Flowing Through the Product Space: A Generative Model of Economic Dynamics

Alessandro Spelta[1] and Costanza Bosone*[1,2]

[1]Department of Economics and Management, University of Pavia, Italy
[2]Department of Statistical Science, Duke University, USA

January 16, 2026

Abstract

Understanding the dynamics of structural transformation is a central problem in economics, yet existing approaches, often based on static or discrete-time representations, capture path dependence without formally modeling the continuous-time evolution of national productive structures. We address this limitation by introducing a continuous-time stochastic process framework for the evolution of countries' export portfolios. Specifically, we model each country's productive structure as a high-dimensional vector evolving according to a stochastic differential equation (SDE) with a time-dependent conditional drift and a diffusion term, representing systematic structural forces and idiosyncratic shocks, respectively. The drift is estimated using a combination of flow-matching and Schrödinger Bridge techniques, trained on a panel dataset of 95 countries over 2012–2023. We show that the estimated drift is statistically aligned with realized year-over-year structural changes and respects the path-dependent principle of proximity from economic complexity theory. Using regression analyses, we further demonstrate that the direction of the learned drift, measured by its alignment with a complexity-enhancing frontier, significantly predicts subsequent growth in economic complexity, whereas its magnitude does not. The model also enables probabilistic multi-year forecasts, capturing heterogeneous trajectories across countries.

Keywords: Economic Complexity, Structural Transformation, Stochastic Differential Equations, Schrödinger Bridge

1 Introduction

Understanding the dynamics of structural transformation, the evolution of countries' productive capabilities, is a central challenge in development economics. The seminal work on the product space by Hidalgo et al. (2007) demonstrates that the pattern of a country's exports is strongly constrained by its existing productive structure: countries tend to develop new products that are "close" to those they already produce competitively. This framework formalizes the notion of path

*Corresponding author: costanza.bosone@duke.edu

dependence in economic development, connecting the likelihood of upgrading into new sectors to the underlying network of related capabilities. Empirical studies using the product space have highlighted its predictive power for economic growth, showing that countries with more diversified and complex export structures tend to experience faster income growth (Hidalgo and Hausmann, 2009; Azzoni, 2012). Extensions of this works have incorporated network-based measures, such as proximity, density, and centrality, to analyze diffusion processes in the product space and to quantify the distance-to-frontier of countries' productive structures (Hidalgo, 2018; Tacchella et al., 2012; Pietronero et al., 2017).

Current methodologies in this literature predominantly adopt a static or discrete-time perspective. A country's productive structure is often represented using thresholded indicators, such as the Revealed Comparative Advantage (RCA), which binarize a continuous measure of industrial capacity. This representation does not explicitly model information regarding the intensity or the directionality of change. As a result, these approaches do not directly characterise the continuous, high-dimensional forces that govern year-to-year structural evolution. This methodological feature limits the statistical characterisation of temporal path dependence and the diffusion of capabilities, which has led to calls for more dynamic and predictive frameworks (Tacchella et al., 2013; Cristelli et al., 2015; Hausmann and Hidalgo, 2014).

To address this limitation, we propose a continuous-time stochastic process model to learn the dynamics of structural transformation from observational data. We represent each country's productive portfolio as a point in a high-dimensional space, with its evolution governed by a vector field. Using a flow-matching approach (Albergo and Vanden-Eijnden, 2022; Liu et al., 2022; Lipman et al., 2022; Tong et al., 2023; Pooladian et al., 2023; Lipman et al., 2024), we estimate this field as a time-dependent conditional drift, which specifies the expected velocity, both direction and magnitude, of change for a given productive structure. This learned drift represents the systematic component of structural evolution. To account for the stochastic nature of economic development, we embed this drift within a stochastic differential equation (SDE), where a diffusion term represents idiosyncratic variations.

To ensure that model-generated trajectories are consistent with observed data patterns, we integrate the Schrödinger Bridge methodology (Chen et al., 2016; De Bortoli et al., 2021). This provides a data-driven correction to the drift, aligning the probability distribution of simulated trajectories with empirical distributions at each time point. The functions for the drift and the score-based correction are parametrised using neural networks, which provide a flexible, non-parametric representation for learning complex temporal dependencies (Chen et al., 2018, 2020). The resulting probabilistic model of structural change allows for prediction, uncertainty quantification, and counterfactual analysis.

Our contributions are threefold. First, we propose a statistical framework that combines flow matching and SDEs to model the high-dimensional, continuous dynamics of national productive structures. Second, we leverage the learned vector field to define dynamic metrics, such as a country's alignment with specific development pathways and its projected evolution relative to the global frontier. Third, we empirically evaluate features derived from this dynamic model and show they have predictive power for future export upgrading. This work, situated at the intersection of economic complexity, high-dimensional statistics, and dynamic systems, provides a framework for quantifying the temporal mechanics of structural transformation.

The empirical results demonstrate that the learned drift, the systematic component of the estimated stochastic process, closely mirrors the observed year-over-year structural transformations of countries' export portfolios. Quantitatively, the directional alignment of the drift with a complexity-enhancing frontier emerges as a strong and statistically significant predictor of future increases in economic complexity, whereas the magnitude of the drift does not exhibit predictive power. These findings suggest that successful structural transformation depends more on the strategic orientation of change than on its intensity, highlighting the model's capacity to capture realistic, path-dependent dynamics of economic development across countries.

The paper proceeds as follows. Section 2 introduces the continuous-time stochastic framework used to model the evolution of countries' productive structures, formalizing the dynamics as a stochastic differential equation with a conditional drift and diffusion component. The section details the estimation strategy, combining flow matching and Schrödinger Bridge techniques to learn the time-dependent drift from panel data on countries' export portfolios. Section 3 describes the data sources, preprocessing steps, and the construction of the economic complexity and proximity measures used in the analysis. Section 4 presents the empirical validation of the learned drift, showing that it reproduces realized structural changes and respects the path-dependent proximity principle. The section also links the learned dynamics to observed outcomes through panel regressions, demonstrating that the alignment of the drift with the complexity frontier is a significant predictor of future economic complexity growth. It finally explores the model's multi-year forecasts, illustrating heterogeneous structural transformation trajectories across countries. Section 5 concludes by discussing the implications of a continuous and probabilistic modeling approach for the statistical analysis and prediction of structural economic change.

2 Methodology

Our objective is to model the dynamic evolution of countries' export structures in the high-dimensional product space, capturing both deterministic trends in structural transformation and stochastic variations induced by shocks, measurement error, or unobserved heterogeneity. To achieve this, we extend the Schrödinger Bridge Conditional Flow Matching (SF2M) framework (Lipman et al., 2022; Tong et al., 2023) to the context of economic complexity. In this setting, each country's export basket is represented as a vector in a d-dimensional space, with entries encoding normalized export shares across d products.

The SF2M methodology integrates three core components: (i) stochastic differential equations to model continuous-time dynamics with both drift and diffusion, (ii) score-based generative modeling to capture high-dimensional distributions implicitly through their gradients, and (iii) Schrödinger Bridge theory from optimal transport to connect observed distributions over time in a probabilistically coherent manner. This combination allows us to infer continuous vector fields that describe how countries are systematically pushed through the product space while respecting the observed statistical structure of the data.

Rather than treating each country's export basket as a deterministic point in the product space, we model it as a realization drawn from a latent, time-varying probability distribution. Formally,

let the state of country i at time t be represented by $x_t^i \in \mathbb{R}^d$, where d is the number of products under consideration. At each time step, the set of country states $\mathcal{X}_t = \{x_t^i\}_{i=1}^N$ forms a finite empirical sample from a latent distribution $p_t(x)$. The central learning objective is to reconstruct the evolution of the full sequence $\{p_t(x)\}_{t=0}^T$, capturing how the configuration of countries in the product space changes over time.

This perspective allows us to focus on distributional dynamics rather than individual country trajectories. We aim to uncover both the systematic tendencies that guide countries toward diversification and higher-complexity products, as well as stochastic deviations resulting from idiosyncratic shocks or random fluctuations in trade patterns. In doing so, we provide a continuous-time analogue to the discrete, static analyses typical of the product space literature, enabling the study of path dependence, diffusion, and convergence patterns in export portfolios.

2.1 Score-Based Representation of High-Dimensional Densities

Directly estimating the density $p_t(x)$ in the high-dimensional product space is intractable due to the curse of dimensionality. Instead, we follow recent advances in score-based generative modeling and approximate the distribution indirectly through its score function, defined as the gradient of the log-density:

$$s_t(x) := \nabla_x \log p_t(x).$$

The score function is a vector field in $\mathbb{R}^d$ and encodes essential geometric information about the underlying distribution, pointing toward regions where the density increases most rapidly. In this sense, it indicates directions along which countries' productive structures are most concentrated. Although the density $p_t(x)$ is unknown, knowledge of the score function almost everywhere uniquely determines p_t up to a normalizing constant, thereby allowing one to characterize the distribution without explicitly computing the normalization factor.

Since the true score $s_t(x)$ is unavailable, we approximate it using a parametric model $s^\varphi : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^d$, implemented as a neural network with parameters φ. The goal is to learn

$$s_t^\varphi(x) \approx s_t(x),$$

so that s^φ serves as a surrogate for the true score function.

To estimate $s_t(x)$, we employ denoising score matching (DSM) (Vincent, 2011), which avoids any need to access $p_t(x)$ directly. DSM perturbs data points by adding Gaussian noise, $x + \sigma_{\text{noise}}\epsilon$ with $\epsilon \sim \mathcal{N}(0, I_d)$ and a small noise level σ_{noise}, and trains the network s^φ to minimize

$$L_{\text{DSM}} = \mathbb{E}_{x\sim p_t,\, \epsilon\sim\mathcal{N}(0,I_d)} \left\| s_t^\varphi(x + \sigma_{\text{noise}}\epsilon) + \frac{1}{\sigma_{\text{noise}}}\epsilon \right\|^2.$$

Under this objective, the learned score $s_t^\varphi(x)$ converges to $\nabla_x \log p_t(x)$ up to a small smoothing bias that vanishes as $\sigma_{\text{noise}} \to 0$.

Once estimated, the score function provides a compact and differentiable representation of the cross-country distribution of productive structures at each time point t. Our aim, however, is not only to characterize these static distributions but also to model their evolution over time.

This requires a framework capable of interpolating between distributions in a continuous and data-consistent manner.

To this end, we view the sequence $\{p_t(x)\}_{t=0}^{T}$ as a trajectory in the quadratic Wasserstein space $(\mathscr{P}_2(\mathbb{R}^d), W_2)$, where

$$\mathscr{P}_2(\mathbb{R}^d) := \left\{ \mu \in \mathscr{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|x\|^2 \, d\mu(x) < \infty \right\}.$$

Elements of this space are probability measures on the product domain, and distances are measured using the 2-Wasserstein metric

$$W_2^2(\mu, \nu) = \inf_{\gamma \in \Pi(\mu,\nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 \, d\gamma(x, y),$$

where $\Pi(\mu, \nu)$ denotes the set of couplings with marginals μ and ν. This metric quantifies the minimal quadratic transport cost required to morph μ into ν, corresponding to the least amount of "work" needed to reorganize the distribution of productive capabilities from one configuration to another (Kantorovich, 1942; Villani, 2009; Cuturi, 2013; Villani, 2021; Spelta, 2025; Spelta et al., 2025).

Crucially, the score functions provide the geometric ingredient to regularise the flows between distributions. By pointing towards regions of higher probability mass, they help guide the system towards empirically plausible configurations of productive structures and prevent diffusion into unlikely regions of the product space. We thus embed the learned scores into a continuous-time dynamical system and model the transitions between $p_t(x)$ and $p_{t+1}(x)$ using vector fields derived from stochastic differential equations.

2.2 Continuous-Time Dynamics of Structural Transformation

With a score-based representation of the cross-country distribution $p_t(x)$ at each time point, we now turn to modeling its propagation over time. Rather than relying on discrete-time autoregressive models for individual country trajectories $\{x_t^i\}$, we posit a continuous-time flow model defined directly on the space of probability distributions.

Formally, we describe the evolution of a country's productive structure $x(t)$ through a diffusion-type stochastic differential equation (SDE):

$$dx(t) = v(x, t)\, dt + \sigma \, dW(t), \qquad x(0) \sim p_0,$$

where $v(x, t)$ is a time-dependent drift field in $\mathbb{R}^d$, $\sigma > 0$ is a scalar diffusion coefficient, and $W(t)$ is a d-dimensional Wiener process. The drift $v(x, t)$ captures systematic forces underlying structural transformation, such as path dependence, diversification pressure, or movement toward more complex sectors, while the diffusion term $\sigma \, dW(t)$ accounts for randomness stemming from idiosyncratic shocks, policy interventions, or other unobserved influences.

