., & López, H. (2008). What is the impact of international remittances on poverty and inequality in Latin America? World Development, 36(1), 89–114. https://doi.org/10.1016/j.worlddev.2007.02.016

Adams, R. H., & Page, J. (2005). Do international migration and remittances reduce poverty in developing countries? World Development, 33(10), 1645–1669. https://doi.org/10.1016/j.worlddev.2005.05.004

Ambrosius, C., & Cuecuecha, A. (2016). Remittances and the use of formal and informal financial services. World Development, 77, 80–98. https://doi.org/10.1016/j.worlddev.2015.08.010

Anderson, J. E., & van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. American Economic Review, 93(1), 170–192. https://doi.org/10.1257/000282803321455214

Banco de México. (2022). Sistema de Información Económica. https://www.banxico.org.mx/SieInternet/

Banco de México (2025). Extracto del Reporte sobre las Economías Regionales Abril – Junio 2025, Recuadro 4, pp. 66-68, documento publicado el 11 de septiembre de 2025.

BBVA (2021). ¿Habrá una reconfiguración en el Mapa Migratorio? BBVA Research. Observatorio Migración México, 1 octubre de 2021.

BBVA (2026). Terminan 11 años consecutivos de crecimiento en las remesas, caen 4.6% en 2025. BBVA Research. Observatorio Migración México, 3 febrero de 2026.

Beine, M., Docquier, F., & Özden, Ç. (2015). Dissecting network externalities in international migration. Journal of Demographic Economics, 81(4), 379–408. https://doi.org/10.1017/dem.2015.14

Clemens, M. A., & McKenzie, D. (2018). Why don't remittances appear to affect growth? Economic Journal, 128(612), F179–F209. https://doi.org/10.1111/ecoj.12463

CEMLA and Banorte (2024). *La emigración mexicana femenina, el ingreso por remesas y la inclusión financiera*. Centro de Estudios Monetarios Latinoamericanos and Banorte, May 2024.

CONEVAL. (2022). Medición de la pobreza en México 2022. https://www.coneval.org.mx

Chhibber, K. (2025). Fintech Corridors and Informal Flows: Remaining Gulf-India Financial Security. Observer Research Foundation, https://www.orfonline.org/expert-speak/fintech-corridors-and-informal-flows-reimagining-gulf-india-financial-security

De Arcangelis, G., Fertig, A., Liang, Y., Srouji, P., & Yang, D. (2023). Measuring remittances. *Journal of development economics*, *161*, 103004.

Demirgüç-Kunt, A., Klapper, L., Singer, D., Ansar, S., & Hess, J. (2018). *The Global Findex Database 2017: Measuring financial inclusion and the fintech revolution*. World Bank. https://doi.org/10.1596/978-1-4648-1259-0

Ratha, D., Plaza, S., & Ju Kim, E. (2022). Bilateral Remittance Matrix (new). In observance of the International Migrants Day, December 18. World Bank Blogs, Published on People Move, https://blogs.worldbank.org/en/peoplemove/bilateral-remittance-matrix-new.

Docquier, F., Rapoport, H., & Salomone, S. (2012). Remittances, migrants' education and immigration policy: Theory and evidence from bilateral data. Regional Science and Urban Economics, 42(5), 817–828. https://doi.org/10.1016/j.regsciurbeco.2012.03.005

Dumont, J.-C., & Lemaître, G. (2005). Counting immigrants and expatriates: A new perspective. OECD Social, Employment and Migration Working Papers No. 25. https://doi.org/10.1787/521408684816

Durand, J., Parrado, E. A., & Massey, D. S. (1996). Migradollars and development: A reconsideration of the Mexican case. International Migration Review, 30(2), 423–444. https://doi.org/10.1177/019791839603000203

Flores, M., & Torre, L. (2024). Regional determinants in origin–destination remittance flows from the US to Mexico: Evidence from the COVID-19 pandemic. Regional Studies, 58(12), 2389–2403. https://doi.org/10.1080/00343404.2024.2409175

Flores, M., & Gutiérrez Bernal, S. (2025). Estimation of state-to-state remittance flows from the US to Mexico. *Quality & Quantity*, 1-20.

