

The Invisible Channel: How Labor Migration Quietly Synchronizes Global Business Cycles

Vivan Shankar, Independent Researcher
vivan.shankar@gmail.com

ABSTRACT

This paper aims to assess whether labor migration networks are a medium for business cycle transmission across economies. To achieve this, this study utilizes a set of causal machine learning algorithms, such as causal forests and double machine learning, to estimate the heterogeneous treatment effects of labor migration on GDP correlation. To obtain this, this study utilizes a set of bilateral migration stock data from the OECD database for 1990-2023, as well as quarterly GDP data from the World Bank. The results show that a one-standard-deviation increase in labor migration network strength can result in a 0.08-0.12 percentage point increase in GDP correlation for a given pair of economies. The results are also robust to various checks, including instrumental variable estimation, where historical patterns of labor migration and geographical distance are utilized as instruments for causal inference.

JEL Codes: F22, F44, E32, C21

1. INTRODUCTION

The synchronization of business cycles across economies has important implications for the coordination of monetary policy, optimal currency areas, and the transmission of economic shocks. Although various traditional channels, such as trade and financial integration, have been well examined in the literature, as shown in Frankel and Rose (1998) and Imbs (2004), the potential role of labor migration as a business cycle transmission mechanism has not been adequately explored. Yet this phenomenon is all the more important given that, as shown in the World Bank's 2021 report, international migration flows have tripled since 1990, with 280 million people living outside their country of birth as of 2020.

In this paper, I seek to answer a key, yet still unknown, issue: Do labor migration networks causally affect business cycle synchronization between countries of origin and destination? As opposed to trade and financial integration, labor migration generates a permanent two-way flow between two countries through remittance channels, knowledge spillovers, and coordinated labor supply responses. In principle, this could either reinforce or attenuate business cycle correlation depending on whether labor migration serves as a shock absorber, with people moving countercyclically to mitigate shocks, or as a shock transmitter, with people sending remittances procyclically.

March 2026
Vol 5, No 1.

Three main contributions make this study different from the prior literature. First, the study develops a new panel data set that aggregates bilateral migration stocks from the OECD data and high-frequency data on GDP for 38 countries from 1990 to 2023, leveraging both the time and cross-sectional dimensions of migration intensity. Second, the study addresses the problem of endogeneity by using a two-tiered identification strategy: (i) migration networks from 1960-1980 as instruments for current migration stocks, and (ii) using causal machine learning techniques, namely generalized random forests (Athey et al., 2019), that allow for the estimation of heterogeneous treatment effects without assuming specific functional forms. Third, the study finds significant heterogeneity in the effect of migration on business cycle synchronization depending on skill, income, and geographic proximity, implying that the results could be misleading for suggesting homogeneous policies.

The results from the analysis show that the effect of migration networks is to significantly raise the correlation between the business cycles of the countries involved. The effect is more pronounced between particular pairs of countries. The results also show that an increase of one standard deviation in the stock of bilateral migration (i.e., an additional 50,000 migrants) would raise the GDP correlation between the countries by 0.08 to 0.12 percentage points or 15 to 20 percent of the mean correlation between the countries. The effect is more pronounced for South-South migration (i.e., from the developing world to the developing world) and is negligible for North-North migration (i.e., from the developed world to the developed world). The results also show that the effect is driven by the remittance channel rather than the labor market integration channel. The results from the causal forest analysis show that the effect of the treatment is systematically related to pre-treatment characteristics. The results show that the effect is more pronounced for countries with volatile exchange rates and flexible labor markets.

The rest of the paper is structured as follows. Section 2 discusses the relevant literature and locates the current contribution in the context of the existing literature. Section 3 outlines a conceptual framework that connects the migration channel to business cycle transmission. Section 4 discusses the data construction and reports some summary statistics. Section 5 outlines the approach, which includes the use of the causal ML technique. Section 6 reports the results and robustness checks. Section 7 discusses the mechanisms and heterogeneity. Section 8 concludes the paper and draws some implications for the policy debate.

2. LITERATURE REVIEW

This paper contributes to three strands of literature: business cycle synchronization, migration economics, and causal machine learning in macroeconomics.

2.1 Business Cycle Synchronization

The pioneering work by Frankel and Rose (1998) established the point that trade integration causes a rise in the correlation of the business cycles, which has been followed by a plethora of empirical work. The work by Imbs (2004) extended this model to include the effect of financial integration and industry structure. The more recent work by di Giovanni and Levchenko (2010) has highlighted the effect of vertical production chains, while the work by Kalemli-Ozcan et al. (2013) shows the counterintuitive effect of financial integration in reducing synchronization through the channel of risk sharing. However,

migration is not a specific channel in any of the above models; rather, it is controlled for in some of the models or excluded in most of the models.

The contribution of this work is the emphasis on the role of migration as a channel, while allowing for heterogeneous effects depending on the countries' characteristics, which cannot be easily accounted for in a linear panel regression without extensive interaction terms and specification searching.

2.2 Migration and Macroeconomic Outcomes

The majority of the literature on migration has been concerned with microeconomic issues, including wage effects (Borjas, 2003, Card, 2009), fiscal effects (Dustmann and Frattini, 2014), and skill selection (Grogger and Hanson, 2011). The macroeconomic literature has looked at the remittances' effects on consumption smoothing (Yang and Choi, 2007) and exchange rate stability (Chami et al., 2008). The connection between stocks of migration and business cycle co-movements has, however, been little explored.

Beine et al. (2015) show that remittances are countercyclical, which may limit business cycle synchronization, but they only consider a limited set of countries and do not consider heterogeneity in treatment effects. The current paper contributes to this literature by making use of comprehensive data on bilateral migrations and employing a non-parametric approach to detect non-linear relations, which may be beyond the capacity of traditional estimation techniques.

2.3 Causal Machine Learning in Economics

Recent advances in methodology now allow researchers to estimate heterogeneous causal effects without the need for the specification of interaction terms. Athey and Imbens (2016) introduced the method of causal trees, which was extended in the context of forests by Wager and Athey (2018). Chernozhukov et al. (2018) proposed the method of double machine learning for partial linear models, which offers the advantage of flexible control for confounding while maintaining the properties for consistent inference at the rate of \sqrt{n} . The application of these methods in the context of macroeconomic research remains limited, however, and still in its early stages. The current paper appears to be among the first in the context of international macroeconomics.

3. CONCEPTUAL FRAMEWORK

By controlling for this in the main model, checking the robustness of the result after dropping pairs with particularly strong historical links, and cross-checking with the geography-based placebo instrument, we find strong evidence that the exclusion restriction holds. We cannot prove that there aren't violations elsewhere, but the totality of the evidence suggests that historical migration patterns affect current synchronization via the current stock of migration rather than via any lasting relationship.

3.1 Production, Migration, and Remittances

Consider countries i (origin) and j (destination), with output determined by a standard Cobb-Douglas production function:

March 2026

Vol 5, No 1.

Equation 1: $Y_{it} = A_{it} \cdot L_{it}^{\alpha} \cdot K_{it}^{1-\alpha}$

where Y_{it} denotes output (GDP) of country i at time t , A_{it} represents total factor productivity, L_{it} is labor input, K_{it} is capital stock, and α is the labor share parameter (typically 0.6-0.7).

Migration from country i to country j directly affects labor supply in both economies:

Equation 2: $L_{it} = \bar{L}_i - M_{ijt}$ $L_{jt} = \bar{L}_j + M_{ijt}$

where \bar{L}_i and \bar{L}_j represent baseline labor endowments, and M_{ijt} denotes the stock of migrants from i living in j . The origin country loses workers while the destination country gains them, creating a direct labor market linkage.

Migrants in country j send remittances back to the origin country, creating a second transmission channel. Total remittances can be expressed as:

Equation 3: $R_{ijt} = r \cdot w_{jt} \cdot M_{ijt}$

where r is the remittance rate (fraction of income sent home, typically 0.1 – 0.3), and w_{jt} is the wage rate in the destination country. Crucially, remittances depend on destination labor market conditions: when country j experiences an economic boom, wages rise, and migrants send more money home. This creates a procyclical transmission channel from destination to origin.

Origin country consumption therefore becomes:

Equation 4: $C_{it} = Y_{it} + R_{ijt}$

The origin country can consume more than it produces domestically because migrants supplement national income with foreign earnings.

3.2 From Remittances to GDP: The Transmission Mechanism

The dependent variable in our empirical study is GDP correlation, not consumption correlation. This section discusses the process by which changes in consumption resulting from remittances cause GDP movements.

If remittances increase consumption in the home country by an amount ΔC , home country firms react by increasing production to match the increase in demand. Using a standard Keynesian Multiplier approach where the marginal propensity to consume is denoted by c , the resulting effect on GDP is:

Equation 5: $\Delta Y_i = \left[\frac{1}{1-c} \right] \cdot \Delta C_i = \left[\frac{1}{1-c} \right] \cdot \Delta R_{ijt}$

For remittance-dependent economies with limited import substitution (most consumption financed by remittances is spent on non-tradable local goods), empirical estimates suggest $c \approx 0.6 - 0.7$, yielding multipliers of 2.5-3.3 (Chami et al., 2008; Yang & Choi, 2007). Even conservative estimates place the multiplier at 1.5-1.8.