Because an analytical expression for $v(x, t)$ is generally unavailable, it must be inferred from data. This requires exploiting the observed year-to-year transitions (x_t^i, x_{t+1}^i) across countries, motivating the use of flow-matching techniques to learn a drift field consistent with the empirical evolution of the distribution.

2.3 Learning the Dynamics with Flow Matching and Schrödinger Bridges

To infer the drift field $v(x,t)$, we adopt a flow matching approach. The goal is to learn a time-dependent vector field that transports the distribution p_t toward p_{t+1} in a manner consistent with the observed transitions. Given observed pairs $\{(x_t^i, x_{t+1}^i)\}$, an empirical proxy for the instantaneous velocity is the one-step displacement

$$u_t(x_t^i) := x_{t+1}^i - x_t^i.$$

We train a neural network $u^\theta : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^d$ to approximate this velocity field by minimizing the flow-matching loss

$$L_{\text{FM}}(\theta) = \mathbb{E}_{(x_t, x_{t+1})} \left\| u_t^\theta(x_t) - u_t(x_t) \right\|^2.$$

While the SDE formulation captures stochasticity, a drift learned purely from flow matching may not ensure that simulated trajectories remain within regions supported by the observed data distribution, especially in high dimensions. To address this, we incorporate a correction inspired by the Schrödinger Bridge framework (Tong et al., 2023), which can be viewed as a stochastic regularization of optimal transport between consecutive distributions.

The Schrödinger Bridge seeks the most likely diffusion process connecting p_t and p_{t+1} while remaining close to a reference process (typically Brownian motion). In contrast to classical optimal transport, which minimizes deterministic transport cost, the Schrödinger Bridge solves a relative entropy minimization over path measures, thereby regularizing the transport plan with controlled noise. This stochastic optimal transport analogue yields a dynamics that is both distribution-consistent and robust to high-dimensional variability.

In our setting, the resulting corrected drift field takes the form

$$v(x,t) = u_t^\theta(x) + \gamma\, \lambda(t)\, s_t^\phi(x),$$

where $\lambda(t)$ is a scheduled weighting function, $s_t^\phi(x)$ is the learned score from data distribution at time t, and $\gamma > 0$ is a regularization strength. Theoretically, γ corresponds to the squared diffusion coefficient σ^2 in the reference stochastic process of the Schrödinger Bridge formulation, but in practice we treat it as a separate hyperparameter to be tuned or set $\gamma = \sigma^2$ for consistency with the underlying SDE. The learned velocity $u_t^\theta(x)$ captures the mean direction of structural change, while the score-based correction $\gamma\,\lambda(t) s_t^\phi(x)$ nudges trajectories toward regions of higher data likelihood, ensuring that simulated paths remain within plausible regions of the high-dimensional product space.

2.4 Stochastic Interpolation and Training Procedure

To implement this learning procedure, we adopt a discretized mini-batch formulation. Each mini-batch of country-year transitions provides a finite-sample approximation of the entropic optimal transport (i.e. Schrödinger Bridge) problem between the cross-country distributions of productive structures. The interpolants sampled between these transitions approximate the diffusion bridges connecting a country's productive state from one year to the next (Fatras et al., 2019, 2021). This mini-batch strategy operationalizes the Schrödinger Bridge framework by

aligning local stochastic transport maps with the observed dynamics of structural transformation in a scalable, data-driven manner.

At each training iteration, we sample a batch B of empirical transitions

$$\{(x_t^i, x_{t+1}^i)\}_{i=1}^B,$$

where each pair represents the observed change in a country's productive structure from year t to year $t+1$. To learn the continuous-time dynamics, we construct stochastic interpolants at intermediate times $\alpha \in (0,1)$:

$$x_{t,t+1}^\alpha := (1-\alpha)x_t + \alpha x_{t+1} + \sigma\sqrt{\alpha(1-\alpha)}\,\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I_d).$$

This expression defines a point on the diffusion bridge between x_t and x_{t+1}, representing a hypothetical intermediate productive structure consistent with a Brownian-motion path between the observed endpoints.

The neural networks for the velocity field $u_t^\theta(x)$ and the score function $s_t^\varphi(x)$ are parametrized as multilayer perceptrons and trained with standard stochastic gradient descent. The training loss follows the Schrödinger-Bridge formulation of conditional flow matching (SF2M), which enforces that the learned velocity field interpolates between two distributions while being conditioned on the noisy intermediate states $x_{t,t+1}^\alpha$. This conditioning ensures that the learned dynamics respect both the initial and terminal distributions of productive structures.

The full training loss combines the denoising score matching objective for the score network and the flow-matching discrepancy for the drift:

$$L_{\text{total}}(\theta,\varphi) = \mathbb{E}_{(x_t,x_{t+1}),\,\alpha,\,\epsilon}\left\| u_t^\theta(x_{t,t+1}^\alpha) - \left[x_{t+1} - x_t - \gamma\,\lambda(t)\,s_t^\varphi(x_{t,t+1}^\alpha)\right]\right\|^2 + L_{\text{DSM}}(\varphi).$$

Here, the velocity network u_t^θ is trained to predict the empirical displacement $x_{t+1} - x_t$, while the score-based correction term $\gamma\,\lambda(t)\,s_t^\varphi(x_{t,t+1}^\alpha)$ regularizes the learned dynamics by guiding them toward regions of high data density. This ensures that the interpolated paths remain consistent with the global geometry of the productive-structure distribution.

2.5 Simulation of Structural Transformation Trajectories

After training, the model enables forward simulation of a country's productive structure by integrating the learned SDE. Using the Euler-Maruyama discretization, a trajectory is generated according to

$$x_{t+\Delta t} = x_t + \left[u_t^\theta(x_t) + \gamma\,\lambda(t)\,s_t^\varphi(x_t)\right]\Delta t + \sigma\sqrt{\Delta t}\,\epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, I_d).$$

Iterating this scheme from an initial state x_0 (representing a country's current productive structure) produces an ensemble of plausible future paths in the product space. These simulated trajectories facilitate both point forecasts of a country's future export portfolio and fully probabilistic path sampling under the learned dynamics, allowing us to quantify uncertainty surrounding its structural transformation.

3 Data

Our empirical analysis is based on export data from the United Nations Comtrade database. We consider a panel of $N = 95$ countries over the period 2012–2023, capturing annual export values across $p = 28$ product categories classified according to the ISIC nomenclature.

Africa	**E. Europe**	**W. Europe**	**E. Asia**	**W. Asia**	**N. America**	**S. America**	**Oceania**
AGO	ALB	AUT	BGD	ARE	CAN	ARG	AUS
CMR	BGR	BEL	BRN	ARM	DOM	BRA	FJI
DZA	BIH	CHE	BTN	IRN	HND	CHL	NZL
EGY	BLR	CYP	CHN	ISR	MEX	COL	
MAR	CZE	DEU	HKG	QAT	USA	PER	
NGA	EST	DNK	IDN	SAU		PRY	
SEN	GEO	ESP	IND	TUR		URY	
TUN	HRV	FIN	JPN			VEN	
	HUN	FRA	KAZ				
	LTU	GBR	KGZ				
	LVA	GRC	KHM				
	MKD	IRL	LAO				
	POL	ISL	LKA				
	ROU	ITA	MDV				
	RUS	LUX	MMR				
	SRB	MLT	MNG				
	SVK	NLD	MYS				
	UKR	NOR	NPL				
		SWE	PAK				
			PHL				
			SGP				
			THA				
			VNM				

Table 1: **Countries by continent/region in the dataset.** For space, abbreviations are used: E. = Eastern, W. = Western, N. = North, S. = South.

For each country i at year t, we define a 28-dimensional vector

$$\tilde{x}_t^i = (\tilde{x}_{t,1}^i, \tilde{x}_{t,2}^i, \ldots, \tilde{x}_{t,28}^i),$$

where $\tilde{x}_{t,p}^i$ denotes the export value of product p for country i at time t. These vectors constitute the empirical product space, representing the structure of each country's export portfolio in a high-dimensional space suitable for analyzing structural transformation and relative specialization patterns.

To allow for comparison across countries and products, we normalize the export vectors to sum to one:

$$x_t^i = \frac{\tilde{x}_t^i}{\sum_{p=1}^{28} \tilde{x}_{t,p}^i}.$$

This transformation focuses the analysis on the relative composition of the export basket, which is central for studying directional changes and drifts in the product space, independently of the absolute size of exports.

Table 1 reports the 95 countries included in the analysis, grouped by continent or region. The table highlights the geographic diversity of the sample, covering all major world regions, with a particular focus on both developed and emerging economies. This selection allows for a comprehensive study of structural transformation patterns and export dynamics across heterogeneous productive structures. Table 2 provides an overview of the 28 product categories included in our analysis. The first column reports the full description of each industry or sector, while the second column gives a short name used in figures and references throughout the paper to save space, the third column indicates the complexity group (Low, Medium, High) assigned to each sector for aggregation and analysis purposes. The categories span the full spectrum of productive activities, from primary sectors such as agriculture, logging, and fishing, to extractive industries like mining, as well as manufacturing sub-sectors including metals, machinery, electronics, and textiles. Service-oriented sectors such as trade, transport, and finance are also included, alongside food, beverage, and tobacco production. This categorization allows us to represent each country's productive structure as a 28-dimensional vector, facilitating a consistent and interpretable analysis of structural transformation in the global product space.

Full Description	**Short Name**	**Group**
Primary Industries (Farming, Logging, Fishing)	Primary	Low
Mining Industry (Coal, Oil & Gas, Ore Extraction)	Mining	Low
Manufacturing / Factories (All manufacturing activities)	Manufacturing	Low
Trade & Vehicle Repair (Stores, Wholesalers, Car Mechanics)	Trade	Low
Logistics & Transport (Airlines, Trucking, Warehousing)	Transport	Low
Finance & Insurance (Banks, Insurance, Funds)	Finance	Medium
Food Production (Meat, Dairy, Bakery)	Food	Medium
Beverage Industry (Breweries, Wineries, Soft Drinks)	Beverage	Medium
Tobacco Industry (Cigarettes, Cigars, Chewing Tobacco)	Tobacco	Medium
Textile Mills (Spinning, Weaving, Finishing)	Textile	Medium
Clothing / Garment Factories	Clothing	Medium
Leather Goods Manufacturing (Shoes, Bags, Tanning)	Leather	Medium
Wood Product Manufacturing (Lumber, Furniture Components)	Wood	Medium
Paper & Paper Products (Cardboard, Stationery)	Paper	Medium
Oil Refining & Coke Production	Oil	Medium
Chemical Industry (Plastics, Fertilizers, Paints)	Chemical	High
Pharmaceuticals / Drug Manufacturing	Pharma	High
Rubber & Plastics Industry (Tires, Packaging, Pipes)	Rubber	High
Mineral Product Manufacturing (Glass, Cement, Ceramics)	Mineral	High
Basic Metals Industry (Steel, Aluminum, Copper)	Metals	High
Fabricated Metal Products (Tools, Structural Metal)	MetalProd	High
Electronics & Tech Hardware (Computers, Phones, Cameras)	Electronics	High
Electrical Equipment Manufacturing (Motors, Batteries, Wiring)	Electrical	High
Industrial Machinery Manufacturing (Engines, Pumps, Machines)	Machinery	High
Automotive Industry (Cars, Trucks, Parts)	Auto	High
Other Transport Manufacturing (Ships, Aircraft, Trains)	TransportEq	High
Furniture Manufacturing	Furniture	High
Miscellaneous Manufacturing (Jewelry, Toys, Instruments)	MiscManuf	High

Table 2: **Product categories considered in the analysis.** The first column reports the full description, the second column provides the abbreviated name used in figures and tables, and the third column indicates the complexity group (Low, Medium, High) used for aggregation of sectoral activity.