Garip, F. (2012). Discovering diverse mechanisms of migration: The Mexico–US stream 1970–2000. American Journal of Sociology, 118(2), 392–433. https://doi.org/10.1086/666483

Garip, F., & Asad, A. L. (2016). Network effects in migration. Annual Review of Sociology, 42, 151–171. https://doi.org/10.1146/annurev-soc-081715-074035

Genoni, M. E., Muñoz Boudet, A. M., & Sahay, S. (2021). *The rise of formal remittances during COVID-19* (Policy Research Working Paper No. 9861). World Bank. https://doi.org/10.1596/1813-9450-9861

Hanson, G. H. (2010). International migration and the developing world. In D. Rodrik & M. Rosenzweig (Eds.), Handbook of Development Economics (Vol. 5, pp. 4363–4414). Elsevier. https://doi.org/10.1016/S1573-4471(10)05015-3

Hanson, G. H., & Woodruff, C. (2003). Emigration and Educational Attainment in Mexico. University of California at San Diego, Mimeographed, 39.

Huerta, A. V., & McLean, L. (2019). Caravanas de migrantes en México. *Revista CIDOB d'afers internacionals*, (122), 163-186.

Massey, D. S. (1987). Understanding Mexican migration to the United States. American Journal of Sociology, 92(6), 1372–1403. https://doi.org/10.1086/228669

Massey, D. S., Durand, J., & Malone, N. J. (2002). Beyond smoke and mirrors: Mexican immigration in an era of economic integration. Russell Sage Foundation.

Massey, D. S., Arango, J., Hugo, G., Kouaouci, A., Pellegrino, A., & Taylor, J. E. (1993). Theories of international migration: A review and appraisal. *Population and development review*, 431-466.

Massey, D. S., & Gentsch, K. (2014). Undocumented migration. Annual Review of Sociology, 40, 107–127. https://doi.org/10.1146/annurev-soc-073012-112430

Massey, D. S., Durand, J., & Pren, K. A. (2016). Why border enforcement backfired. American Journal of Sociology, 121(5), 1557–1600. https://doi.org/10.1086/684200

Mora-Rivera, J., García-Amador, M. C., & Van Gameren, E. (2024). A regional analysis of the impact of remittances on health expenditures: evidence from Mexico. *Review of Economics of the Household*, 1-32.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the US labor market. Quarterly Journal of Economics, 118(2), 549–599. https://doi.org/10.1162/003355303321675455

Rapoport, H., & Docquier, F. (2006). The economics of migrants' remittances. In S.-C. Kolm & J. M. Ythier (Eds.), Handbook of the Economics of Giving, Altruism and Reciprocity (Vol. 2, pp. 1135–1198). Elsevier. https://doi.org/10.1016/S1574-0714(06)02017-3

Torre Cantalapiedra, E. (2022). El estudio de las caravanas migrantes en México. *Norteamérica*, *17*(2), 67-89.

Woodruff, C., & Zenteno, R. (2007). Migration networks and microenterprises in Mexico. Journal of Development Economics, 82(2), 509–528. https://doi.org/10.1016/j.jdeveco.2006.03.006

World Bank. (2022). Migration and Development Brief 37. https://www.worldbank.org

Yang, D. (2008). International migration, remittances and household investment: Evidence from Philippine migrants' exchange rate shocks. Economic Journal, 118(528), 591–630. https://doi.org/10.1111/j.1468-0297.2008.02134.x

Yang, D. (2011). Migrant remittances. Journal of Economic Perspectives, 25(2), 129–152. https://doi.org/10.1257/jep.25.2.129