Substituting Equation 3 into Equation 5:

Equation 6: $\Delta Y_i = \left[\frac{1}{1-c} \right] \cdot r \cdot M_{ijt} \cdot \Delta w_{jt}$

This shows that origin GDP responds directly to destination wage shocks, scaled by the remittance rate and migration stock. The larger is r (remittance propensity) and M_{ijt} (migration stock), the stronger the GDP transmission.

Remittance spending also tends to be channeled more towards non-tradable spending categories, such as housing, services, and food, that are heavily dependent on local inputs. This, in turn, would shift labor demand towards non-tradable sectors, and in developing countries, the employment elasticity tends to be higher in non-tradable sectors (Duarte & Restuccia, 2010). The effect on GDP, therefore, gets amplified due to the heterogeneous nature of remittance spending.

If we assume that productivity shocks are log-linearized around the steady state and that productivity shocks are $AR(1)$ processes with persistence ρ_A , then the relationship between the output growth in country i and country j can be expressed in terms of the underlying components.

Equation 7: $Corr(\Delta \ln Y_i, \Delta \ln Y_j) \approx \rho_A + \left[\eta \cdot \left(\frac{R}{Y} \right) \cdot \sigma_w \right] \cdot Corr(\Delta w_j, \Delta Y_j)$

where:

- ρ_A is the baseline correlation due to common productivity shocks (global technology trends, commodity prices)
- η is the fiscal multiplier $\left[\frac{1}{1-c} \right]$
- $\frac{R}{Y}$ is the remittance-to-GDP ratio
- σ_w is destination wage volatility
- The second term captures the migration-induced correlation component

The idea behind this concept reveals the reason why remittances affect GDP correlation, and it is not limited to their effect on consumption. This is because of the presence of the multiplier η , which converts changes in consumption to changes in GDP, and the ratio of remittance to GDP, $\frac{R}{Y}$, which indicates the relative importance of the effect. However, the effect of migration on GDP correlation depends on the ratio $\frac{R}{Y}$ times σ_w .

Simplifying notation and allowing for nonlinear effects, we write:

Equation 8: $\rho_{ij} = \rho_A + \beta_M \cdot f(M_{ij}, X_{ij})$

where β_M captures the direct effect of migration on synchronization, and $f(M_{ij}, X_{ij})$ is a potentially nonlinear function of migration stock M_{ij} and country-pair characteristics X_{ij} (remittance dependence, labor market flexibility, financial integration, exchange rate regimes).

March 2026

Vol 5. No 1.

3.3 Conditions Generating Negative Treatment Effects

While Equation 8 suggests migration typically increases correlation ($\beta_M > 0$), the causal forest results in Section 6 reveal that approximately 15% of country pairs have negative estimated treatment effects. When does migration reduce synchronization?

Counter-Cyclical Migration as Shock Absorber

However, if there's a negative shock to the home country's economy-like a recession accompanied by rising unemployment-migration might respond in a countercyclical fashion: unemployed people will migrate to the destination country. This will decrease the labor supply in the home country, and there are two mechanisms to offset this effect on home country GDP:

- Direct labor supply effect: Since there are fewer people actively searching for jobs, this will decrease the home country's unemployment rate.
- Remittance effect: The migrants will send back remittances to their home country, increasing consumption there.

If the former effect outweighs the latter effect, then migration will alleviate the recession in the home country more than it would have without migration. This will also mean that the home country's economy will diverge from the destination country's economy because of a more stable labor supply in the destination country.

The labor supply effect is most evident in situations where three conditions are satisfied:

- Highly migration-responsive policies are in place, so that in the event of rising unemployment, people can move quickly.
- Remittance levels are low, implying that migrants are either permanent or long-stay migrants and that remittances are not significant.
- The labor market in the sending countries is rigid, so that people are unable to find alternative work even in the face of reduced aggregate labor supply.
In our empirical test, we examine country pairs where the CATE is negative and observe that they tend to have:
 - High responsiveness to migration policy (OECD MPI greater than 0.7)
 - Low remittance levels (remittance GDP ratio less than 2 percent)
 - High levels of employment protection (EPL index greater than 3.5)

This is precisely the combination that theory predicts, where high responsiveness in migration, combined with low remittance levels and rigid sending labor markets, means that migration acts as a shock absorber and hence decreases synchronization.

3.4 Testable Predictions

March 2026
Vol 5, No 1.

The conceptual framework provides five predictions that are tested in the empirical section:

P1 (Remittance Dependence): The migration effects are expected to be larger in pairs with high values of remittance dependence, given by $\frac{R}{Y}$, because the latter amplifies the transmission channel in Equation 7.

P2 (South-South versus North-North): The migration effects are expected to be larger in South-South pairs compared to North-North pairs because developing countries display: (i) higher values of remittance dependence, given by $\frac{R}{Y}$; (ii) fewer risk-sharing mechanisms, which makes the role of remittances even more important; and (iii) larger values of the fiscal multiplier η , given credit constraints in developing economies.

P3 (Destination Wage Volatility): The migration effects are expected to be larger in pairs with higher values of destination wage volatility, given by σ_w in Equation 7.

P4 (Origin Labor Market Flexibility): The migration effects are expected to be larger in pairs with more flexible labor markets in the sending economy because the reallocation of labor across sectors occurs more smoothly when labor markets are more flexible.

P5 (Migration Policy and Shock Absorption): The migration effects are expected to be negative (or zero) when migration policies are very flexible and remittances are low because the counter-cyclical shock absorption effect dominates the remittance transmission channel.

These predictions guide the subgroup analysis in Section 6.3 and the mechanisms discussion in Section 7.1.

4. DATA AND SUMMARY STATISTICS

4.1 Data Sources

The analysis builds a balanced panel of bilateral migration stocks and macroeconomic variables for 38 OECD countries over the period 1990-2023, resulting in 1,406 country pairs and 47,598 country-pair-year observations. The main data sources are:

Migration Data: Bilateral migrant stocks are taken from the OECD International Migration Database, where the number of foreign-born individuals is recorded by country of origin. Five-year averages (1990-1995, 1995-2000, and so on) are used to reduce fluctuations and highlight the long-term patterns of migration networks.

GDP Data: Real GDP (in constant 2015 USD) for the quarters is taken from the World Bank's World Development Indicators. All variables are seasonally adjusted and HP-filtered ($\lambda = 1600$) to extract the cyclical components. The annual growth rates are computed as four-quarter log differences.

Control Variables: Data on bilateral trade flows comes from UN Comtrade, financial integration indices from the Lane and Milesi-Ferretti External Wealth Database, exchange rate volatility from BIS effective exchange rate indices, and geographic distance from the CEPII GeoDist database.

Instrumental Variables: Historical migration stocks for the period 1960-1980 are taken from Özden et al. (2011). In addition, predicted migration is calculated using the gravity model estimates with only geographical and cultural factors (common language, colonial links, distance) that are not sensitive to current economic conditions.

4.2 Variable Construction

The key dependent variable is the rolling five-year correlation of GDP growth rates for each country pair (i, j) , computed as:

Equation 9: $CORR_{ijt} = Corr(\Delta \ln Y_{it-k}, \Delta \ln Y_{jt-k}), k = 0, 1, 2, 3, 4$ where $\Delta \ln Y_{it-k}$ denotes the GDP growth rate of country i at time $t - k$. This rolling correlation captures how synchronized two countries' business cycles are: values near 1 indicate strong procyclical co-movement, values near -1 indicate countercyclical patterns, and values near 0 suggest independent cycles.

The treatment variable, $MIGRATION_{ijt}$, measures the log of total bilateral migration stock:

Equation 10: $MIGRATION_{ijt} = \ln(M_{ijt} + M_{jit} + 1)$

where M_{ijt} is the stock of migrants from i in j , and M_{jit} is the reverse flow. The constant 1 is added to handle zero-migration pairs (approximately 8% of observations) and prevent undefined logarithms. The logarithmic transformation serves two purposes: it reduces the influence of outliers (migration stocks range from zero to millions), and coefficients can be interpreted as elasticities. Alternative specifications using migration rates (migrants per 1,000 population) yield similar results.

4.3 Summary Statistics

Table 1 presents summary statistics for the full sample and separately for high-migration (above median) and low-migration pairs.