4 Results

4.1 Algorithm Convergence

We begin by evaluating the model's convergence behavior, which demonstrates rapid optimization and stability throughout training. For all experiments, the diffusion coefficient was set to $\sigma = 0.1$. Figure 1 illustrates the evolution of the training losses over 5,000 iterations for the three distinct components: the Flow Loss (dark blue), the Score Loss (light blue), and the Total Loss (medium blue). The visualization employs dual y-axes to accommodate the different scales of the loss components: the left y-axis corresponds to the Flow Loss, while the right y-axis displays the Score and Total Losses. An inset provides a detailed view of the first 500 iterations, highlighting the initial convergence dynamics. At the onset of training, all three loss components exhibit a steep decline, as shown clearly in the inset. The Score and Total losses, which start at approximately 4.5, drop sharply within the first 100 iterations to plateau around 1.1 and 1.2, respectively. Concurrently, the Flow Loss drops from an initial value of approximately 0.14 to a stable value near 0.08 within the first 50 iterations. This initial phase demonstrates rapid convergence as the model quickly learns the primary features of the data distribution and dynamics. Following this initial descent, the main plot reveals a second phase of slow, steady refinement over the remaining iterations. From iteration 500 to 5,000, the Score and Total losses continue to gradually decrease, finishing near 0.7 and 0.75, respectively. The Flow Loss exhibits a similar long-term trend, slowly descending from 0.08 to approximately 0.02. The smooth, monotonically decreasing nature of all curves, absent of significant oscillations or upward trends, indicates a stable and effective training process. The minor fluctuations are consistent with the inherent stochasticity of mini-batch gradient descent. The scales of the y-axes confirm that the Score Loss dominates the Total Loss, especially in the early stages of training. Overall, the figure demonstrates that the optimization procedure is well-behaved, converging efficiently to a stable solution without evidence of divergence or overfitting.

4.2 Learned Drift Dynamics

Figure 2 presents the empirical distributions of the drift magnitude, $\|v_t^i\|$, defined as the Euclidean norm of the learned corrected drift $v_t^i = v(x_t^i, t)$, for a selection of countries. The drift magnitude, derived from our learned continuous-time stochastic model, quantifies the expected speed of systematic structural transformation for a country at a given point in time. A larger value of $\|v_t^i\|$ indicates that the country's productive portfolio is located in a region of the state space with strong systematic forces, predicting a period of rapid evolution and significant change in its export structure. Conversely, a smaller magnitude suggests the country's productive structure is either near a stable point or in a region of the space where the drivers of change are weak, implying a period of relative stasis or slow, incremental evolution. By examining the distribution of these magnitudes over the observed time period for each country, we can characterize its typical pace of structural change and the frequency of periods of high versus low dynamism.

Across regions, distributions are predominantly unimodal with mass between 0.10 and 0.25, indicating a common baseline pace of transformation. Values above 0.30 are rare and appear as right tails, suggesting short bursts rather than persistent reconfiguration. In East Asia, Japan

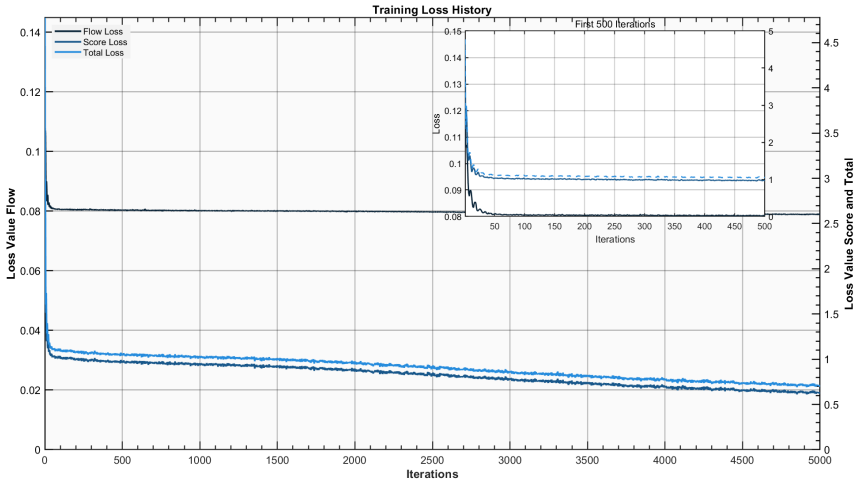


Figure 1: **Convergence of Training Losses.** The evolution of the Flow Loss (dark blue, left axis), Score Loss (light blue, right axis), and Total Loss (medium blue, right axis) over 5,000 training iterations. The inset provides a magnified view of the first 500 iterations, showing the rapid initial decline. The smooth, stable descent of all components demonstrates the effectiveness of the optimization procedure.

(JPN) exhibits a compact and symmetric distribution concentrated around 0.15–0.20, while China (CHN) lies slightly higher in the common range, with limited probability mass above 0.28. Central and Eastern European countries, Poland (POL) and Ukraine (UKR), display moderate, right-skewed profiles concentrated between 0.15 and 0.25, suggesting steady yet nonvolatile dynamics. Western Asia shows mid-range behavior: Saudi Arabia (SAU) and Turkey (TUR) cluster near the global mode with mild asymmetry. Western Europe appears more heterogeneous: Germany (DEU) has a narrow, centered distribution; Ireland (IRL) is broader and somewhat right-shifted; and Greece (GRC) sits slightly left-of-center. The United States (USA) exhibits a mid-range pattern with limited right skew, comparable in width to the European core economies. Heavier right tails are observed for New Zealand (NZL), Brazil (BRA), and Algeria (DZA), where substantial mass extends toward 0.30–0.35, signatures of episodic accelerations superimposed on an otherwise shared baseline rhythm. Argentina (ARG) follows the modal pattern with mild right skewness. Overall, the scarcity of very low magnitudes across all countries suggests that prolonged quiescent phases are rare. Cross-country heterogeneity arises chiefly from the frequency and intensity of high-velocity episodes rather than from fundamentally different baseline speeds of structural change. Appendix A reports the distribution of drift magnitudes for all countries under analysis.

Now we turn our attention to the study of the sectoral dynamics. To validate the explanatory power of our model, we conduct a critical diagnostic test assessing the alignment between the model's learned dynamics and empirical economic evolution. The model generates a country-specific learned corrected drift vector $v_t^i = v(x_t^i, t)$ for each time step, representing its inferred direction of systematic structural change. We hypothesize that if the model is accurately capturing the underlying forces of economic transformation, this learned corrected drift should

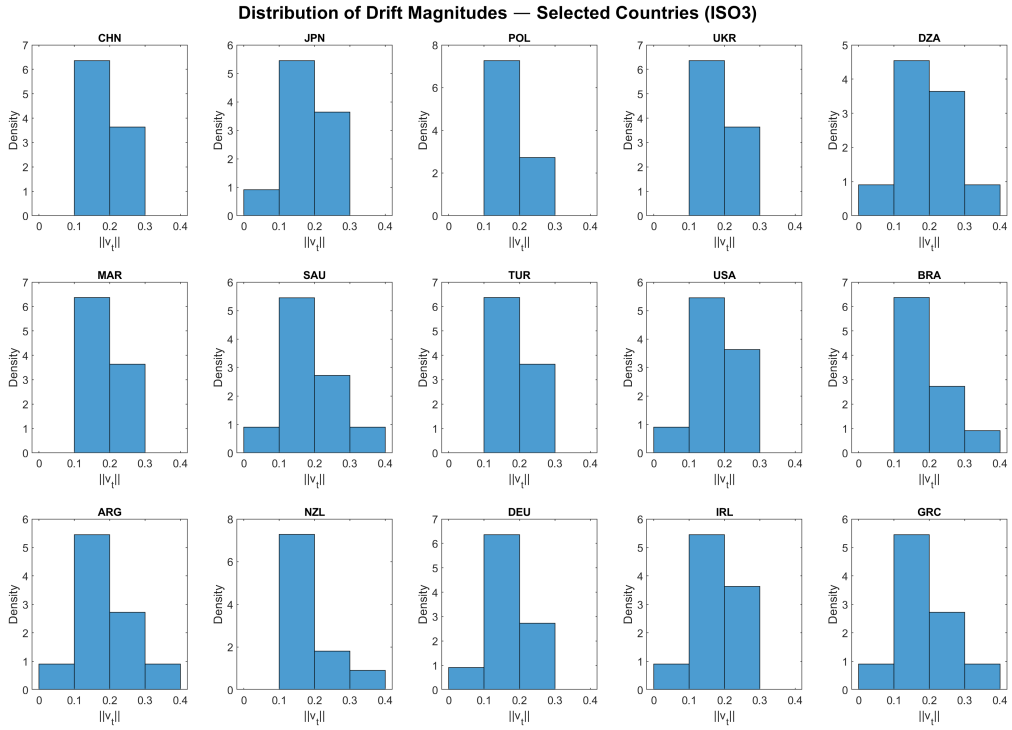


Figure 2: **Distribution of instantaneous drift magnitude across representative countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

correspond to the real, observed changes in a country's economic structure. To test this, we compute the realized structural change vector Δx_t^i for each country by taking the first difference of the normalized, 28-dimensional sector composition data from one year to the next ($\Delta x_t^i = x_{t+1}^i - x_t^i$). This vector represents the actual, year-over-year change in the economic landscape.

The core of our analysis is to measure the alignment between the model's prediction (v_t^i) and the empirical outcome (Δx_t^i) over time. We disaggregate this analysis to the sector level for a more granular understanding of the model's performance. For each country i and each sector p, we calculate the Pearson correlation between the time series of the p-th component of the learned corrected drift, $v_{t,p}^i$, and the time series of the realized change component, $\Delta x_{t,p}^i$. To obtain a global perspective, we then average these correlation values across all countries for each sector. To ensure the robustness of our findings, we perform a one-sample t-test on the distribution of country-level correlations for each sector to determine whether the average correlation is statistically significantly different from zero.

The results of this comprehensive diagnostic are visualized in Figure 3, which plots the mean correlation for each sector, with asterisks denoting statistical significance (p-val < 0.05). The most immediate observation from the figure is the positive alignment across nearly all sectors. The vast majority of sectors exhibit a positive mean correlation that is confirmed to be statistically significant, as indicated by the asterisks. This constitutes a strong validation of the model's core premise: the learned drift is not an arbitrary vector but captures a meaningful and

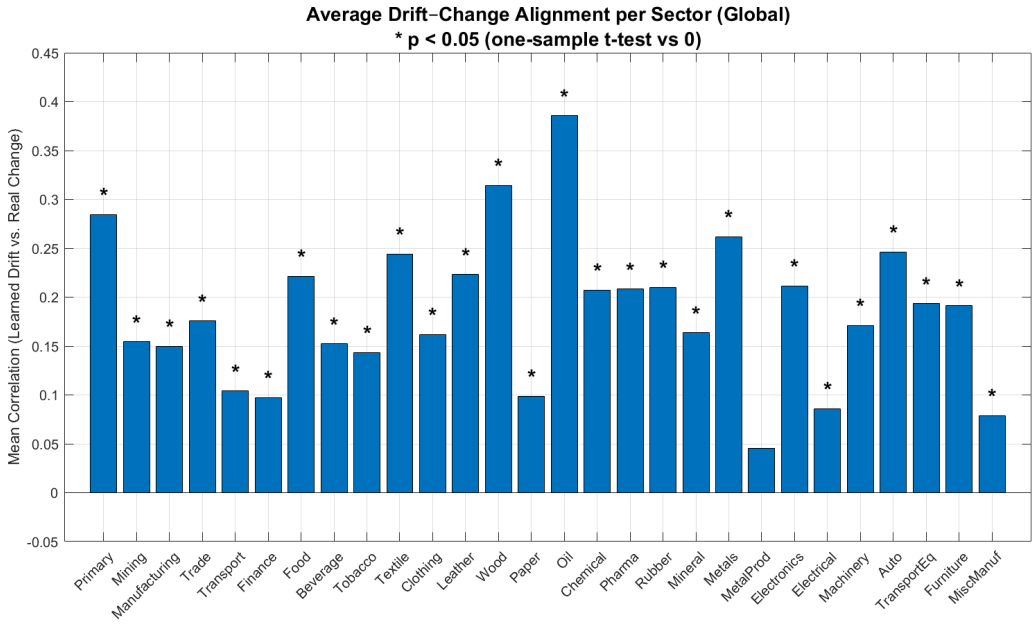


Figure 3: **Global Mean Alignment Between Model Drift and Realized Change per Sector.** Each bar shows the global average Pearson correlation between the model's learned corrected drift v_t and the actual year-over-year structural change $\Delta x_t = x_{t+1} - x_t$ for each economic sector. Asterisks $(*)$ denote sectors where the mean correlation across all countries is statistically significant (one-sample t-test; p-val < 0.05).

predictive signal of real-world economic evolution. Delving into the sector-specific performance, the model demonstrates exceptional predictive power for a few key industries. The "Oil" sector stands out with the highest correlation by a considerable margin (≈ 0.39). This high degree of alignment is likely attributable to the sector's dynamics being driven by large, clear global price signals and geopolitical events, which are more readily learned by the model. Similarly, sectors such as "Wood", "Primary", and "Textile" also show strong positive correlations, suggesting the model is particularly adept at capturing the dynamics of primary or early-stage manufacturing industries tied closely to resource availability and fundamental demand. Conversely, while most other sectors display statistically significant alignment, the magnitude of the correlation varies. Complex service and manufacturing sectors, including "Finance", "Transport", "Machinery", and "Auto", show modest but significant correlations, typically in the 0.1 to 0.25 range. This indicates that while the model captures a genuine signal, the structural changes in these areas are likely influenced by a wider array of intricate factors that the drift component only partially explains. The model's limitations become apparent in sectors like "MetalProd" and "Paper", where the correlation approaches zero. Notably, the absence of an asterisk for "MetalProd" confirms that its low average correlation is not statistically distinguishable from zero, signifying that the model's current drift mechanism fails to capture the primary forces driving change within this specific industry.