**Appendix A**

**Table A1. Annual $R^2$ by Allocation Matrix, 2013–2025**

| Year | Matrículas 2018 | Matrículas 2022 | ENADID 2018 | ENADID 2023 | Transacciones |
|---|---|---|---|---|---|
| 2013 | 0.882 | 0.851 | 0.635 | 0.586 | 0.324 |
| 2014 | 0.879 | 0.840 | 0.669 | 0.585 | 0.334 |
| 2015 | 0.885 | 0.848 | 0.748 | 0.626 | 0.330 |
| 2016 | 0.889 | 0.845 | 0.738 | 0.605 | 0.338 |
| 2017 | 0.875 | 0.831 | 0.757 | 0.621 | 0.349 |
| 2018 | 0.868 | 0.823 | 0.785 | 0.628 | 0.360 |
| 2019 | 0.871 | 0.825 | 0.776 | 0.629 | 0.357 |
| 2020 | 0.861 | 0.807 | 0.758 | 0.587 | 0.358 |
| 2021 | 0.858 | 0.804 | 0.740 | 0.591 | 0.378 |
| 2022 | 0.836 | 0.808 | 0.694 | 0.637 | 0.398 |
| 2023 | 0.773 | 0.763 | 0.604 | 0.602 | 0.418 |
| 2024 | 0.783 | 0.774 | 0.576 | 0.564 | 0.406 |
| 2025 | 0.811 | 0.810 | 0.620 | 0.637 | 0.414 |

*Note:* $R^2$ from OLS regression of imputed on actual Banco de México annual state-level receipts ($n = 32$ per year).

**Table A2. Top 15 Fastest-Growing Bilateral Corridors, 2013–2025**

| # | U.S. State | Mexican State | Region | $w_{ij}$ 2018 (%) | Flow 2013 (USD M) | Flow 2025 (USD M) | Δ Flow (USD M) |
|---|---|---|---|---|---|---|---|
| 1 | Ohio | Veracruz | South-Southeast | 5.6% | $12M | $65M | $53M |
| 2 | Ohio | Oaxaca | South-Southeast | 13.9% | $29M | $162M | $133M |
| 3 | Ohio | Guanajuato | West | 13.2% | $28M | $154M | $126M |
| 4 | Ohio | Estado de México | Central | 8.3% | $17M | $97M | $80M |
| 5 | Ohio | Michoacán | West | 14.2% | $30M | $165M | $135M |
| 6 | Ohio | Chiapas | South-Southeast | 7.1% | $15M | $83M | $68M |
| 7 | Ohio | CDMX | Central | 6.6% | $14M | $77M | $63M |
| 8 | Iowa | Michoacán | West | 13.7% | $14M | $62M | $48M |
| 9 | Iowa | Jalisco | West | 11.2% | $11M | $51M | $40M |
| 10 | New Jersey | Veracruz | South-Southeast | 4.2% | $17M | $67M | $50M |
| 11 | New Jersey | Puebla | Central | 35.9% | $150M | $572M | $422M |
| 12 | New Jersey | Estado de México | Central | 5.0% | $21M | $79M | $58M |
| 13 | New Jersey | CDMX | Central | 5.8% | $24M | $92M | $68M |
| 14 | New Jersey | Oaxaca | South-Southeast | 26.8% | $112M | $428M | $316M |
| 15 | New Jersey | Guerrero | South-Southeast | 7.4% | $31M | $118M | $87M |

Notes: Under the fixed *MCAS* 2018 matrix, all corridors from the same U.S. state grow at identical rates equal to that state's aggregate outflow growth. Source: Author´s calculations.