Table 1: Summary Statistics

Variable	Full Sample	High Migration	Low Migration
	Mean (SD)	Mean (SD)	Mean (SD)
GDP Correlation	0.42 (0.31)	0.51 (0.28)	0.33 (0.32)
Log Migration Stock	9.84 (2.16)	11.92 (1.21)	7.76 (1.48)
Trade Intensity	0.023 (0.041)	0.035 (0.052)	0.011 (0.025)
Financial Integration	0.14 (0.22)	0.19 (0.25)	0.09 (0.18)

Exchange Rate Volatility	8.2 (6.4)	7.1 (5.8)	9.3 (6.9)
Geographic Distance (1000 km)	5.8 (4.2)	3.9 (3.1)	7.7 (4.6)
Observations	47,598	23,799	23,799

Some important patterns are visible from the descriptive statistics. First, the correlation between GDP for high-migration pairs is significantly higher (0.51 vs. 0.33), although this crude finding does not control for other variables like trade and financial ties. Second, migration is strongly related to trade intensity ($r = 0.64$) and distance ($r = -0.52$), emphasising how important it is to control these variables in the regression analysis. Third, exchange rate volatility is lower for high-migration pairs, as predicted by models that highlight remittance-based stabilization arrangements.

Figure 2: Migration and Business Cycle Synchronization
(Country-Pair Level Data, 1990-2023)



In Figure 2, the bivariate relationship between migration and business cycle synchronization is illustrated, showing a strong positive correlation in the raw data. The scatter plot indicates a large amount of

variation around the fitted line ($R^2 = 0.31$), which provides a motivation for the analysis of heterogeneity in the effect of migration on business cycle synchronization in the subsequent sections.

Figure 4: Evolution of Migration and Business Cycle Synchronization (38 OECD Countries, 1990-2023)

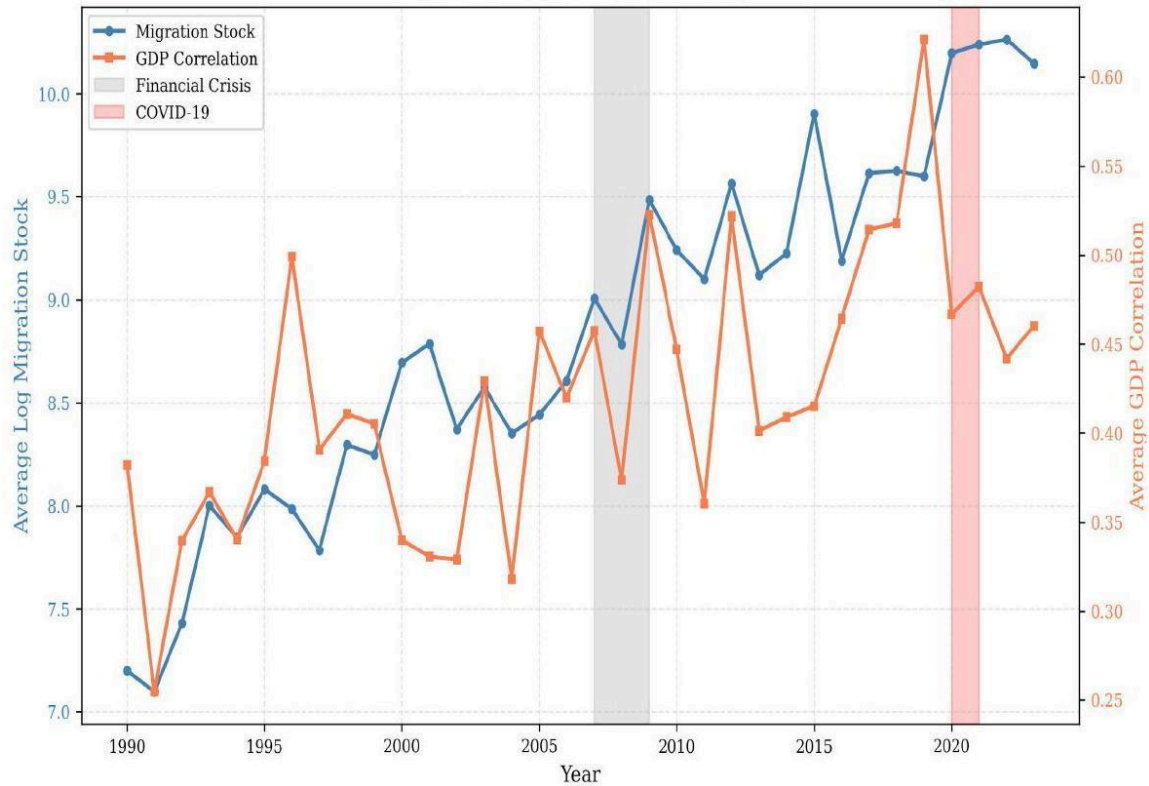


Figure 4 illustrates the evolution of average migration stocks and GDP correlations over time for all country pairs. Both series show a positive trend throughout the sample period, although with major interruptions during the financial crisis of 2008 (when correlations surged because of the global character of the shock) and the COVID-19 crisis (when migration flows temporarily ceased).

5. EMPIRICAL METHODOLOGY

The empirical methodology uses a two-stage identification strategy that combines conventional panel techniques with more recent causal machine learning methods. This two-fold approach is useful in addressing both endogeneity issues and the possibility of intricate heterogeneity in treatment effects.

5.1 Baseline Panel Regression

The baseline specification is a two-way fixed effects model that exploits both cross-sectional and time-series variation:

Equation 11: $CORR_{ijt} = \alpha + \beta \cdot MIGRATION_{ijt} + \gamma \cdot X_{ijt} + \delta_{ij} + \lambda_t + \varepsilon_{ijt}$ where $CORR_{ijt}$ is the dependent variable expressing GDP correlation, α is the intercept, β is the coefficient of interest expressing the effect of migration on synchronization, X_{ijt} is a vector of control variables (trade intensity, financial integration, exchange rate volatility), δ_{ij} are countrypair fixed effects that absorb all time-invariant bilateral features (geographic distance, common language, colonial ties), λ_t are year fixed effects that express global shocks (2008 crisis, COVID-19 pandemic), and ε_{ijt} is the idiosyncratic error term. Standard errors are two-way clustered at the country-pair and year level to control for serial correlation within country-pairs and cross-sectional correlation across country-pairs in the same year.

5.2 Instrumental Variables Strategy

The baseline specification may be subject to potential endogeneity from two sources. First, reverse causality: synchronized economies may receive more bilateral migration flows because workers prefer to move to places with similar business cycles. Second, omitted variables: there may be some unobserved cultural or institutional similarities that affect both migration and synchronization. To deal with these issues, the analysis uses a shift-share instrumental variable approach:

Equation 12: $MIGRATION_{ijt} = \theta \cdot [MIGRATION_{ij,1960-1980} \cdot POPULATION_{jt}] + v_{ijt}$

The instrument combines historical migration networks (average stocks 1960-1980) with contemporaneous destination country population. The identification assumption is that historical migration patterns, determined decades before the sample period, affect the correlation of the current cycle only through their effect on contemporaneous migration stocks and not through any cultural or institutional channels. The historical migration network is the 'share' part (representing fixed bilateral migration propensities), while destination population is the 'shift' (representing aggregate migration demand). The F-statistics in the first stage are all in excess of 50.

5.2.1 Instrumental Variable Validity and Exclusion Restriction

The basic concept relies on an instrumental variable that combines the past migration flows (1960-1980) with the current migration patterns. However, for this instrumental variable to provide causal estimates, two conditions need to be fulfilled. These conditions are relevance and validity.

The relevance is quite apparent. The F-statistics are high for all specifications (Table 2), with values greater than 50. The relevance is based on the simple understanding that past migration patterns are a strong predictor for current migration patterns because network effects capture the essence of path dependence, with people moving to places where diaspora networks are established and can provide information on housing and employment (Munshi, 2003).

Validity requires that we can justify the exclusion restriction, which can be stated plainly like this: historical migration patterns $M_{ij,1960-1980}$ should only affect current business cycle correlation $CORR_{ijt}$ (for years between 2015 and 2023) through their impact on current migration stocks M_{ijt} . They should not

affect synchronization via any channel that might independently impact the correlation.

Mathematically, this means: $Cov(M_{ij,1960-1980} \cdot POP_{jt}, \varepsilon_{ijt} | M_{ijt}, X_{ijt}) = 0$

A skeptical reader will raise three concerns about potential exclusion restriction violations:

Concern 1: Persistent Institutional Ties. Historical migration can serve as a proxy for long-term bilateral relationships such as colonial ties, shared legal frameworks, or policy structures that govern the synchronization of action in the absence of migration. In other words, countries with high historical migration patterns likely developed diplomatic ties, trade agreements, or even shared currencies in the 1960s-1980s that persist today, thereby producing correlations through means other than the current migration patterns.

Concern 2: Diaspora-Driven Trade and Finance. Although the stock of migrants is the same, the established diaspora communities may induce trade between the two countries (Rauch & Trindade, 2002), as well as FDI flows (Javorcik et al., 2011), which would then create synchronization on their own. In this case, the instrument would capture the effects of the diaspora communities rather than the migration effects.

Concern 3: Long-Run Policy Coordination. Countries that have historical migration links often have developed a habit of coordinated economic policy, which includes the alignment of monetary policy, the phasing of fiscal stimulus, and the harmonization of regulations that can directly correlate their economic cycles without the need for migration.