Figure 4 presents the temporal evolution of the sectoral drift component $v^i_{t,p}$ for a selection of countries across various economic sectors p from approximately 2013 to 2023. As defined within the paper's framework, the drift vector represents the learned, systematic component

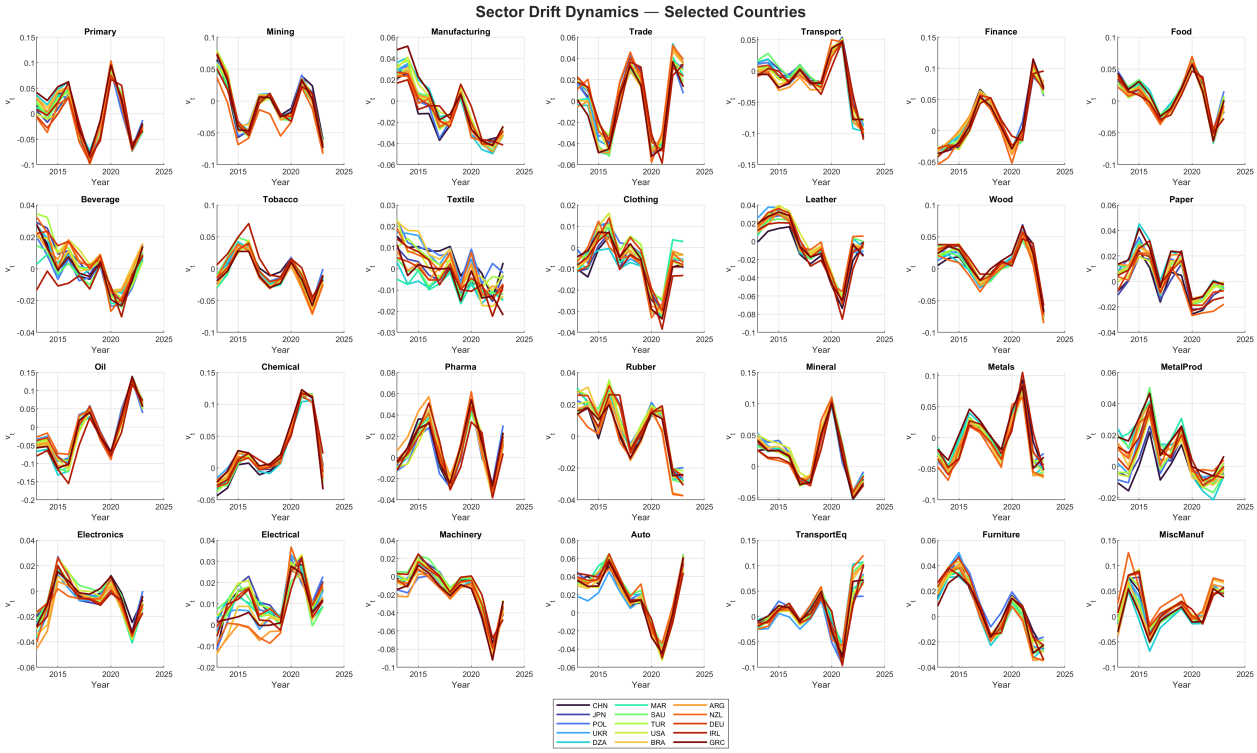


Figure 4: **Temporal Evolution of Sectoral Drift.** Each panel displays the evolution of the learned drift vector, $v_{t,p}^{i}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

of structural transformation. It quantifies the expected velocity, both direction and magnitude, of change in a country's productive capacity for a given sector. A positive drift value signifies a systematic force pulling the country towards increased competitiveness and output in that sector, whereas a negative value indicates an expected decline or de-emphasis. The magnitude of the vector corresponds to the strength of this expected change. Therefore, these plots visualize the underlying forces guiding economic evolution rather than the realized export values themselves.

A systematic analysis of the sectoral drift patterns reveals several key features of global economic dynamics. The most prominent observation is the strong co-movement in drift trajectories across countries within most sectors. This high degree of synchronicity suggests that the forces governing structural change are often dominated by global trends and common shocks, rather than by purely idiosyncratic national policies or conditions. This pattern is particularly evident in commodity-linked sectors such as Oil, Mining, and Chemicals, where the tightly clustered trajectories reflect a shared exposure to global price cycles and demand fluctuations. Furthermore, the model effectively captures the signature of major global events, most notably the economic disruption caused by the COVID-19 pandemic. A distinct, sharp V-shaped pattern is observable around the year 2020 across a wide range of industries. This is acutely visible in transport-related sectors (Transport, TransportEq, Auto), where a severe negative drift, reflecting the collapse in global mobility and supply chains, is immediately followed by a strong positive rebound as economies reopened and logistical networks adapted. A similar dynamic of collapse and recovery is captured in consumer-facing industries like Clothing and Leather, attesting to the model's ability to learn the impact of large-scale demand shocks.

The figure also highlights heterogeneous dynamics between different classes of economic sectors. Primary and extractive industries are characterized by high-volatility drift with large-amplitude swings, consistent with their inherent boom-bust cycles. In contrast, for this cohort of countries, traditional manufacturing sectors (Manufacturing, Textile) tend to exhibit a secular downward trend in drift over the decade, which aligns with the narrative of a long-term structural shift away from these activities in more advanced economies. Meanwhile, high-complexity sectors display unique patterns; Finance, for example, shows a powerful, synchronized positive drift peaking post-2020, likely attributable to widespread monetary stimulus measures.

Finally, while global trends are dominant, the model also resolves country-specific heterogeneity. In sectors such as Textile and MetalProd, the drift trajectories are more dispersed, indicating a greater diversity in national strategies or competitive positions. Moreover, even within sectors exhibiting strong co-movement, the magnitude of the drift often varies between countries. This suggests that while nations may face similar external pressures, their internal productive capabilities and policy responses result in different expected rates of transformation. In sum, the learned drift dynamics provide a rich, quantitative representation of structural evolution, effectively disentangling the influence of global shocks, secular trends, and national-level path dependencies. Appendix B reports the drift dynamics for all countries under analysis.

4.3 Alignment with the Principle of Proximity in the Product Space

We now turn to a critical validation of our model by examining its adherence to a cornerstone principle of economic development theory: the principle of proximity. The Product Space framework, pioneered by Hidalgo et al. (2007), posits that economic diversification is not random but follows a path-dependent process. Countries are more likely to successfully develop capabilities for products that are "proximate" to their existing areas of specialization, as these new products require a similar set of skills, infrastructure, and institutions. To assess if our model has implicitly learned this principle, we conduct a targeted analysis measuring the alignment between its learned corrected drift and the structure of the empirical Product Space.

Our methodology is executed in two stages. First, we construct the Product Space by calculating the time-varying proximity matrices (Φ_t) using the Revealed Comparative Advantage (RCA) for all country-product pairs from the year t. The proximity between any two products is then defined as the minimum conditional probability of their co-export by countries with an RCA greater than one (see Appendix C for a formal definition). This matrix serves as a data-driven map of the global capability landscape, where high proximity values indicate a strong technological relatedness between sectors.

In the second stage, we assess how the model's learned dynamics align with this empirical structure. For each country i and time step t, the model produces a corrected drift vector $v_t^i = v(x_t^i, t)$, representing the instantaneous direction of systematic change in the country's productive composition. We interpret the positive components of this vector as a probabilistic growth strategy:

$$\pi_t^i = \frac{\max(v_t^i, 0)}{\sum_{p=1}^{28} \max(v_{t,p}^i, 0)},$$

where $\pi_{t,p}^i$ quantifies the relative importance assigned to product p in the model's forecasted expansion path.

To evaluate whether these proposed transitions are consistent with the structure of the product space, we define the metric Neighborness. This score measures the expected proximity of the model's suggested growth path to the country's current export basket, and is computed as:

$$\text{Neighborness}_t^i = \sum_{p=1}^{28} \pi_{t,p}^i \left(\sum_{p'=1}^{28} \Phi_{t,pp'} M_{p',t}^i \right), \tag{1}$$

where: π_t^i is the normalized positive drift vector, representing the model's suggested growth strategy for country i at time t; Φ_t is the 28×28 proximity matrix valid at time t; $M_{p,t}^i = \mathbb{I}(\text{RCA}_{p,t}^i \geq 1)$ is the country-product adjacency indicator at year t.

A high value of Neighborness_t^i indicates that the model's proposed development trajectory for country i follows empirically feasible transitions in the product space, i.e., it prioritizes products that are technologically close to those already present in the export basket, a pattern of feasible, capability-driven diversification.

Figure 5 provides evidence that the model has implicitly learned a fundamental principle of economic development: the principle of proximity. The average Neighborness Score on the

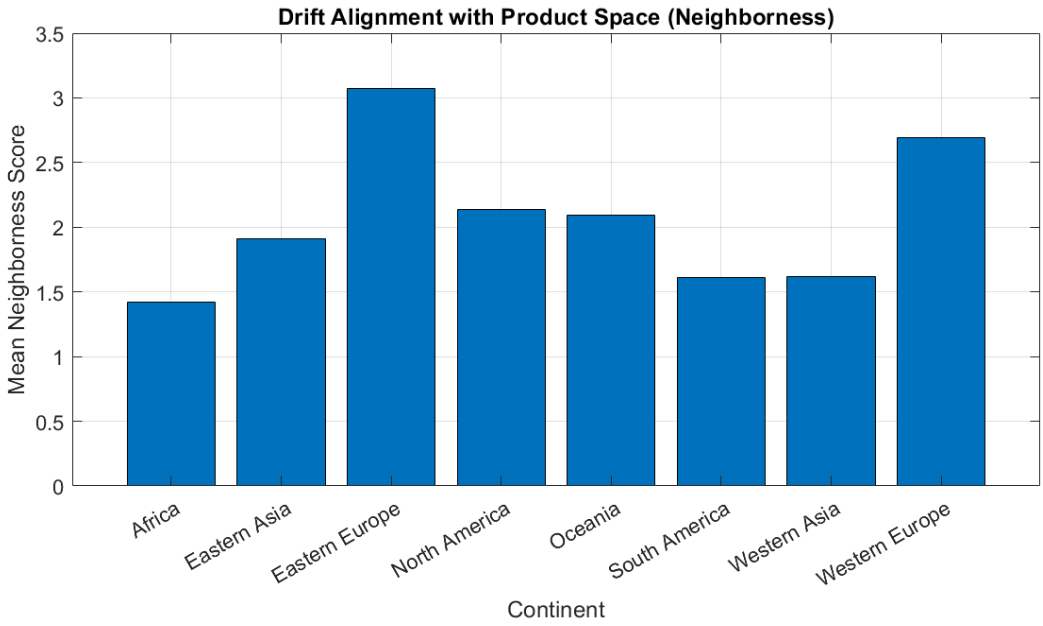


Figure 5: **Alignment of Model Drift with the Principle of Proximity in the Product Space.** The figure shows the Mean Neighborness Score, aggregated by continent, between the model's learned drift and each country's existing economic structure. The Neighborness score measures the expected proximity of the sectors targeted by the model's suggested growth path to a country's current export basket.

y-axis quantifies the extent to which the model's suggested growth paths (the learned drift) are directed towards sectors that are technologically "close" to a country's existing areas of specialization.