**Figure A3. Changes in Bilateral Allocation Weights Between the Pre- and Post-COVID Periods: ENADID, 2018–2023**

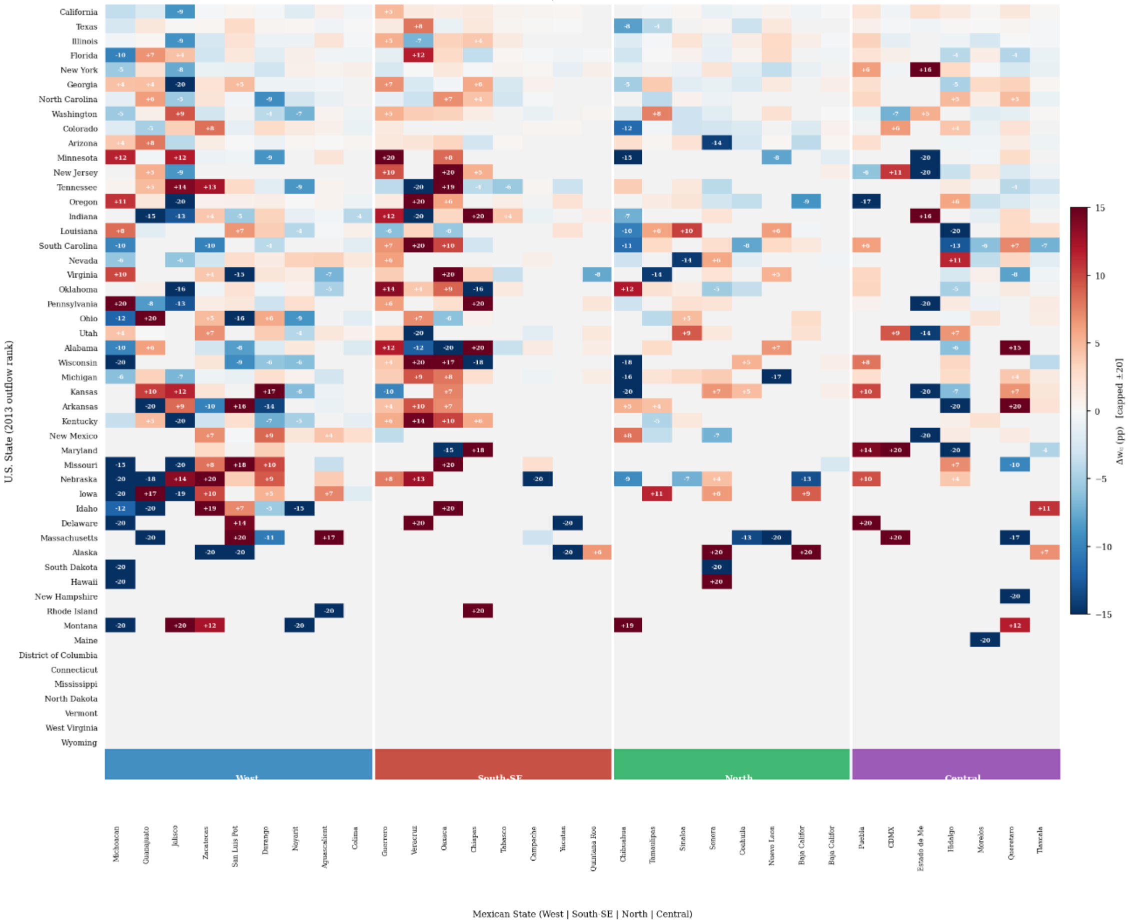


Source: Author´s calculations.

## Appendix B. Robustness of the Corridor-Proportionality Assumption

This appendix evaluates the corridor-proportionality assumption underlying the bilateral imputation strategy. Under this assumption, migrant-origin shares within each U.S. state proxy for the destination shares of remittance outflows from that state. The appendix addresses two related concerns: whether transaction data reveal systematic departures from proportionality, and whether weakly fitted post-COVID corridors could mechanically generate the positive structural effect for South-Southeast Mexico reported in Section 5.2.

## B.1 Robustness of the Corridor-Proportionality Assumption: Evidence from Transaction-Level Data

The core identifying assumption is that, within a given U.S. state, migrants from different Mexican origin states remit similar amounts on average. This allows migrant-origin shares to serve as valid proxies for remittance allocation shares. To assess this assumption, we compare the preferred *MCAS* 2018 matrix with the Banorte transaction matrix, which records the volume-weighted distribution of approximately 9.7 million formal wire transfers in 2021–2022.

Table B.1 reports the main fit statistics. To improve comparability, the table keeps the same metric definitions throughout. Pearson $r$ measures linear association across corridor shares; average row-wise $R^2$ is the simple mean of within-U.S.-state fit across non-missing rows; weighted MAPE reports percentage deviations using the *MCAS* 2018 corridor shares as weights; and the signed mean deviation reports the mean difference between transaction shares and *MCAS* shares for South-Southeast corridors, where negative values indicate that *MCAS* allocates relatively more than transactions.

The transaction comparison yields two conclusions. First, the corridor-proportionality assumption performs best in the traditional backbone corridors that carry the largest remittance volumes. Second, the principal deviations are concentrated in weaker or institutionally distorted corridors and do not provide evidence that the preferred matrix spuriously inflates South-Southeast Mexico. If anything, the direction of the deviation is consistent with the view that transaction records understate Southern destinations because formal remittance channels are less prevalent there.