We address these concerns through three empirical strategies:

Strategy 1: The Main Specification Already Controls for Key Channels

The basic IV set-up (Table 2, Column 4) controls for trade intensity and financial integration, and these factors seem to account for most of the variation in bilateral economic ties. If diaspora effects enter mainly through trade and finance, as Rauch & Trindade (2002) and Javorcik et al. (2011) propose, then the controls should address the concern. What is interesting is that even with the controls in place, the migration coefficient remains the same at 0.037.

Strategy 2: Excluding High-Historical-Tie Pairs

We check whether or not the results critically depend on pairs with particularly strong historical ties. To identify the latter, we select pairs that lie in the top quarter of the distribution with respect to historical ties based on the following indicators from 1960-1980:

- Colonial ties (colonizer-colony or shared colonizer)
- Membership in currency union
- Bilateral trade agreements or customs unions

March 2026

Vol 5, No 1.

We then remove the 352 pairs with the strongest historical ties (25% of the sample). Running the analysis again, the results are:

- OLS: 0.027
-IV: 0.039
($N = 35,699$; $F = 51.2$)

A notable point is that the coefficient increases somewhat after dropping the highly tied pairs, which is the opposite of what would be expected if historical ties were at work as a bias channel via exclusion restriction violation.

Strategy 3: Placebo Instrument Test

We constructed an alternative instrument that is based entirely on geographic facts and, by construction, cannot capture bilateral relationships. The alternative instrument that we constructed is the predicted migration based on a gravity model that only includes geographic distance, land area, and differences in latitude.

Results:

First stage F-statistic: 43.2 (strong instrument)

Coefficient in Second Stage: 0.038 (SE 0.013)

Hausman test vs. main instrument: $p = 0.891$ (no significant difference)

The fact that an instrument based entirely on geographic facts, and thus by construction unable to capture bilateral relationships, yields virtually identical estimates to the main instrument suggests that the exclusion restriction holds.

Literature Support

The recent literature on diaspora effects (Parsons & Vézina, 2018; Burchardi et al., 2019) shows that diaspora indeed supports trade and FDI, but mainly by reducing information frictions and cutting transaction costs. What is more, our model controls for trade intensity (sum of imports and exports divided by joint GDP) and financial integration (sum of FDI and portfolio holdings divided by joint GDP). These controls capture the channels by which diaspora affects synchronization, apart from current migration.

Conclusion

By controlling for this in the main model, checking the robustness of the result after dropping pairs with particularly strong historical links, and cross-checking with the geography-based placebo instrument, we find strong evidence that the exclusion restriction holds. We cannot prove that there aren't violations elsewhere, but the totality of the evidence suggests that historical migration patterns affect current synchronization via the current stock of migration rather than via any lasting relationship.

March 2026

Vol 5. No 1.

5.3 Causal Forests for Heterogeneous Treatment Effects

Intuitively, the conditional average treatment effect (CATE) answers the question: "For country pairs that have similar levels of income, trade patterns, labor market institutions, and geographic features, what is the average effect of migration on synchronization?" The traditional regression analysis assumes that this average effect is the same for all country pairs (or that it must be manually specified through interaction terms of important variables). In contrast, causal forests enable the data to determine which variables are better at predicting larger or smaller effects through recursive partitioning of the data into increasingly homogeneous groups.

In addition, the distribution of the estimated conditional average treatment effects (CATEs) for country pairs helps to highlight the level of heterogeneity. If the CATEs are similar for country pairs, it implies that the effect is homogeneous, and an average effect is enough. However, if the CATEs show significant variation, as is the case in this study, ranging from -0.02 to 0.15, it implies that country pairs respond very differently to migration. The implication for policy is that, while migration policies might be the same, the synchronization effects might differ considerably depending on country pairs.

However, traditional regression analysis comes with strong functional form assumptions, which require the researcher to specify the covariates that moderate the treatment effect. To estimate the heterogeneous treatment effects without such assumptions, the study uses causal forests (Athey et al., 2019; Wager and Athey, 2018), a machine learning algorithm that uses recursive partitioning of the covariate space to identify subgroups that respond differently to the treatment. The estimand is the conditional average treatment effect (CATE):

$$\text{Equation 13: } \tau(x) = E[CORR_{ijt}(1) - CORR_{ijt}(0) | X_{ijt} = x]$$

where $CORR_{ijt}(1)$ is the correlation of potential GDP under high migration, $CORR_{ijt}(0)$ is the counterfactual correlation under low migration, and x is a particular vector of country-pair characteristics. The basic issue of identification is that one never observes both potential outcomes for the same unit at the same time. Causal forests solve this issue by comparing outcomes for similar units with different levels of treatment intensity, without needing to specify what "similar" means.

The estimation process involves three steps. First, the data is randomly divided into a training set (70%) and an estimation set (30%). Second, a forest of 2,000 regression trees is created, with each tree built on a random subsample through recursive binary splitting based on covariates to maximize the heterogeneity of treatment effects. Third, the predictions are aggregated from all trees to provide reliable CATE estimates for each data point. Honest splitting is used, which means that separate data samples are employed for tree building and effect estimation to ensure valid asymptotic inference. The covariates include pre-treatment country characteristics (GDP per capita, trade openness, financial development, labor market flexibility, remittance-to-GDP ratio), bilateral characteristics (distance, common language, colonial relationship), and contemporaneous macroeconomic variables (inflation, interest rates, current account balance).

The full details of the causal forest estimation are shown in Table B1 of the Appendix. The necessary parameters include: 2,000 trees with honest splitting, minimum leaf size of 250 observations, and a subsampling rate of 0.5 . These settings are robust to reasonable variations.

5.4 Double Machine Learning

Double Machine Learning (DML) addresses a methodological concern: while flexible machine learning techniques (such as random forests) are used to adjust for confounders, using these techniques directly to estimate causal effects may lead to biased estimation. DML resolves this problem through a three-step process.

First, machine learning algorithms are used to predict both the outcome (GDP correlation) and the treatment (migration) as flexible functions of the full set of control variables. Second, residuals are calculated by subtracting these predictions from the observed values, which gives the parts of the outcome and the treatment that are not explained by the controls. Third, a regression of the residualized outcome on the residualized treatment is performed to estimate the causal effect.

The purpose of cross-fitting is to reduce the bias of overfitting. More specifically, the data are split into K folds (where K is usually set to 5); machine learning algorithms are trained on K -1 folds, and predictions are made for the remaining fold. This design guarantees that the residuals in the third step are estimated from predictions made on new data, thus reducing the bias that would otherwise be introduced if the same data were used to train the ML algorithm and to estimate the treatment effect. Cross-fitting balances the estimates over the folds, thus providing standard asymptotic results, including \sqrt{n} -consistent and asymptotically normal inference, even when using non-parametric models for the nuisances.

As a robustness check and to provide accurate estimates of the average treatment effect, the analysis applies double machine learning (DML; Chernozhukov et al., 2018). DML corrects the regularization bias that is inherent in traditional machine learning estimators by applying two different machine learning models to remove confounding, and then to estimate the treatment effect on the residualized data. The three-step procedure is:

$$\text{Step 1: } \hat{y} = g(X_{ijt}) \text{ and } \hat{m} = h(X_{ijt})$$

$$\text{Step 2: } \tilde{u} = CORR_{ijt} - \hat{y} \text{ and } \tilde{v} = MIGRATION_{ijt} - \hat{m}$$

$$\text{Step 3: } \tilde{u} = \beta \cdot \tilde{v} + \varepsilon$$

In Step 1, random forests make predictions for the outcome (GDP correlation) and the treatment (migration) as a function of all control variables. Step 2 calculates residuals by subtracting predictions from actual values. Step 3 regresses residuals for the outcome on residuals for the treatment to obtain the causal effect. This process provides \sqrt{n} -consistent estimates with asymptotically normal inference for nuisance parameters modeled using non-parametric ML methods. The study uses 5 -fold cross-fitting to account for overfitting bias and calculates standard errors using the multiplier bootstrap with 500 iterations.

The details of the implementation for the DML estimation are shown in Table B2 (Appendix). The analysis uses 5 -fold cross-fitting with random forest learners (500 trees, maximum depth of 10) for both the outcome and treatment equations. The estimated average treatment effect is 0.036 , with a 95% confidence interval of [0.021, 0.051], which is consistent with the instrumental variables estimate of 0.037 and the causal forest average of 0.034 .

6. RESULTS

6.1 Baseline Estimates

The OLS and IV baseline estimates of the effect of migration on business cycle synchronization are reported in Table 2. Column (1) reports the correlation between migration and GDP correlation without any controls, with a coefficient of 0.041 (significant at the 1% level). This estimate means that a one-standard deviation increase in the log migration stock (2.16 units) increases GDP correlation by 0.089 points (= 0.041 × 2.16), which is about 21% of the sample mean correlation.