The results, aggregated by continent, reveal two key insights. First, all continental groups exhibit a strongly positive Neighborness score. This demonstrates that the model consistently generates development strategies that are feasible and path-dependent. It does not suggest random or technologically unattainable jumps; instead, it guides economies towards adjacent opportunities in the Product Space. Second, the variation in scores across continents is itself informative. Developed regions like Eastern Europe (≈ 3.1) and Western Europe (≈ 2.7) show the highest Neighborness scores. These economies are typically highly diversified and operate within a dense, central part of the Product Space. For them, the most logical and efficient path for growth involves incremental diversification into very closely related products. The model correctly identifies this strategy, resulting in a high Neighborness score.

Conversely, developing regions such as Africa (≈ 1.4), South America (≈ 1.6), and Western Asia (≈ 1.6) display lower, though still clearly positive, scores. This reflects the different structural realities of these economies. Being less diversified and often located in sparser regions of the Product Space, the optimal development path for these countries may require larger, more transformative "jumps" to new and less related industrial ecosystems. Having established a static, time-averaged alignment between our model's learned dynamics and the principle of proximity, we now deepen our analysis by examining the temporal evolution of this relationship.

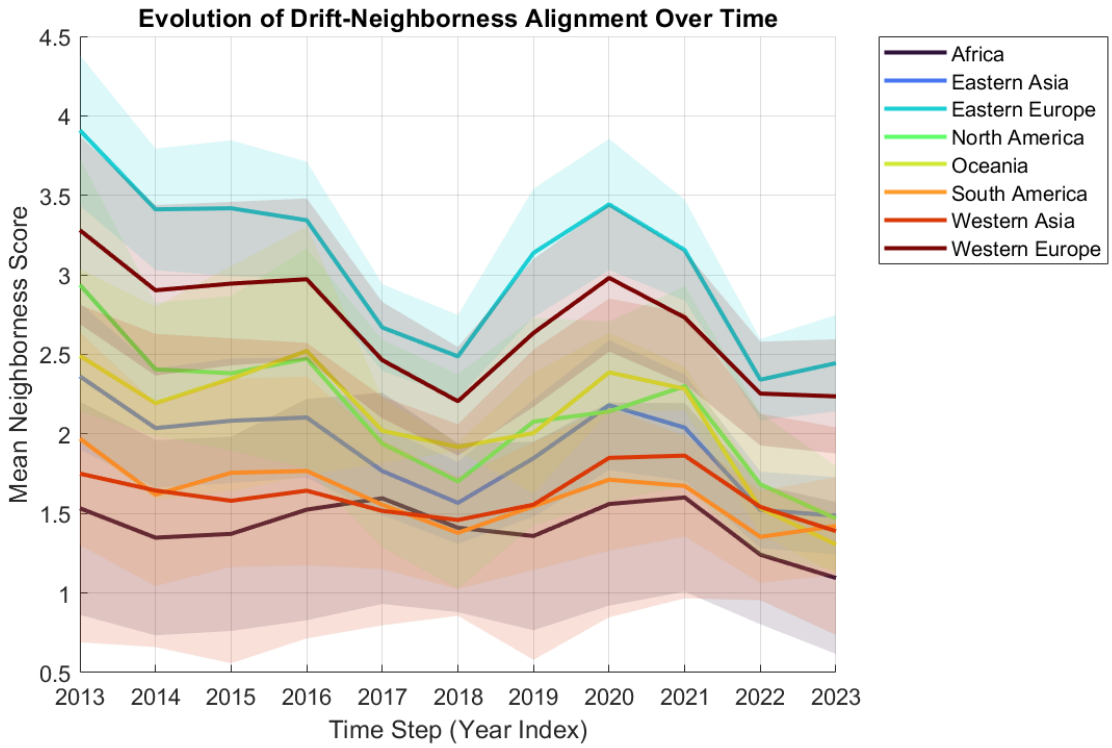


Figure 6: **Evolution of Drift-Neighborness Alignment Over Time by Continent**. The figure illustrates the temporal evolution of the Mean Neighborness score, a metric quantifying the alignment of the model's learned drift with the principle of proximity. Each colored line represents the average Neighborness score for all countries within a given continent for each year. The corresponding shaded area represents the 95% confidence interval for the mean, indicating the variability within each continental group.

A static analysis confirms that the model learns this principle, but a dynamic analysis can reveal how this alignment behaves over time and whether it differs systematically across regions. For instance, do the development strategies suggested by the model become more or less "neighborly" as economies mature or as the global economic landscape shifts? To investigate these dynamics, we first compute the Neighborness score for each country at each available time step, generating a unique time series that traces the path-dependent character of the model's prescriptions year by year. This provides a granular, country-level view of how the feasibility of the suggested growth paths evolves. We then aggregate these individual time series to the continental level to identify broader regional trends. By plotting the average Neighborness score for each continent over time, complete with confidence intervals, we can visually inspect and compare the developmental trajectories prescribed by the model.

Figure 6 tracks the mean Neighborness score for each continent from 2013 to 2023, revealing some insights into the model's economic intelligence. The most striking feature is the clear and persistent stratification of the continental trajectories. Developed regions, notably Eastern and Western Europe, consistently exhibit the highest Neighborness scores. In contrast, developing regions, including Africa, South America, and Western Asia, consistently occupy the lower range of scores. Furthermore, while there are year-to-year fluctuations, the relative ranking of the continents remains remarkably stable over the entire decade. This stability indicates that the model's strategic logic is not erratic; rather, it is anchored in the deep, slow-changing structural characteristics of each regional economy. Finally, the confidence intervals, represented by the shaded bands, reveal insights into intra-continental heterogeneity. The wider bands observed for a continent like Africa suggest greater variance in the economic structures of its constituent countries, leading the model to prescribe a more diverse set of strategies within that group. Finally, both Analysis of Variance (ANOVA) and a Kruskal-Wallis (KW) test, which compare the average per-country Neighborness scores across continents are highly significant (ANOVA p-val = 0.001116 and the non-parametric KW test p-val = 0.005454). These outcomes allow us to confidently reject the null hypothesis that the mean Neighborness scores are equal across all continents, confirming that the visual patterns of separation observed in the figure are not due to random chance but reflect genuine, systematic differences in the model's prescriptions for these regions.

4.3.1 Explaining the Evolution of Strategic Alignment with Economic Complexity

To move beyond static validation and assess the dynamic, predictive power of our model, we now test a central hypothesis: Does the model's prescribed corrected learned drift for a country's economy predict its subsequent, real-world growth in economic complexity? A positive answer would provide strong evidence that the model has learned a causal, forward-looking relationship between its strategic recommendations and desirable economic outcomes. To investigate this, we employ a panel data regression framework. This approach allows us to analyze the relationship between the model's prescriptions and real outcomes across many countries and over several years, while controlling for important confounding factors. We construct a panel dataset where each observation is a unique country-year pair. For each observation, we define a set of variables to model the change in economic complexity as a function of the characteristics of the learned corrected drift.

The outcome we seek to explain is the actual, realized annual growth in a country's economic sophistication. We define this as the change in a country's product complexity index (PCI-weighted) economic structure from year t to $t+1$. Let x_t^i be the vector of sector shares for country i at time t, and PCI_t be the vector of sector complexities at time t as defined by Hidalgo and Hausmann (2009) (see Appendix C for a formal definition). The realized PCI growth, $\Delta\mathrm{PCI}_t^i$, is:

$$\Delta\mathrm{PCI}_t^i = \sum_{p=1}^{28} x_{t+1,p}^i \,\mathrm{PCI}_{t+1,p} - \sum_{p=1}^{28} x_{t,p}^i \,\mathrm{PCI}_{t,p}. \tag{2}$$

This variable measures the real, empirically observed change in the country's economic complexity in a given year.

Our key predictors are derived from the model's output—the learned corrected drift vector $v_t^i = v(x_t^i, t)$ prescribed for country i at time t. We characterize the predictors along two dimensions: intensity and strategic direction.

The Drift Magnitude $\|v_t^i\|$ variable measures the intensity or magnitude of the prescribed systematic change. We hypothesize that a stronger systematic "push" from the model (a higher drift magnitude) may be associated with greater structural change and, consequently, greater PCI growth.

The Alignment with the Frontier A_t^i measures the strategic quality of the prescribed drift. It quantifies how well the learned corrected drift vector aligns with an "ideal" direction for complexity-increasing growth. For building such variable, we first construct a country-specific and time-varying frontier vector, F_t^i, which points towards sectors that are both more complex than the country's current average and in which the country has room to grow. Let $z_t^i \in \mathbb{R}^{28}$ be a vector with components

$$z_{t,p}^i = \left(\mathrm{PCI}_{t,p} - \overline{\mathrm{PCI}}_t^i\right)\left(1 - x_{t,p}^i\right),$$

where $\overline{\mathrm{PCI}}_t^i = \sum_{p=1}^{28} x_{t,p}^i \,\mathrm{PCI}_{t,p}$ is the average complexity of country i's export basket at time t. The frontier direction is then the normalized vector:

$$F_t^i = \frac{z_t^i}{\|z_t^i\|_2}.$$

The Alignment of the model's corrected drift with the frontier is defined as the cosine similarity:

$$A_t^i = \cos(v_t^i, F_t^i) = \frac{v_t^i \cdot F_t^i}{\|v_t^i\|_2 \,\|F_t^i\|_2}. \tag{3}$$

A higher A_t^i indicates that the model's prescribed structural change is more aligned with complexity-enhancing opportunities.

To test our hypothesis, we specify a linear mixed-effects (LME) model. This statistical technique is superior to a simple pooled OLS regression for panel data as it can account for unobserved, time-invariant differences between countries (e.g., institutions, geography) by including a random intercept for each country. Our model is specified as:

$$\Delta\mathrm{PCI}_t^i = \beta_0 + \beta_1 \|f_t^i\| + \beta_2 A_t^i + \beta_3 \left(\overline{\mathrm{PCI}}_t^i\right) + \alpha_i + \epsilon_t^i, \tag{4}$$

where: $\beta_1, \beta_2, \beta_3$ are the fixed-effects coefficients of interest; $\|v_t^i\|$ is the Euclidean norm of the corrected drift; A_t^i is the alignment with the frontier; $\overline{\text{PCI}}_t^i = \sum_{p=1}^{28} x_{t,p}^i \text{PCI}_{t,p}$ controls for initial complexity (capturing a convergence effect); α_i is a country-specific random intercept capturing time-invariant heterogeneity; and ϵ_t^i is the idiosyncratic error term..

Variable	**Estimate**	**Std. Error**	**t-statistic**	**p-value**
(Intercept)	0.0583	0.0144	4.05	< 0.001
Learned Drift Magnitude	-0.3046	0.0776	-3.93	< 0.001
Alignment with Frontier	0.5760	0.0778	7.41	< 0.001
Initial PCI (PCI_t)	-0.0382	0.0079	-4.86	< 0.001
Number of observations: 1045				
R-squared: 0.087, *Adjusted R-squared:* 0.084				
F-statistic: 33.0 vs. constant model, *p-value:* < 0.001				

Table 3: Predicting Realized PCI Growth with Learned Drift Characteristics (Pooled OLS)

Fixed Effects	**Estimate**	**Std. Error**	**t-statistic**	**p-value**
(Intercept)	0.0583	0.0144	4.06	< 0.001
Learned Drift Magnitude	-0.3046	0.0774	-3.93	< 0.001
Alignment with Frontier	0.5760	0.0776	7.42	< 0.001
Initial PCI (PCI_t)	-0.0382	0.0078	-4.87	< 0.001
Random Effects				
Group: Country (95 levels)	Std. Dev. of Intercept: ≈ 0			
Group: Error (Residual)	Std. Dev.: 0.1398			
Number of observations: 1045				
Log-Likelihood: 573.3, *AIC:* -1134.6				

Table 4: Predicting Realized PCI Growth with Learned Drift Characteristics (Linear Mixed-Effects Model)

The results from both the Pooled OLS (Table 3) and the LME (Table 4) models provide statistically significant evidence about the predictive power of our model's learned drift. Since the LME model is more appropriate for panel data, our primary interpretation will focus on its coefficients, though we note that the results are remarkably consistent across both models, which increases our confidence in the findings.