**Table B1. Fit Statistics for Population-Based and Transaction-Based Allocation Matrices**

| Source Matrix | Pearson r | Avg. Row-R² | wMAPE | Signed Δ (share pts) |
|---|---|---|---|---|
| *MCAS* 2018 | 0.573 | 0.375 | 64.9% | −0.004 (SSE) |
| *MCAS* 2022 | 0.558 | 0.373 | — | — |
| ENADID 2023 | 0.377 | 0.237 | — | — |
| ENADID 2018 | 0.250 | 0.194 | — | — |

Notes: wMAPE = population-share-weighted mean absolute percentage error, computed for *MCAS* 2018 only. Signed Δ = mean(w^Trans − w^*MCAS*18) for South-Southeast corridors, reported for *MCAS* 2018 only; negative values indicate that *MCAS* over-attributes relative to transactions. Source: Authors' calculations.

Table B.2 decomposes the same comparison by Mexican region. This regional view is useful because it separates general noise from economically meaningful bias. The evidence confirms that proportionality holds best in the backbone regions of the remittance system and that the main deviations are interpretable in terms of transaction-channel structure rather than implausible behavioral differences in remitting propensity.

**Table B2. Regional Deviations Between the Matrículas 2018 and Transaction Matrices**

| Region | wMAPE | Signed Mean Δ | Interpretation |
|---|---|---|---|
| West (traditional) | 57.1% | −0.012 | Near-zero bias; assumption holds well in backbone corridors |
| Central | 48.9% | −0.014 | Slight over-attribution; assumption holds; no systematic concern |
| South-Southeast | 79.9% | −0.004 | *MCAS* over-attributes relative to transactions — southward shift is a lower bound on the true effect |
| North | 127.2% | +0.001 | Driven by Nuevo León / Banorte network concentration in Monterrey; channel-selection effect, not a propensity difference |

Notes: Signed mean Δ = mean(w_ji^Trans − w_ji^Mat18) across all corridors in the region. A negative value indicates that *MCAS* over-attributes relative to Banorte transactions. wMAPE is population-share-weighted mean absolute percentage error, with weights equal to the *MCAS* 2018 corridor share, focusing the diagnostic on economically significant flows. Source: Authors' calculations.

### B.2. Stress-Test of the Southward-Shift Finding: Targeted Matrix Perturbations

The next question is whether the observed deviations could nevertheless combine in a way that mechanically produces the positive structural effect for South-Southeast Mexico. To address this concern, Section B.2 applies a set of targeted perturbations to the post-COVID allocation matrix, deliberately shifting weights away from South-Southeast destinations in the weakest-fitting corridors. The purpose of these exercises is to test not whether the baseline matrix is exact, but whether plausible misspecification can overturn the main result.

All scenarios begin from the *MCAS* 2018 and 2022 matrices used in the main Oaxaca–Blinder decomposition. Scenario 1 targets the five emerging U.S. sending states identified as having the weakest within-state proportionality fit (New York, Georgia, North Carolina, Florida, Illinois;

mean row-wise $R^2 = 0.164$). For each of these states, the South-Southeast allocation weights in the 2022 matrix are reduced by 25 percent (Scenario 1a) and 50 percent (Scenario 1b), with the freed mass redistributed proportionally to West and Central destinations.

Scenario 2 targets the Nuevo León corridors, where *MCAS* systematically underweights Northern Mexico, and thus overweights South-Southeast ,by mechanically doubling (Scenario 2a) and tripling (Scenario 2b) the Nuevo León allocation weights in all U.S. states where the baseline share exceeds a minimal threshold (0.5 percent), capped at 25 percent, with the freed mass drawn proportionally from South-Southeast.

Scenario 3 operates on all seven U.S. states with row-wise $R^2$ below 0.15 (New York, Georgia, North Carolina, South Dakota, Massachusetts, Missouri, Montana). For each of these states, the South-Southeast allocation weights in the 2022 matrix are reduced by 25 percent (Scenario 3a) and 50 percent (Scenario 3b), with the freed mass redistributed to West and Central. This is the most sweeping perturbation: it simultaneously reduces South-Southeast allocation for all poorly-fit states, regardless of whether the mismatch reflects genuine propensity differences or Banorte network effects.