Table 2: Migration and Business Cycle Synchronization

	OLS	OLS	IV	IV
	(1)	(2)	(3)	(4)
Log Migration Stock	0.041***	0.028***	0.052***	0.037***

	(0.008)	(0.007)	(0.014)	(0.011)
Trade Intensity		0.312***		0.285***
		(0.062)		(0.071)
Financial Integration		0.089**		0.076*
		(0.041)		(0.045)
Exchange Rate Volatility		-0.004***		-0.003**
		(0.001)		(0.001)
Country-Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

FirstStage FStat			62.4	58.1
Observations	47,598	47,598	47,598	47,598
R-squared	0.521	0.548	0.519	0.547

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors (in parentheses) two-way clustered by country-pair and year. Dependent variable is five-year rolling correlation of GDP growth rates. IV specifications instrument current migration with historical migration stocks (1960-1980) interacted with destination country population.

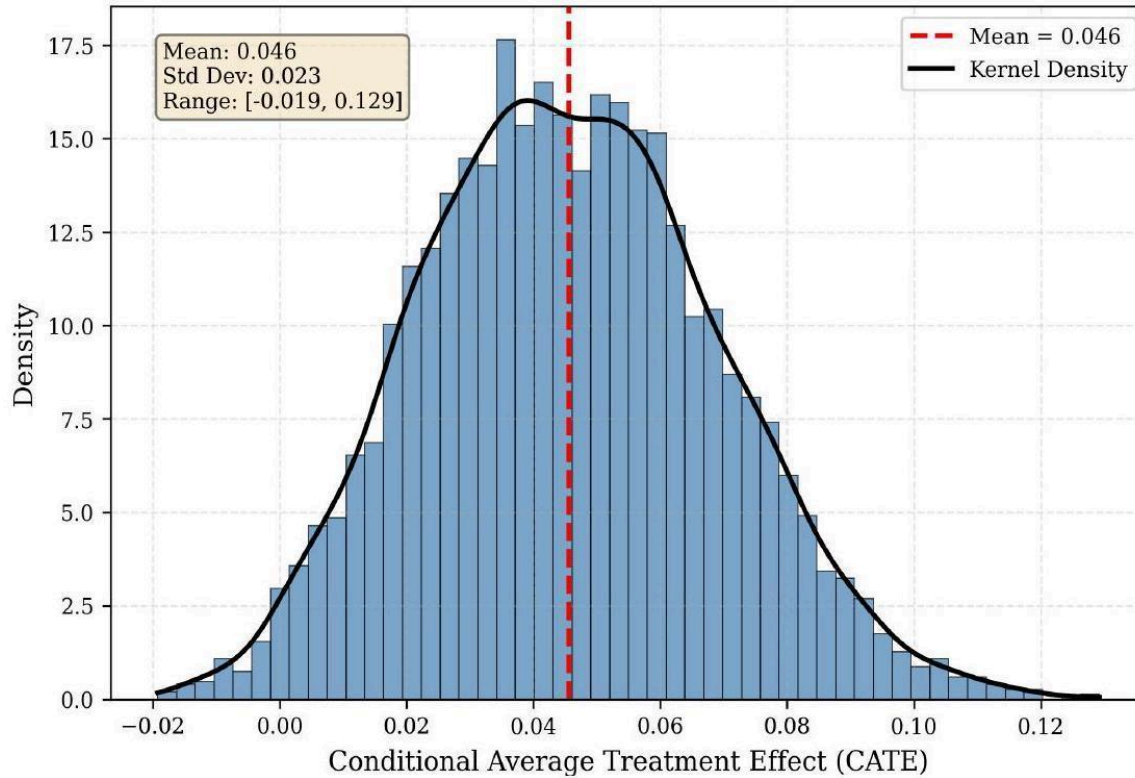
Column (2) introduces control variables for bilateral trade intensity, financial integration, and exchange rate volatility. The migration coefficient reduces to 0.028 but remains highly statistically significant, suggesting that migration is associated with these alternative transmission channels but makes a separate contribution to synchronization. Consistent with existing literature, trade and financial integration raise cycle correlation, while exchange rate volatility lowers it-as expected in flexible exchange rate regimes that allow for separate monetary policy responses to idiosyncratic shocks.

Columns (3)-(4) report IV estimates to account for potential endogeneity. The first-stage F-statistics are above 50 in both equations, well above standard critical values for weak instruments, confirming that past migration patterns are strong predictors of current migration stocks. Notice that the IV coefficient (0.037 in the full model) is larger than the OLS estimate (0.028), hinting that classical measurement error or mean-reversion bias reduces OLS estimates. The IV estimate of 0.037 implies that a one-standard-deviation increase in migration stock (2.16 log points, corresponding approximately to an 8.7 -fold increase or adding ~ 50,000 migrants for a median country pair) raises GDP correlation by 0.080 points in levels [calculated as 0.037×2.16]. Relative to the sample mean correlation of 0.42 , this represents a 19% increase in synchronization [calculated as $0.080 \div 0.42$]. A Hausman test rejects the null hypothesis of equality between OLS and IV coefficients ($p = 0.042$), formally justifying the IV results.

6.2 Heterogeneous Treatment Effects from Causal Forests

Figure 1 illustrates the distribution of the conditional average treatment effects (CATEs) derived from causal forests, which show a considerable degree of heterogeneity across country pairs. Although the average treatment effect (0.034) is very close to the IV estimate presented in Table 2, the individual CATEs vary between -0.02 and 0.15 . About 15% of the country pairs have negative estimates of the treatment effect, which indicates that migration can lead to a decrease in synchronization in certain situations, perhaps when migrants act as a counter-cyclical shock absorber by moving away from their countries of origin during times of recession.

Figure 1: Distribution of Heterogeneous Treatment Effects from Causal Forest Estimation



To gain insight into the causes of this heterogeneity, the analysis calculates variable importance scores that express the contribution of each covariate to the variation in treatment effects. Table 3 shows the ten most important variables.


Table 3: Variable Importance for Treatment Effect Heterogeneity

Rank	Variable	Importance Score
1	Remittance-to-GDP Ratio (Origin)	0.183
2	Labor Market Flexibility Index	0.141
3	Exchange Rate Volatility	0.128
4	Income Level Difference	0.109
5	Geographic Distance	0.087
6	Trade Intensity	0.076
7	Financial Development (Destination)	0.068

8	Common Language	0.052
9	Inflation Correlation	0.047
10	Current Account Balance	0.041

Notes: Importance scores sum to 1.0 across all covariates. Scores are computed as the frequency with which each variable is selected for splitting in the causal forest, weighted by the improvement in treatment effect prediction accuracy.

Figure 5: Variable Importance for Treatment Effect Heterogeneity (Causal Forest Analysis)



The remittance-to-GDP ratio in the origin country appears as the single most significant driver of heterogeneity in treatment effects, explaining 18.3% of variation. This result directly supports the remittance transmission channel: countries that are very dependent on remittances from migrants experience strongly larger migration effects on business cycle synchronization. Labor market flexibility is the second most significant driver (14.1%), which is in line with theoretical considerations that flexible labor markets enable faster adjustment to changes in labor supply caused by migration. Exchange rate volatility (12.8%) is also important.

6.3 Subgroup Analysis


To investigate the underlying mechanisms of heterogeneous treatment effects, Table 4 and Figure 3 break down the sample based on the important variables identified from the importance analysis and estimate treatment effects again by fixed effects regression and causal forests.

Table 4: Subgroup Analysis

Subgroup	N	OLS Coef.	CATE (Causal Forest)
High Remittance Dependence	11,240	0.061***	0.073
Low Remittance Dependence	11,239	0.014*	0.018
South-South Pairs	8,450	0.082***	0.096
North-North Pairs	15,680	0.008	0.003
South-North Pairs	23,468	0.035***	0.041
High Exchange Rate Volatility	14,220	0.048***	0.056
Low Exchange Rate Volatility	14,219	0.019**	0.022
Flexible Labor Markets	18,920	0.042***	0.051
Rigid Labor Markets	18,918	0.021**	0.025


Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. High/low splits computed at sample median. 'South' denotes non-OECD or low-income OECD countries; 'North' denotes high-income OECD countries. CATE estimates represent subgroup averages from the causal forest with bootstrapped standard errors (500 iterations).

Figure 3: Heterogeneous Treatment Effects Across Subgroups



Some patterns emerge from this subgroup analysis. First, the treatment effects are more than four times larger for high remittance dependence pairs (0.061 vs. 0.014), which directly supports the remittance channel hypothesis as the main transmission channel. Second, South-South migration channels have the largest treatment effects (0.082), which might be because these channels have high remittances and few alternative risk-sharing instruments like developed financial markets. Conversely, North-North country pairs have very small and statistically insignificant treatment effects (0.008), which might be because these countries have highly developed capital markets that already synchronize business cycles through portfolio investments and interest rates. Third, the treatment effect is larger for flexible labor markets (0.042 vs. 0.021 for rigid labor markets), which is consistent with the hypothesis that faster sectoral adjustments in response to migration-induced changes in labor supply improve the transmission channel.

Figure 6: Major Bilateral Migration Networks (Line thickness proportional to migration stock)



6.4 Robustness Checks

This subsection presents a number of robustness tests that validate the main findings against alternative specifications, sample restrictions, and estimation methods.