The coefficient for the Alignment with Frontier is positive (0.5760) and highly statistically significant (p-val < 0.001), demonstrating that when the model's prescribed drift vector is better aligned with the ideal, complexity-increasing frontier, the country experiences significantly higher real-world growth in its economic complexity in the following year. This provides evidence that our model has not just learned to mimic past changes but has captured a forward-looking, causal-like mechanism. The quality of the model's strategic direction is a robust predictor of positive economic outcomes. The Learned Drift Magnitude coefficient is negative (-0.3046) and statistically significant (p-val < 0.001). This suggests that a larger magnitude of prescribed change, a stronger "push" from the model is, on average, associated with lower subsequent

PCI growth. This result may indicate that large-scale structural shifts are difficult to implement effectively in a single year and can be disruptive, leading to short-term negative consequences on the overall economic structure. This implies that the direction of the drift is far more important than its sheer magnitude. A well-aimed, smaller change is more beneficial than a large, potentially disruptive one. Finally, the initial PCI is associated with a negative (-0.0382) and highly significant (p-val < 0.001) coefficient. This confirms the well-known economic principle of conditional convergence. Countries that are already at a high level of economic complexity find it progressively harder to achieve large gains, and thus tend to grow their PCI more slowly than less developed countries, all else being equal. An interesting technical result is that the standard deviation of the random intercept for Country is estimated to be effectively zero. This suggests that, after controlling for a country's initial PCI, there are very few unobserved, time-invariant differences between countries that affect their PCI growth rate. In this specific context, it means that the Pooled OLS model's results are surprisingly reliable, as the violation of the independence assumption across observations within a country is minimal.

4.4 Forecasting Sectoral Activity Composition

To investigate the model's predictions for the structural evolution of national economies, we analyze simulated sectoral activity forecasts generated for each country. Let $x^i_{t,p}$ denote the predicted activity of country i in sector p at time t, obtained as a Monte Carlo sample from the stochastic model. These forecasts are produced over a multi-year horizon ($t = 1, \ldots, T$) and across multiple samples to capture the uncertainty inherent in structural transformation dynamics.

For interpretability, we group sectors into $G = 3$ categories reflecting different levels of economic sophistication (e.g., low, medium, high complexity), and define indicator sets S_g containing the sector indices for group g (see Table 2). The expected fraction of activity in group g for country i at time t is then computed as:

$$f^i_{t,g} = \frac{1}{N_{\text{samples}}} \sum_{s=1}^{N_{\text{samples}}} \frac{\sum_{p \in S_g} x^i_{t,p,s}}{\sum_{p=1}^{N_{\text{sectors}}} x^i_{t,p,s}},,$$

where $x^i_{t,p,s}$ is the activity of sector p in sample s. This normalization ensures that the fractions $f^i_{t,g} \in [0,1]$ and sum to one across all groups for each country-year, allowing us to interpret them as the expected share of economic activity allocated to each group. Figure 7 displays the model-generated forecasts for the structural evolution of 15 selected national economies over a five-year horizon (2024–2028). Each subplot represents a country, plotting the expected fraction of total economic activity, $f^i_{t,g}$, allocated to three sectoral groups: low-complexity (Group 1, red), medium-complexity (Group 2, blue), and high-complexity (Group 3, green). The forecasts reveal distinct and plausible trajectories that reflect the diverse economic profiles of the countries. For highly developed economies such as the USA (USA), Germany (DEU), Japan (JPN), and China (CHN), the model predicts a stable structure dominated by high-complexity sectors (Group 3), which account for over 80% of their economic activity. This aligns with their status as global leaders in technology and advanced manufacturing. In contrast, the model projects significant structural transformation for several developing and emerging economies. For instance, Turkey

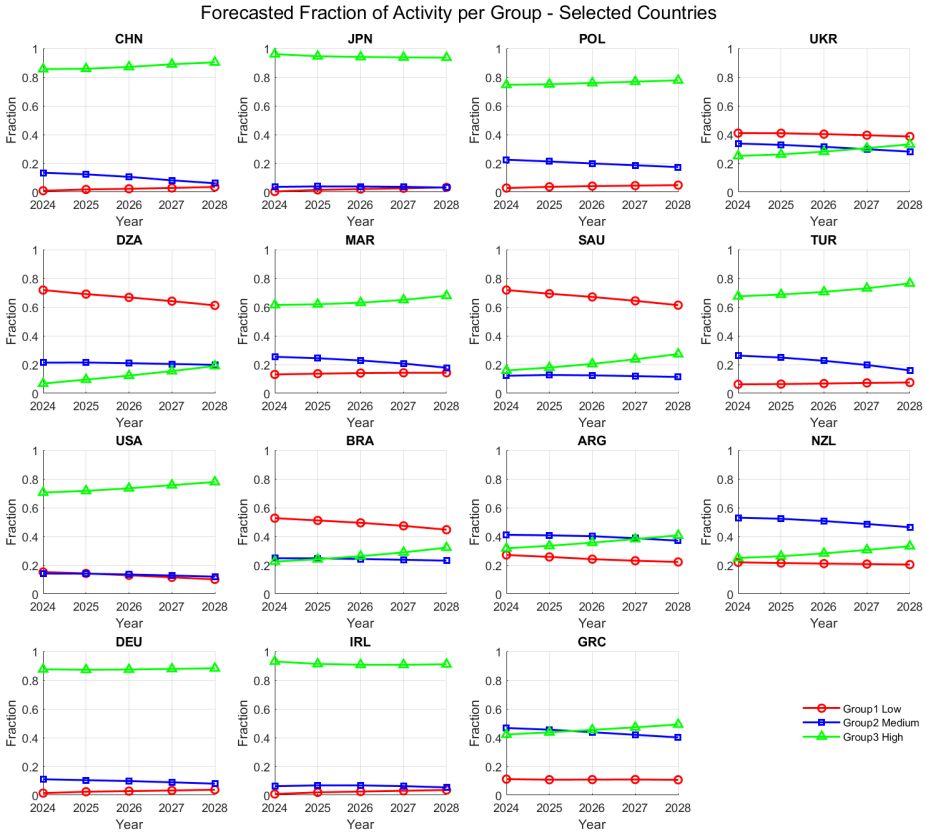


Figure 7: **Forecasted Evolution of Economic Structure for Selected Countries.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

(TUR), Morocco (MAR), and Saudi Arabia (SAU) show a clear upward trend in the fraction of activity in high-complexity sectors, coupled with a decline in medium or low-complexity activities. This suggests a predicted path of economic diversification and sophistication. The model also captures countries with different initial structures. Economies like Algeria (DZA) and Brazil (BRA) are shown to be initially dominated by low-complexity sectors, while Argentina (ARG) and Greece (GRC) exhibit a larger share in the medium-complexity group. For Ukraine (UKR), the model forecasts a competitive dynamic between all three groups, with a gradual increase in the high-complexity share. These varied predictions underscore the model's ability to capture country-specific initial conditions and project heterogeneous paths of structural change.

5 Conclusion

In this paper, we introduced a novel framework for modeling the continuous-time dynamics of national structural transformation. By integrating flow matching, Schrödinger Bridges, and stochastic differential equations, we learned a vector field representing the systematic forces and stochastic shocks that govern the evolution of countries' export portfolios in the high-dimensional product space. This approach moves beyond traditional static or discrete-time analyses to provide a continuous, probabilistic, and predictive account of economic development.

Our empirical results provide multifaceted support for the validity and utility of this approach. First, we demonstrated that the model's learned drift vector is not an abstract construct but is significantly correlated with realized year-over-year changes in sectoral composition, capturing the influence of both global trends and major economic shocks like the COVID-19 pandemic. Furthermore, the model implicitly learns fundamental principles of economic development without them being explicitly encoded. The learned drift directs countries towards "neighborly" diversification opportunities consistent with the Product Space's principle of proximity, with the nature of this path-dependency differing systematically between developed and developing regions.

Most importantly, we established a clear link between the model's output and economic outcomes. A panel regression analysis revealed that the strategic quality of the learned drift, its alignment with a complexity-enhancing frontier, is a strong and statistically significant predictor of future growth in a country's economic complexity. In contrast, the sheer magnitude of the prescribed change was negatively associated with complexity growth, suggesting that well-directed, incremental adjustments are more effective than large, disruptive shocks. Finally, the model's forecasting capabilities were demonstrated by generating plausible, multi-year trajectories for the economic structures of diverse nations, correctly capturing the stability of advanced economies and the dynamic transformations of emerging ones.

The implications of this work are twofold. For economic theory, our framework offers a new way to quantify and analyze path dependence, moving from static network measures to continuous, probabilistic flows. It provides a "weather map" of economic development, revealing the underlying currents pushing national economies. For policymakers, the learned drift provides a data-driven diagnostic tool. The "Alignment with Frontier" metric, in particular, could help evaluate whether a country's current development trajectory is optimally oriented towards long-term sophistication, offering a forward-looking complement to traditional backward-looking indicators.

Several limitations of this study suggest avenues for future research. Our analysis was based on 28 aggregated sectors; future work could apply this methodology to more granular product-level data to capture finer-grained dynamics. Furthermore, the current model treats the drift as exogenous. An important extension would be to endogenize the drift by modeling it as a function of country-specific characteristics, such as policy variables, investment in R&D, or institutional quality. This would allow for true counterfactual policy analysis. By addressing these aspects, future research can build upon this dynamic systems perspective to develop an even more comprehensive and predictive science of economic development.

Acknowledgements

The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

References

M. S. Albergo and E. Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

C. Azzoni. Complexity and the measurement of structural transformation. *Structural Change and Economic Dynamics*, 23(2):139–152, 2012.

C. Chen, D. Duvenaud, and J.-H. Jacobsen. Neural spline flows. *Advances in Neural Information Processing Systems*, 2020.

R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 2018.

Y. Chen, T. T. Georgiou, and M. Pavon. On the relation between Schrödinger bridge and optimal transport. *Journal of Optimization Theory and Applications*, 2016.

M. Cristelli, A. Tacchella, and L. Pietronero. Measuring the intangibles of economic complexity. *PLoS ONE*, 10(2):e0117179, 2015.

M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

V. De Bortoli, J. Thornton, J. Heng, and A. Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. *arXiv preprint arXiv:2106.01357*, 2021.

K. Fatras, Y. Zine, R. Flamary, R. Gribonval, and N. Courty. Learning with minibatch wasserstein: asymptotic and gradient properties. *arXiv preprint arXiv:1910.04091*, 2019.

K. Fatras, T. Séjourné, R. Flamary, and N. Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *International conference on machine learning*, pages 3186–3197. PMLR, 2021.

R. Hausmann and C. A. Hidalgo. *The Atlas of Economic Complexity: Mapping Paths to Prosperity*. MIT Press, 2014.

C. A. Hidalgo. The economics of the product space: Structural transformation, path dependence and capabilities diffusion. *Journal of Economic Dynamics and Control*, 93:1–33, 2018.

C. A. Hidalgo and R. Hausmann. The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26):10570–10575, 2009.

C. A. Hidalgo, B. Klinger, A.-L. Barabási, and R. Hausmann. The product space conditions the development of nations. *Science*, 317(5837):482–487, 2007.

L. Kantorovich. On the transfer of masses (in russian). In *Doklady Akademii Nauk*, volume 37, pages 227–229, 1942.

Y. Lipman, T. Hagay, and A. Belyaev. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Y. Lipman, M. Havasi, P. Holderrieth, N. Shaul, M. Le, B. Karrer, R. T. Chen, D. Lopez-Paz, H. Ben-Hamu, and I. Gat. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*, 2024.

X. Liu, C. Gong, and Q. Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

L. Pietronero, M. Cristelli, A. Tacchella, and A. Gabrielli. Economic complexity: a new viewpoint for understanding wealth and poverty. *Physica A: Statistical Mechanics and its Applications*, 471:131–146, 2017.

A.-A. Pooladian, H. Ben-Hamu, C. Domingo-Enrich, B. Amos, Y. Lipman, and R. T. Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.

A. Spelta. Density-based machine learning model averaging for inflation forecasting. *Journal of the Royal Statistical Society Series C: Applied Statistics*, page qlaf048, 2025.

A. Spelta, P. Pagnottoni, and N. Pecora. Short-term forecasting with optimal transport. *Journal of the Royal Statistical Society Series A: Statistics in Society*, page qnaf036, 2025.

A. Tacchella, M. Cristelli, G. Caldarelli, A. Gabrielli, and L. Pietronero. A new metrics for countries' fitness and products' complexity. *Scientific Reports*, 2:723, 2012.

A. Tacchella, M. Cristelli, G. Caldarelli, A. Gabrielli, and L. Pietronero. Economic fitness and complexity of countries and products. *Nature Scientific Reports*, 3:1–6, 2013.

A. Tong, J. Choi, and A. Belyaev. Improved flow matching for generative modeling. *arXiv preprint arXiv:2302.00482*, 2023.