Table B.3 reports the structural effect for South-Southeast Mexico under each scenario alongside the unperturbed baseline. All six individually motivated scenarios preserve a positive structural contribution. The finding is robust to considerable misspecification in the weakest corridors of the system. Scenario 1b — the most aggressive single-mechanism perturbation, halving South-Southeast allocation from all five emerging states — reduces the structural effect from \$1.15 billion to \$0.49 billion (57 percent reduction). The effect remains positive and economically meaningful: South-Southeast still captures a disproportionately large share of the COVID-era remittance surge that cannot be attributed to scale alone. Scenario 2b, which triples Nuevo León weights and forces the freed mass out of South-Southeast, reduces the structural effect to \$0.69 billion (60 percent of baseline). Scenario 3b, which halves South-Southeast allocation for all low-fit states simultaneously, yields \$0.26 billion (23 percent of baseline), still positive.

A falsification threshold calculation reveals that the structural effect is reversed only if the *MCAS* 2022 matrix overstates South-Southeast allocation by 65 percent for every low-fit U.S. state simultaneously. Given that the signed mean deviation in shows *MCAS* actually over-attributes to South-Southeast by only 0.008 share points on average, and that the direction of this bias is explained by lower formal-channel penetration in Southern Mexico rather than by matrix misspecification, a 65 percent overstatement is not empirically plausible.

**Table B.3. Stress Test of the South-Southeast Structural Effect Under Targeted Matrix Perturbations**

| Scenario | Perturbation | Structural Effect ($B) | % of Baseline | Positive? |
|---|---|---|---|---|
| Baseline | *MCAS* 2018→2022, unperturbed | +$1.15B | 100% | Yes |
| Scenario 1a | Emerging states (NY, GA, NC, FL, IL): reduce $mat_{22}$ SSE weights −25%, redistribute to West+Central | +$0.80B | 70.03% | Yes |
| Scenario 1b | Emerging states: reduce $mat_{22}$ SSE weights −50%, redistribute to West+Central | +$0.46B | 40.05% | Yes |
| Scenario 2a | Nuevo León weights doubled (×2), reduce SSE proportionally | +$0.83B | 72.10% | Yes |
| Scenario 2b | Nuevo León weights tripled (×3, capped at 25%), reduce SSE proportionally | +$0.51B | 44.19% | Yes |
| Scenario 3a | Low-fit states (row-$R^2 \leq 0.15$): reduce $mat_{22}$ SSE weights −25%, redistribute to West+Central | +$0.83B | 72.04% | Yes |
| Scenario 3b | Low-fit states: reduce $mat_{22}$ SSE weights −50%, redistribute to West+Central | +$0.51B | 44.8% | Yes |
| Falsification threshold | SSE reduction in $mat_{22}$ required for low-fit states to drive structural effect to zero | 65% — implausible | — | — |

Notes: All perturbations modify the post-COVID allocation matrix (*MCAS* 2022) and reduce South-Southeast weights, with freed mass redistributed proportionally to West and Central destinations. Source: Authors' calculations.

The falsification-threshold result is especially informative. Reversing the sign of the structural effect would require the *MCAS* 2022 matrix to overstate South-Southeast allocation by 65 percent simultaneously across all low-fit U.S. states. That degree of joint misspecification is far beyond the range suggested by the transaction diagnostics. It is also inconsistent with the signed-deviation evidence from Section B.1, which indicates that the preferred matrix already allocates slightly more to South-Southeast than the Banorte transaction data record, and likely does so for reasons related to formal-channel undercoverage in Southern Mexico rather than to true over-allocation.

The stress tests therefore reinforce the main interpretation of the present paper. The positive structural contribution to South-Southeast Mexico is not produced by a handful of noisy corridors, nor does it disappear when the post-COVID matrix is distorted in the most unfavorable plausible direction. The result remains positive under all individually motivated perturbations and is overturned only under an extreme and empirically implausible degree of coordinated misspecification.

As discussed above, Sections B.1 and B.2 establish that the paper's main finding does not depend on an implausibly strong version of the corridor-proportionality assumption. The assumption holds most closely in the backbone corridors that account for the majority of remittance volume, while the deviations that appear in weaker corridors are economically interpretable and quantitatively insufficient to overturn the results. The positive structural contribution to South-Southeast Mexico therefore remains best understood as a genuine feature of the post-COVID remittance reallocation rather than as a by-product of the imputation procedure.