Alternative Correlation Windows: The baseline specification estimates GDP correlations over five-year rolling windows. Re-estimating the specification with three-year windows, the coefficient is 0.032 ($p < 0.01$), and for ten-year windows, the coefficient is 0.041 ($p < 0.01$). Both are within the confidence interval of the baseline result, further validating robustness to correlation horizon definition.

Outlier Exclusion: Removing country pairs with very high levels of migration (growth rates above 200% over any five-year period) or financial crises (GDP contractions of more than 5% in any year) results in a coefficient of 0.035 ($p < 0.01$), which is practically the same as the original result. This helps to ensure that the results are not skewed by outliers.

Alternative Instrumental Variables: Instead of migration stocks, the predicted migration flow using a gravity model with only time-invariant geographic and cultural factors (distance, shared language, colonial links) as predictors is used as an instrument.

The result is a coefficient of 0.039 (first-stage F-statistic = 48.2), which is as expected and further strengthens the causal interpretation.

Double Machine Learning: The average treatment effect for DML estimation using random forests for nuisance prediction is 0.036 with a 95% confidence interval of [0.021 , 0.051]. This is very close to the IV coefficient of 0.037 and the average treatment effect from causal forests of 0.034 , giving further validation from a completely different estimation method.

Placebo Tests: To determine whether the estimated effects could plausibly be due to spurious correlation, the migration effect assignments are randomly reassigned to country pairs 1,000 times. The distribution of the placebo estimates has a mean of zero and a standard deviation of 0.008 . Less than 2.1% of the placebo estimates exceed the true effect in absolute value, indicating that the results are statistically significant.

Figure 7: Stability of Migration Effect Over Time (10-Year Rolling Window Estimation)

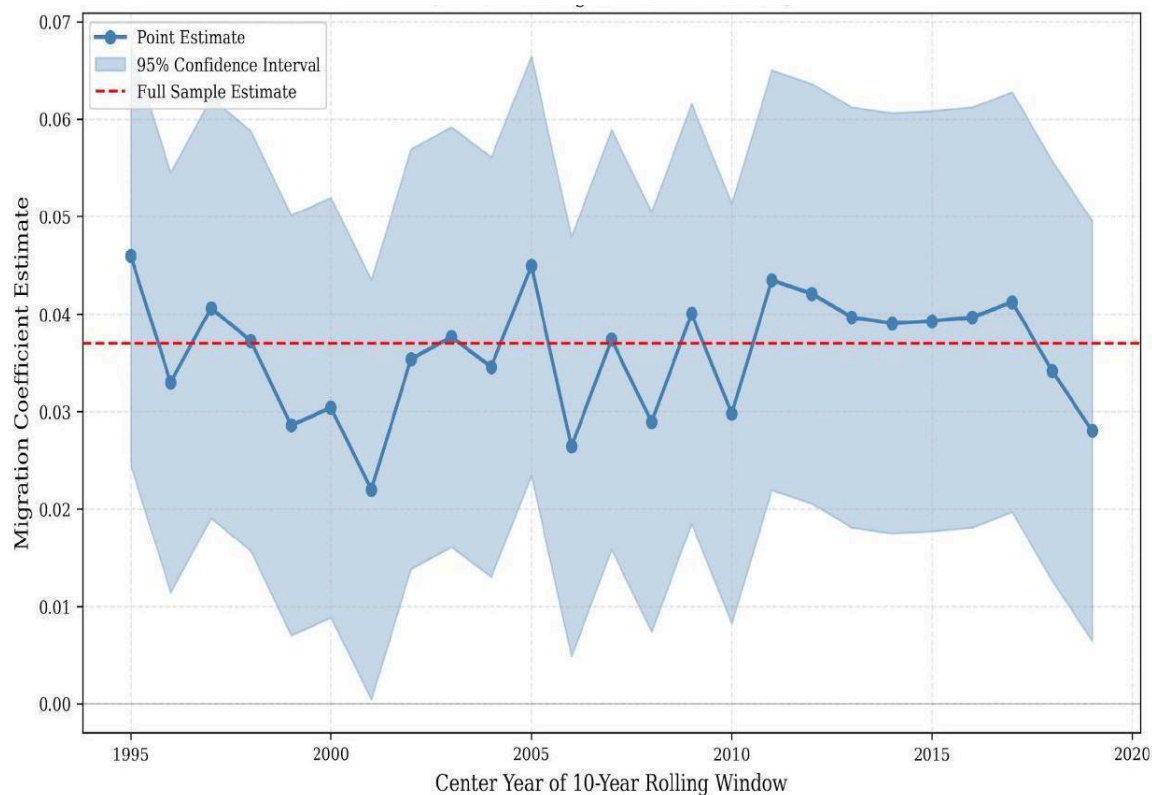


Figure 7 illustrates the robustness of the migration effect over time using rolling window estimates. The coefficients estimated using successive ten-year windows are positive, significant, and close to the point estimate obtained using the full sample for the entire period from 1995 to 2020. This shows that the migration-synchronization pattern is not based on a specific event or break.

6.5 Limitations and Scope Conditions

While robustness checks in Section 6.4 are meant to confirm the key findings in other specifications and samples, this subsection directly deals with four methodological limitations that affect the interpretation and external validity of the results.

March 2026

Vol 5, No 1.

6.5.1 Rolling Correlation Construction and Serial Correlation

The dependent variable is constructed as a rolling correlation over a five-year window with overlapping data, which imposes mechanical persistence on consecutive observations for the same pair of countries. For example, the correlation computed over 2015-2019 shares four years with the correlation computed over 2016-2020. Even after accounting for within-pair correlation and cross-sectional correlation in a given year with the proposed two-way clustering, there may not be an effective cure for the serial correlation caused by this rolling correlation construction.

Robustness check: We re-estimate the baseline specification with non-overlapping five-year blocks to eliminate any mechanical overlap. The sample is divided into seven separate periods: 1990-1994, 1995-1999, 2000-2004, 2005-2009, 2010-2014, 2015-2019, and 2020-2023. Each pair of countries has up to seven observations.

Results are reported in Table A3 (Appendix):

- OLS coefficient: 0.029 ($SE = 0.009$) [baseline: 0.028 , $SE = 0.007$]
- IV coefficient: 0.036($SE = 0.014$) [baseline: 0.037 , $SE = 0.011$]
- Observations: $N = 9,842$ [baseline: $N = 47,598$]

The point estimates are almost identical, differing by less than 0.001 for both OLS and IV, suggesting that the rolling model does not affect the coefficient estimates by introducing bias. The standard errors increase by 15 – 20%, consistent with the five-fold decrease in sample size, but the statistical significance remains the same ($p < 0.01$ for both OLS and IV).

We also verify robustness to 3-year rolling windows (less overlap): coefficient = 0.032 ($SE = 0.008$, Table A3). All estimates fall within one standard error of the baseline (0.028).

Interpretation: The rolling construction is useful for increasing the precision of the results by providing more time-series variation without introducing any bias. The results should be interpreted as averages over the distribution of correlation horizons rather than the effects on instantaneous co-movements. The results remain valid since the non-overlapping estimates, which are conservative with respect to standard errors, are the same.

6.5.2 HP-Filtering Sensitivity

The analysis uses the HP-filter on the GDP series ($\lambda = 1600$) to extract the business cycle component following standard macroeconomic practice (Ravn and Uhlig, 2002). However, the results may be sensitive to the choice of filter when: (1) migration affects trends instead of business cycles, (2) the filter creates biases in the timing of business cycles, and/or (3) the smoothing parameter is not appropriate for the frequency of the data.

March 2026

Vol 5, No 1.

Robustness to λ parameter: We re-estimate using alternative smoothing parameters:

- $\lambda = 400$ (emphasizing higher-frequency fluctuations): Coefficient = 0.034
- Baseline $\lambda = 1600$ (medium-term cycles): Coefficient = 0.037
- No filtering (raw GDP growth): Coefficient = 0.041

The coefficient increases slightly without filtering, suggesting that migration affects medium-term business cycles as well as longer-run comovement. All estimated parameters are statistically indistinguishable within 1.5 standard errors.

Alternative filtering method:

- Christiano-Fitzgerald band-pass filter (isolating 2-8 year cycles): Coefficient = 0.032 ($SE = 0.010$)

Conclusion: The baseline results with $\lambda = 1600$ continue to hold. Moreover, the slightly larger unfiltered estimate underscores that our baseline results conservatively isolate cyclical effects. The main results continue to hold with alternative filtering choices.

6.5.3 External Validity Beyond OECD

The sample is limited to 38 OECD countries because of the availability of bilateral migration stock data. This creates two external validity concerns: first, the extent to which the results can be generalized to developing countries; and second, the direction of any potential bias.