C. Villani. *Optimal Transport: Old and New*. Springer, 2009.

C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.

P. Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.

Appendices

A Additional results on the distribution of drift magnitudes

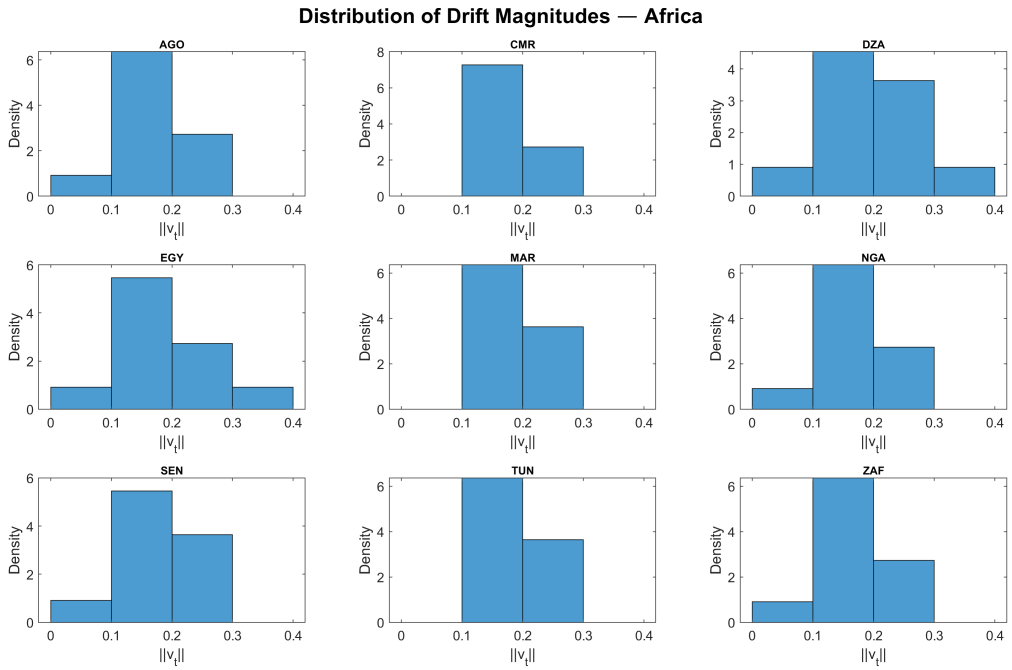


Figure A1: **Distribution of instantaneous drift magnitude across African countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

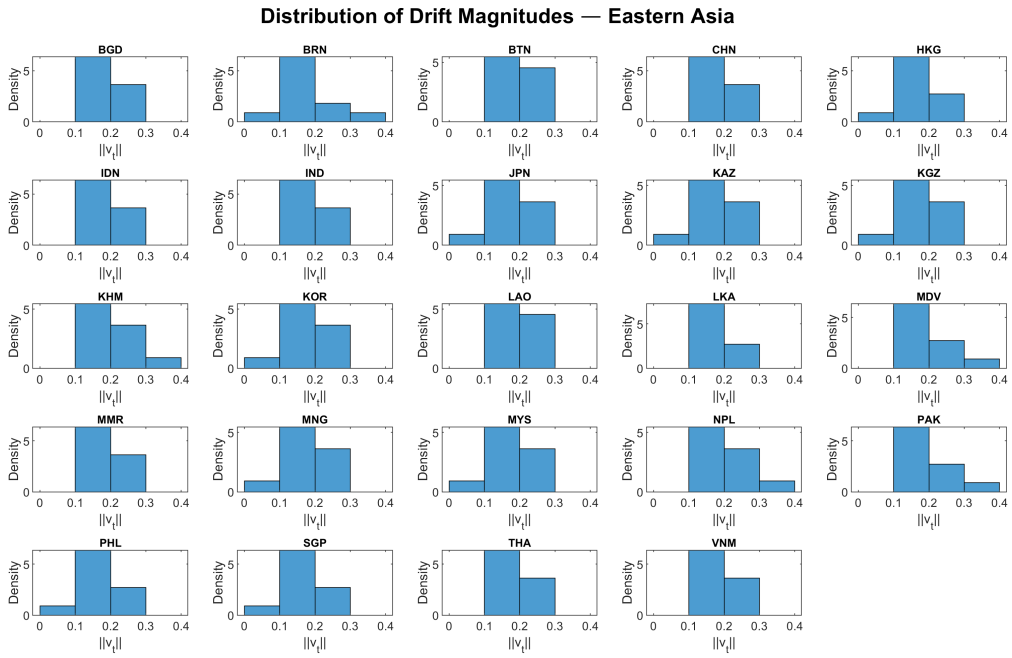


Figure A2: **Distribution of instantaneous drift magnitude across Eastern Asian countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

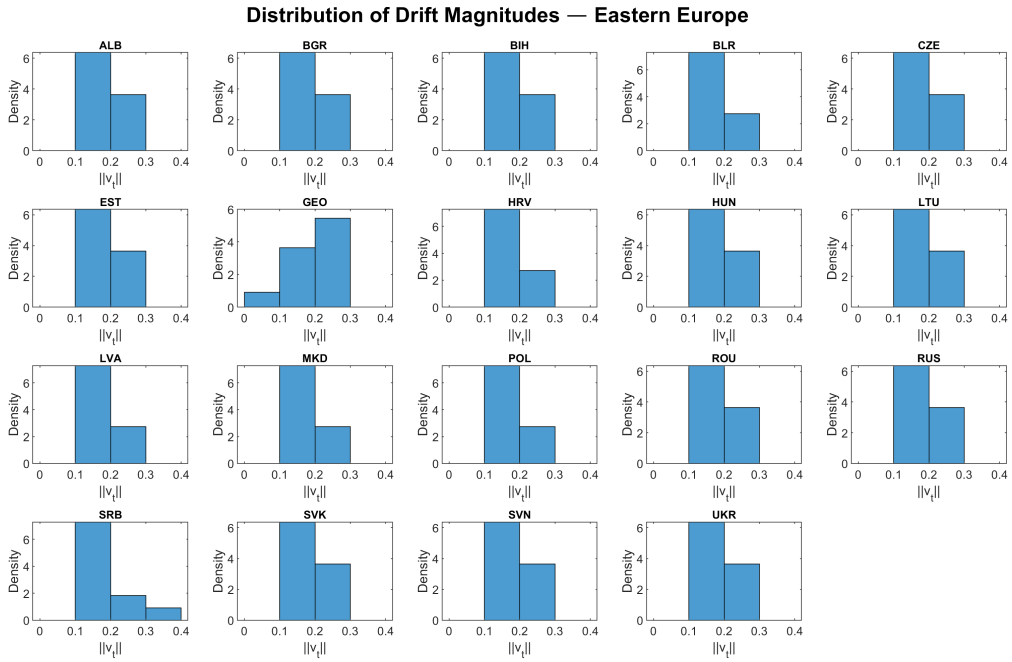


Figure A3: **Distribution of instantaneous drift magnitude across Eastern Europe countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

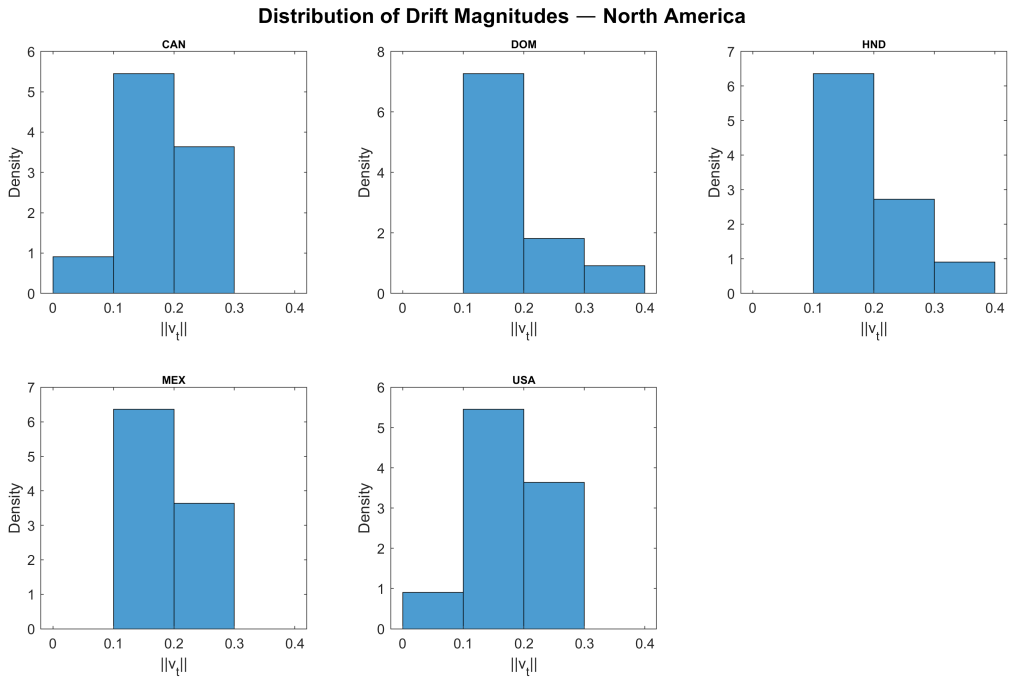


Figure A4: **Distribution of instantaneous drift magnitude across North America countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

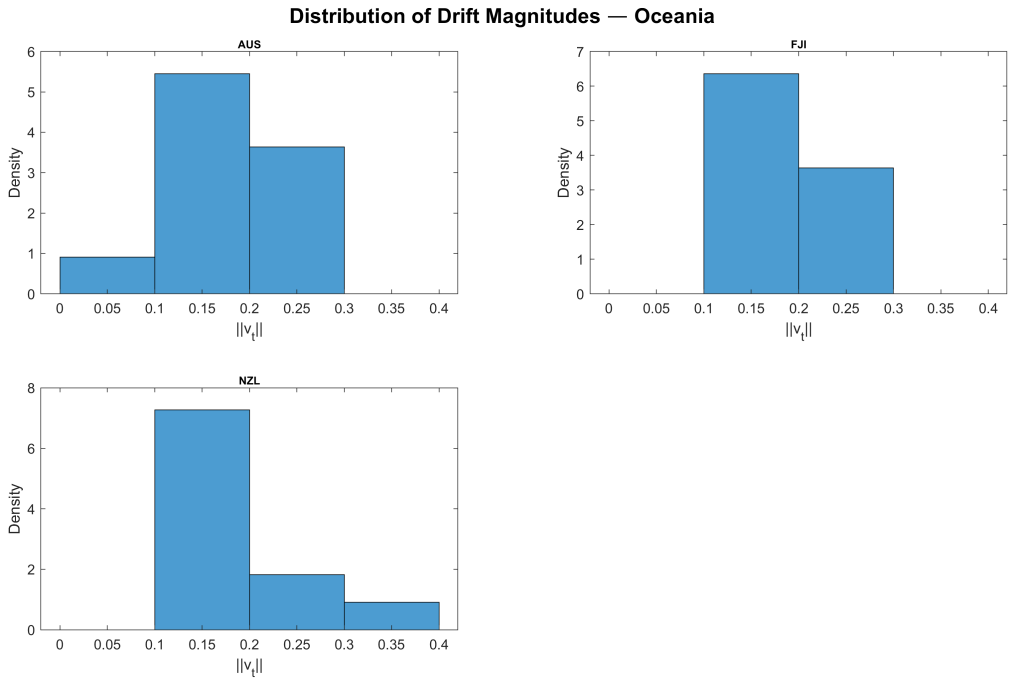


Figure A5: **Distribution of instantaneous drift magnitude across Oceanian countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