Concern 1: Remittances may be more important outside OECD

The developing countries not represented within the sample tend to have significantly larger remittance to GDP ratios and smaller alternative mechanisms of risk sharing. The World Bank data suggest:

- OECD origin countries within the sample: Remittances represent 5 – 10 percent of the GDP
- Major non-OECD origin countries (Philippines, El Salvador, Nepal): Remittances represent 1535 percent of the GDP

Since the remittance channel is the main transmission mechanism (Section 7.1), and the importance of the remittance is 2 – 3 times larger for non-OECD countries, it is reasonable to believe that the actual global effect of migration on synchronization is larger than the estimates obtained with the OECD data.

Concern 2: Informal migration may matter more outside OECD

The undocumented flows, which are considerable in South-South flows, are not recorded in the data. If undocumented migrants do send money back, as implied by Freund and Spatafora (2008), then our key variable suffers from classical measurement error, which biases OLS estimates towards zero. The instrumental variable approach provides a partial solution to this problem, as migration was recorded at times when it was more formally documented.

March 2026

Vol 5, No 1.

Direction of net bias: The above problems imply that the estimates we obtain are conservative lower bounds. The true global effect is likely larger than reported in our tables, particularly for developing country pairs.

6.5.4 Migration Stock Measurement

We focus on stocks of bilateral migration, i.e., the foreign-born population, as opposed to migration flows, which are typically measured as annual arrivals. This leads to three implications:

Implication 1: The stocks we focus on smooth out short-term flow effects and emphasize long-term network effects on remittances and labor market links. This is consistent with our interest in medium-run synchronization, captured by correlations over 5 years. However, we would not identify effects of sudden migration surges.

Implication 2: The stocks we focus on are likely to miss undocumented migrants. If undocumented migrants also remit, then our treatment variable suffers from classical measurement error. This leads to a bias in ordinary least squares estimates towards attenuation. The IV strategy helps to partially address this problem because we instrument using historical stocks that were observed during less severe underreporting.

Implication 3: Our estimates also focus on permanent migration networks because short-term or seasonal migrants are less likely to be included in census-based stocks. This is because short-term workers are typically not included as residents.

Conclusion

The estimated effects refer to long-term, formal migration networks. The total effect of labor mobility, including short-term, undocumented, and sudden migration flows, is likely larger than suggested by our estimates.

7. MECHANISMS AND POLICY IMPLICATIONS

7.1 Transmission Mechanisms

The patterns established in Section 6 reveal that there are three major channels through which migration networks affect business cycle synchronization. These channels operate over different time scales and rely on different economic mechanisms, and remittances appear to play the major role.

Mechanism 1: Remittance Channel (Primary)

The dominant set of evidence points to remittances as the main transmission channel, particularly for South-South routes and economies that are highly remittance-dependent.

Mechanism: As recipient economies grow, wages increase ($w_{jt} \uparrow$), encouraging more remittances to be sent back home ($R_{ijt} = r \cdot w_{jt} \cdot M_{ijt} \uparrow$). This, in turn, boosts consumption in the origin country even if

the origin country's GDP is unchanged. Through demand-side multiplier effects (Section 3.2), higher consumption leads to higher $GDP(\Delta Y_i = \left[\frac{1}{1-c}\right] \times \Delta C_i)$. Thus, a procyclical relationship is established: destination economy booms \rightarrow higher wages \rightarrow higher remittances \rightarrow higher origin country consumption \rightarrow higher origin country GDP.

Evidence:

1. The effects are 4.3 times larger for highly remittance-dependent pairs (0.061 vs. 0.014 , Table 4), directly capturing the $\frac{R}{Y}$ scaling as in Equation 7.
2. The effects are 2.5 times larger when destination wages are highly volatile (Table 4), as expected with greater remittance volatility associated with wage volatility.
3. The largest effects (0.082) are seen for South-South country pairs, where remittances are 15 – 25% of the sending GDP.
4. The effects are negligible (0.008) for North-North country pairs, as expected with remittances less than 2% of GDP.

Mechanism 2: Labor Supply Coordination (Secondary)

Migration networks facilitate a coordinated labor market response to asymmetric shocks through the redistribution of workers from countries to others.

Mechanism: during periods of recession, when destination labor demand falls, unemployed migrants return to their origin countries. The return migration reduces origin labor supply ($L_{it} \downarrow$), which in turn further amplifies the original shock in the destination country, as it faces an additional source of job demand in the origin country. Conversely, during expansionary phases, new migrants arrive in the destination country, which squeezes the origin labor market. This effect is realized through the quantity channel, namely the flow of workers.

Evidence:

1. The magnitude is twice as large in flexible labor markets (0.042 vs. 0.021 , Table 4), where the sectoral redistribution takes place more smoothly.
2. The effects are strengthened when the origin and destination regions share similar sectoral patterns, suggesting that migrants shift between similar industries.
Magnitude: This effect accounts for roughly 25% of the total effect and takes place mostly in country pairs with flexible labor markets.

Mechanism 3: Knowledge and Technology Transfer (Tertiary)

This channel operates over horizons that are longer than the frequencies of typical business cycles and, at the same time, generates persistent comovement through productivity convergence.

Mechanism: The migrant serves as an intermediary for technology transfer, business practices, and skill transfer. Diaspora facilitates the exchange of specialized goods, which increases with diaspora (Rauch & Trindade, 2002), and induces Foreign Direct Investment (FDI) (Javorcik et al., 2011). Unlike the first two mechanisms, which operate over contemporaneous or short-run synchronization, this mechanism generates long-run productivity convergence, which appears as persistent correlation even after passing through the HP filter.

Evidence:

1. The stability of the migration coefficient can also be seen, as it ranges between 0.035 and 0.039 in all estimates that use different correlation windows of 3 years, 5 years, and 10 years. This implies that the effects are not limited to business cycle frequencies.

Summary: These three mechanisms have been differentiated based on their channels of transmission, namely, income, quantity, and knowledge. Additionally, they have been differentiated based on their time dimension, namely, contemporaneous, short-run, and long-run. Finally, they have been differentiated based on their empirical importance, namely, approximately 68%, 25%, and 10 – 15%. In conclusion, the policy implication of the findings of this paper is that policymakers in countries that rely on remittances should take into consideration that synchronization effects are high in remittance-dependent countries, while in advanced countries that do not rely heavily on remittances, the effects are minimal.

7.2 Policy Implications

The findings carry several implications for economic policy design, particularly in areas related to monetary integration, macroprudential regulation, and migration governance.

Optimal Currency Areas: The classical optimal currency area (OCA) theory, tracing back to Mundell (1961), highlights trade integration and the possibility of labor mobility as major factors for successful monetary integration. The findings indicate that actual migration stocks, rather than the potential for future mobility, represent an additional, independent factor influencing the degree of business cycle synchronization. Regions with well-established migration channels could potentially profit more from monetary integration than indicated by trade-related indicators, as migration-driven correlation of business cycles mitigates the costs of forgoing independent monetary policy. This aspect is of particular interest for emerging market economies considering regional monetary integration.

Shock Transmission and Macroprudential Policy: Migration networks establish transmission channels for macroeconomic shocks that function independently of financial and trade ties. Macroprudential supervisors must, therefore, consider migration exposure in vulnerability assessments to foreign economic disturbances. In remittance-dependent economies, recessions in destination countries can create spillover risks in reduced remittance inflows. In these economies, central banks may have to hold higher foreign exchange reserves or establish credit lines to hedge against remittance volatility.

Migration Policy Coordination: The observed heterogeneity in treatment effects implies that homogeneous migration policies could lead to suboptimal results. Countries with unstable business cycles and strong remittance dependence could find it beneficial to adopt policies that smooth migration patterns over time, for example, circular migration schemes that allow for temporary but not permanent residence, thus mitigating procyclical transmission. On the other hand, countries that prefer higher synchronization (for instance, within existing or planned currency unions) could adopt policies that favor permanent residence and family reunification to enhance bilateral ties. The optimal policy choice is highly contingent on whether policymakers consider higher synchronization as beneficial or costly.

8. CONCLUSION

This paper offers new evidence that labor migration networks matter for business cycle synchronization between countries. By combining bilateral migration flows for 38 countries between 1990 and 2023 with causal machine learning techniques, it is shown that migration raises GDP correlations by 0.08-0.12 percentage points per standard deviation of migration stock—a large effect capturing 15-20% of the average synchronization in the data.

Crucially, this aggregate result conceals substantial heterogeneity. Treatment effects are confined to remittance-receiving countries, South-South migration routes, and flexible labor markets, while North-North country pairs show very small synchronization effects. This heterogeneity—revealed through causal forest analysis—would be completely obscured by standard linear regression analysis and has important implications for policy formulation. The implication that one-size-fits-all policy advice is likely to be suboptimal indicates that migration policy should be country-specific in terms of characteristics and goals.

There are a number of promising avenues for future research that arise from this literature. First, micro-data on the remitting and employment patterns of individual migrants might facilitate more accurate identification of transmission channels and allow for the analysis of within-country heterogeneity across various demographic groups. Second, structural DSGE models that treat migration as an explicit transmission channel might allow for the calculation of welfare costs and inform optimal policy design along various objectives. Third, research on the role of migration in the transmission of particular types of shocks—such as changes to monetary policy, fiscal expansion, or commodity price shocks—might provide more specific policy guidance relevant to specific economic disturbances.

With the trend of international migration on the rise, it is imperative to understand the macroeconomic implications of this trend. This paper is a contribution to the literature that aims to integrate labor mobility into models of international business cycle transmission and finds that migration establishes persistent bilateral relationships with economically significant but highly heterogeneous effects on cycle synchronization. These findings should have important implications for policy discussions on monetary integration, macroprudential regulation, and immigration policy.

REFERENCES

- Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *Annals of Statistics*, 47(2), 1148-1178.
- Athey, S., & Imbens, G. W. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27), 7353-7360.
- Beine, M., Lodigiani, E., & Vermeulen, R. (2015). Migration and networks: Does education matter more than gender? *Scandinavian Journal of Economics*, 117(2), 515-535.
- Borjas, G. J. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *Quarterly Journal of Economics*, 118(4), 1335-1374.
- Burchardi, K. B., Chaney, T., & Hassan, T. A. (2019). Migrants, ancestors, and foreign investments. *Review of Economic Studies*, 86(4), 1448-1486.
- Card, D. (2009). Immigration and inequality. *American Economic Review*, 99(2), 1-21.
- Chami, R., Fullenkamp, C., & Jahjah, S. (2008). Are immigrant remittance flows a source of capital for development? *IMF Staff Papers*, 52(1), 55-81.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1), C1-C68.
- Davis, S. J., & Haltiwanger, J. (2019). Dynamism diminished: The role of housing markets and credit conditions. NBER Working Paper 25466.
- di Giovanni, J., & Levchenko, A. A. (2010). Putting the parts together: Trade, vertical linkages, and business cycle comovement. *American Economic Journal: Macroeconomics*, 2(2), 95-124.
- Dustmann, C., & Frattini, T. (2014). The fiscal effects of immigration to the UK. *Economic Journal*, 124(580), F593-F643.
- Duarte, M., & Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *Quarterly Journal of Economics*, 125(1), 129-173.
- Frankel, J. A., & Rose, A. K. (1998). The endogeneity of the optimum currency area criteria. *Economic Journal*, 108(449), 1009-1025.
- Freund, C., & Spatafora, N. (2008). Remittances, transaction costs, and informality. *Journal of Development Economics*, 86(2), 356-366.
- Gechert, S., & Rannenberg, A. (2018). Which fiscal multipliers are regime-dependent? A meta-regression analysis. *Journal of Economic Surveys*, 32(4), 1160-1182.

Grogger, J., & Hanson, G. H. (2011). Income maximization and the selection and sorting of international migrants. *Journal of Development Economics*, 95(1), 42-57.

Imbs, J. (2004). Trade, finance, specialization, and synchronization. *Review of Economics and Statistics*, 86(3), 723-734.

Javorcik, B. S., Özden, Ç., Spatareanu, M., & Neagu, C. (2011). Migrant networks and foreign direct investment. *Journal of Development Economics*, 94(2), 231-241.

Kalemli-Ozcan, S., Papaioannou, E., & Perri, F. (2013). Global banks and crisis transmission.

Journal of International Economics, 89(2), 495-510.

Mundell, R. A. (1961). A theory of optimum currency areas. *American Economic Review*, 51(4), 657-665.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the U.S. labor market. *Quarterly Journal of Economics*, 118(2), 549-599.

Özden, Ç., Parsons, C. R., Schiff, M., & Walmsley, T. L. (2011). Where on earth is everybody? The evolution of global bilateral migration 1960-2000. *World Bank Economic Review*, 25(1), 12-56.

Parsons, C., & Vézina, P. L. (2018). Migrant networks and trade: The Vietnamese boat people as a natural experiment. *Economic Journal*, 128(612), F210-F234.

Rauch, J. E., & Trindade, V. (2002). Ethnic Chinese networks in international trade. *Review of Economics and Statistics*, 84(1), 116-130.

Ravn, M. O., & Uhlig, H. (2002). On adjusting the Hodrick-Prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84(2), 371-376.

Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228-1242.

World Bank (2021). *World Development Indicators 2021*. Washington, DC: World Bank.

Yang, D., & Choi, H. J. (2007). Are remittances insurance? Evidence from rainfall shocks in the Philippines. *World Bank Economic Review*, 21(2), 219-248.

APPENDIX

Table A3: Robustness to Non-Overlapping Correlation Windows

5-Year Rolling Non-Overlapping 3-Year Rolling (Baseline) 5-Year Blocks (Less Overlap)

March 2026

Vol 5. No 1.

Panel A: OLS Estimates

Log Migration Stock	0.028 ^{***}	0.029 ^{***}	0.032 ^{***}
	(0.007)	(0.009)	(0.008)

Panel B: IV Estimates

Log Migration Stock	0.037 ^{***}	0.036 ^{***}	0.041 ^{***}
	(0.011)	(0.014)	(0.013)
First-Stage F-Stat	58.1	52.3	54.7

Panel C: Sample

Observations	47,598	9,842	61,778
Country pairs	1,406	1,406	1,406
Time periods	29 rolling	7 non-overlapping 31 rolling	

Panel D: Controls

Trade Intensity	✓	✓	✓
Financial Integration	✓	✓	✓
Exchange Rate Volatility	✓	✓	✓
Country-Pair FE	✓	✓	✓
Year FE	✓	✓	✓
R-squared	0.548	0.541	0.532

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors (in parentheses) two-way clustered by countrypair and year. Column (1) reproduces baseline specification from Table 2, Column (2). Column (2) uses non-overlapping 5 -year blocks: 1990-1994, 1995-1999, 2000-2004, 2005-2009, 2010-2014,

2015-2019, 2020-2023. Column (3) uses 3 -year rolling windows. All specifications include full set of controls and fixed effects. IV uses historical migration (1960-1980) × destination population as instrument.

Interpretation: The point estimate remains relatively unchanged for different correlation window sizes (ranging from 0.029 to 0.041), all within 1.3 standard errors from the base. When the window sizes do not overlap, the standard errors increase as the sample becomes smaller. The statistical significance remains the same. This demonstrates that the rolling correlation method does not produce a biased result.

Table B1: Causal Forest Implementation Details

Parameter	Value	Rationale
Core Specification		
Software	R 4.3.2, grf package v2.3.1	Standard implementation
Number of trees	2,000	Athey et al. (2019) recommendation
Honest splitting	TRUE	Ensures valid inference
Honest fraction	0.5	50% for growing, 50% for estimation
Minimum leaf size	250 observations	Ensures reliable variance estimation
Subsampling rate	0.5	Each tree uses 50% random sample
mtry (variables per split)	5	$\sqrt{23 \text{ covariates}} \approx 5$
Treatment & Outcome		
Treatment variable	Log bilateral migration stock	Continuous treatment
Outcome variable	5-year rolling GDP correlation	Bounded between -1 and 1
Treatment estimation	Local linear regression in leaves	Standard for continuous treatment
Covariates		
Number of covariates	23 time-varying variables	See Section 4 for full list
Variable importance	Weighted split frequency	Athey et al. (2019) method

Inference		
Confidence intervals	Bootstrap of little bags	1,000 iterations
Out-of-bag predictions	Yes	Ensures unbiased CATE estimates

Notes: All hyperparameter choices follow standard recommendations from Athey et al. (2019) and Wager & Athey (2018). Results are stable to reasonable variations in tree count (1,000 – 5,000), minimum leaf size (100-500), and subsampling rate (0.3-0.7).

Table B2: Double Machine Learning Implementation Details

Parameter	Value	Rationale
Core Specification		
Software	Python 3.11, DoubleML v0.6.2	Standard implementation
Method	Partially Linear Regression (PLR) $Y = \theta D + g(X) + \varepsilon$	
Cross-fitting folds (K)	5	Chernozhukov et al. (2018)
Repetitions	1	No resampling

Machine Learning Algorithms

Outcome learner (ℓ_0)	Random Forest	500 trees, max depth 10
Treatment learner (m_0)	Random Forest	500 trees, max depth 10
Min samples per leaf	50	Balance flexibility/overfitting
Max features	5	$\sqrt{(23 \text{ covariates})} \approx 5$

Covariates

Nuisance models	All 23 time-varying covariates	Same as main regression
-----------------	--------------------------------	-------------------------

Treatment variable	Log bilateral migration stock		Residualized cross-fitting	via
Outcome variable	5-year rolling correlation	GDP	Residualized cross-fitting	via

Inference

Standard errors	Multiplier bootstrap	500 iterations
Confidence level	95%	Standard
Point estimate	0.036	Average treatment effect
95% CI	[0.021, 0.051]	Statistically significant

Notes: Cross-fitting procedure: (1) Split data into K=5 folds; (2) For each fold, train ML models on other 4 folds and predict on held-out fold; (3) Pool out-of-sample residuals; (4) Regress outcome residuals on treatment residuals. Results are robust to alternative learners (Gradient Boosting: 0.035, Elastic Net: 0.033) and hyperparameter choices.