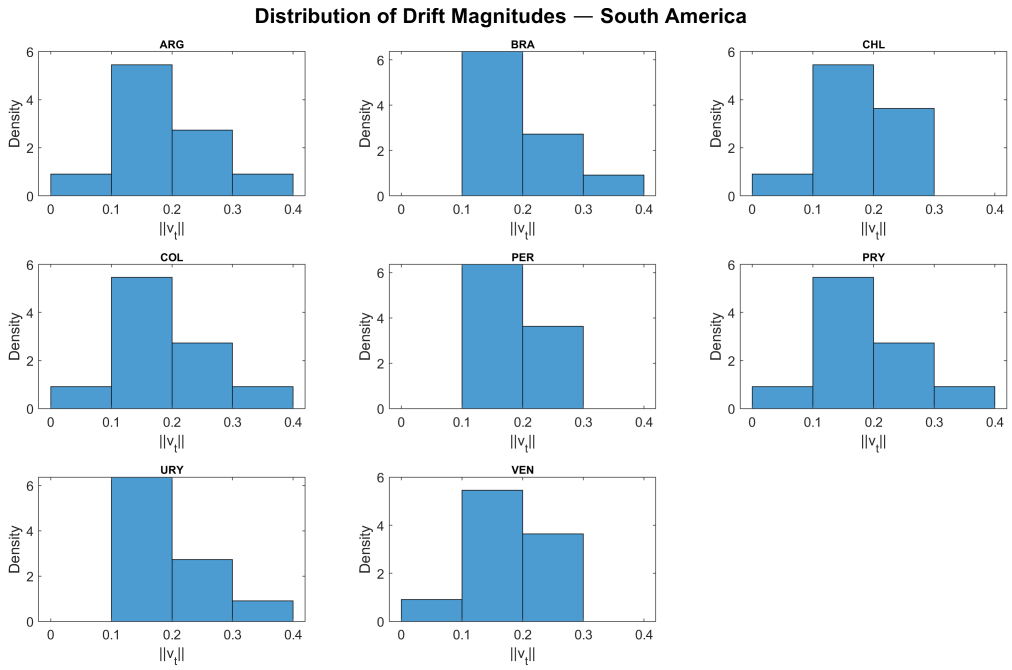


Figure A6: **Distribution of instantaneous drift magnitude across South America countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

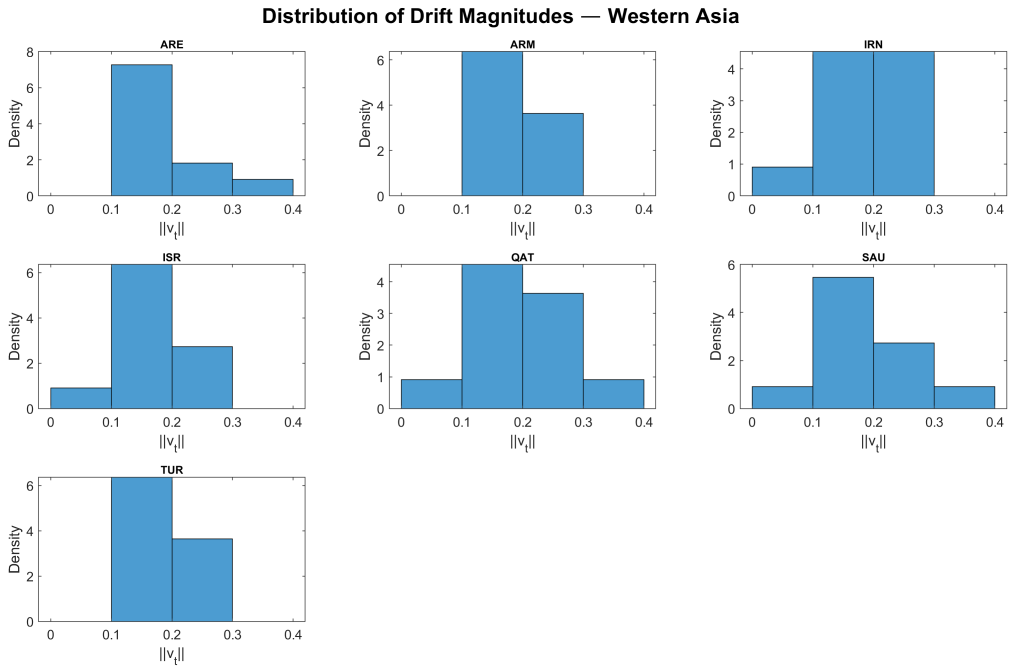


Figure A7: **Distribution of instantaneous drift magnitude across Western Asia countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

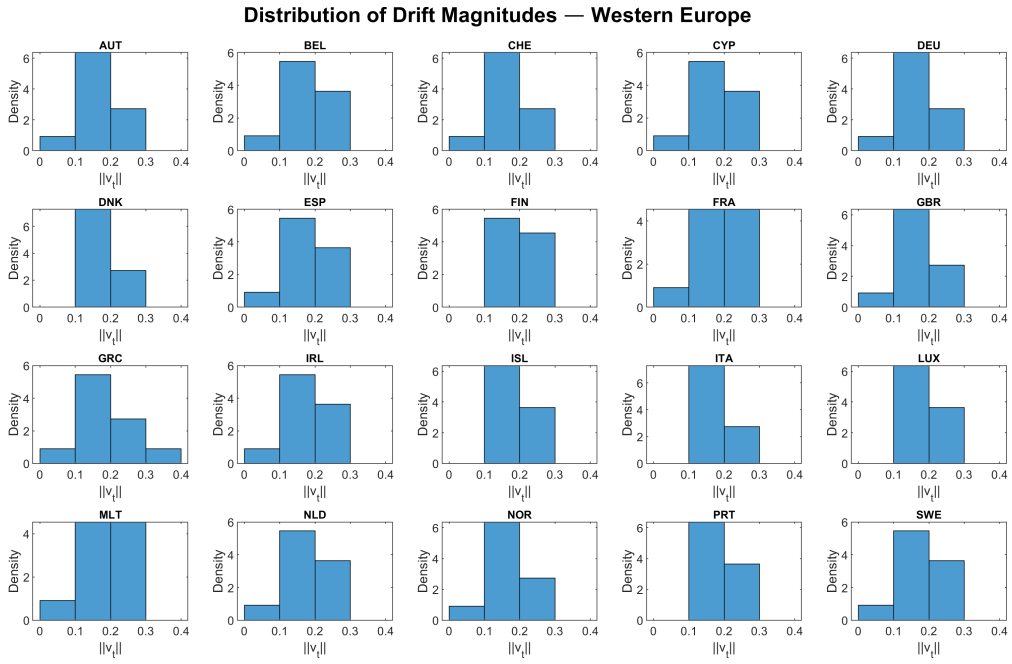


Figure A8: **Distribution of instantaneous drift magnitude across Western Europe countries.** Each panel shows the empirical distribution of $||v_t^i||$, the norm of the learned, time-dependent drift of a country's export portfolio in product space, estimated via flow matching and refined with a Schrödinger Bridge. The magnitude indexes the expected speed of systematic structural change per unit time; larger values indicate faster reconfiguration but not direction or upgrading.

B Additional results on the temporal evolution of sectoral drift

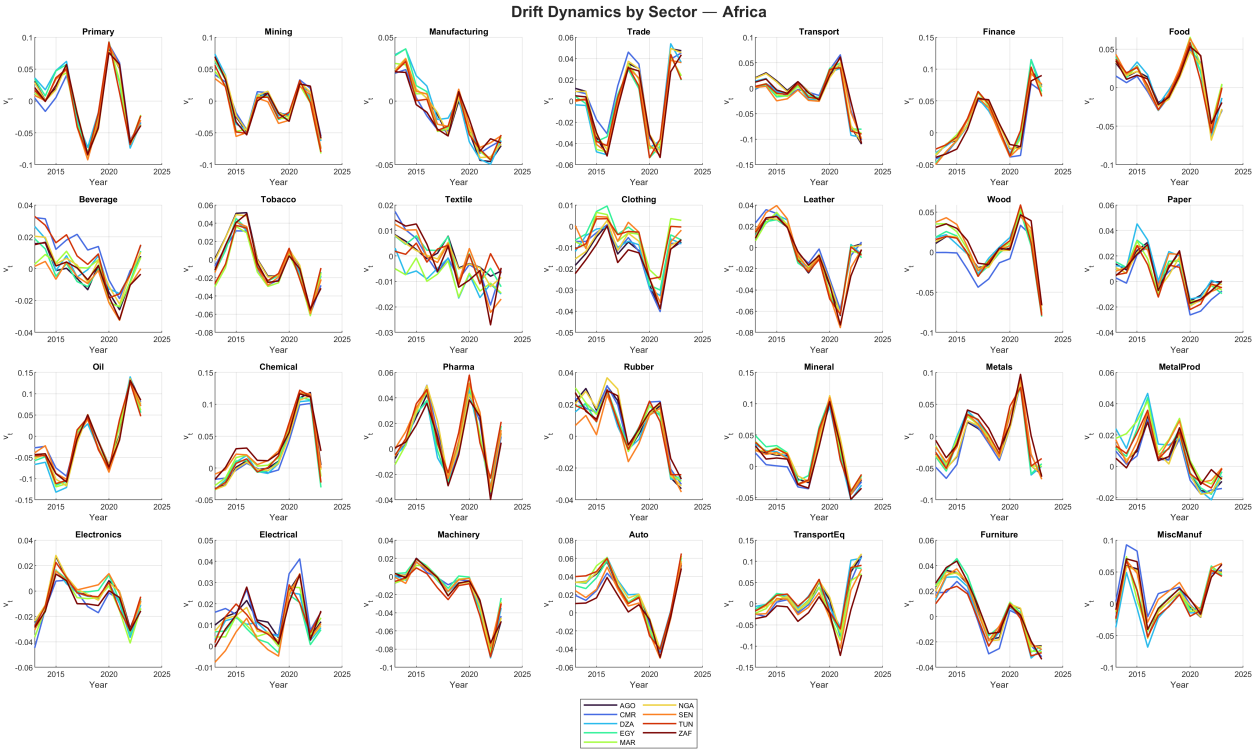


Figure B1: **Temporal evolution of sectoral drift for African countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

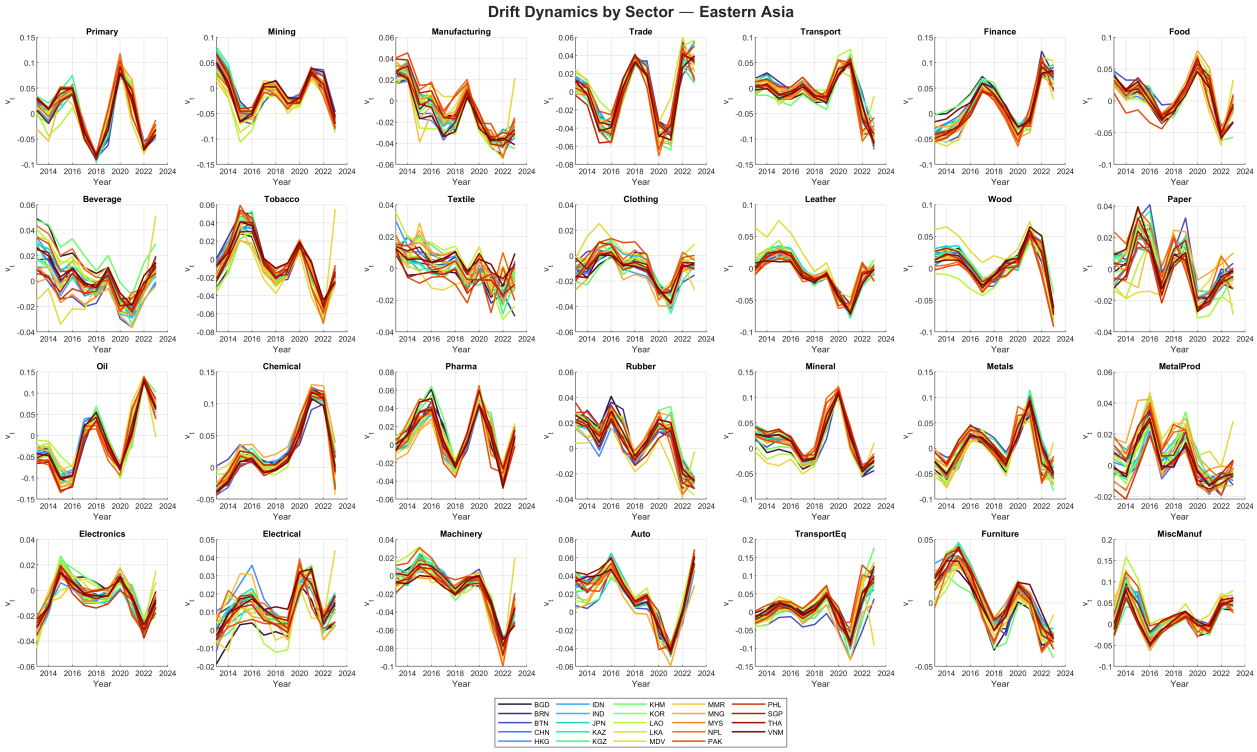


Figure B2: **Temporal evolution of sectoral drift for Eastern Asia countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

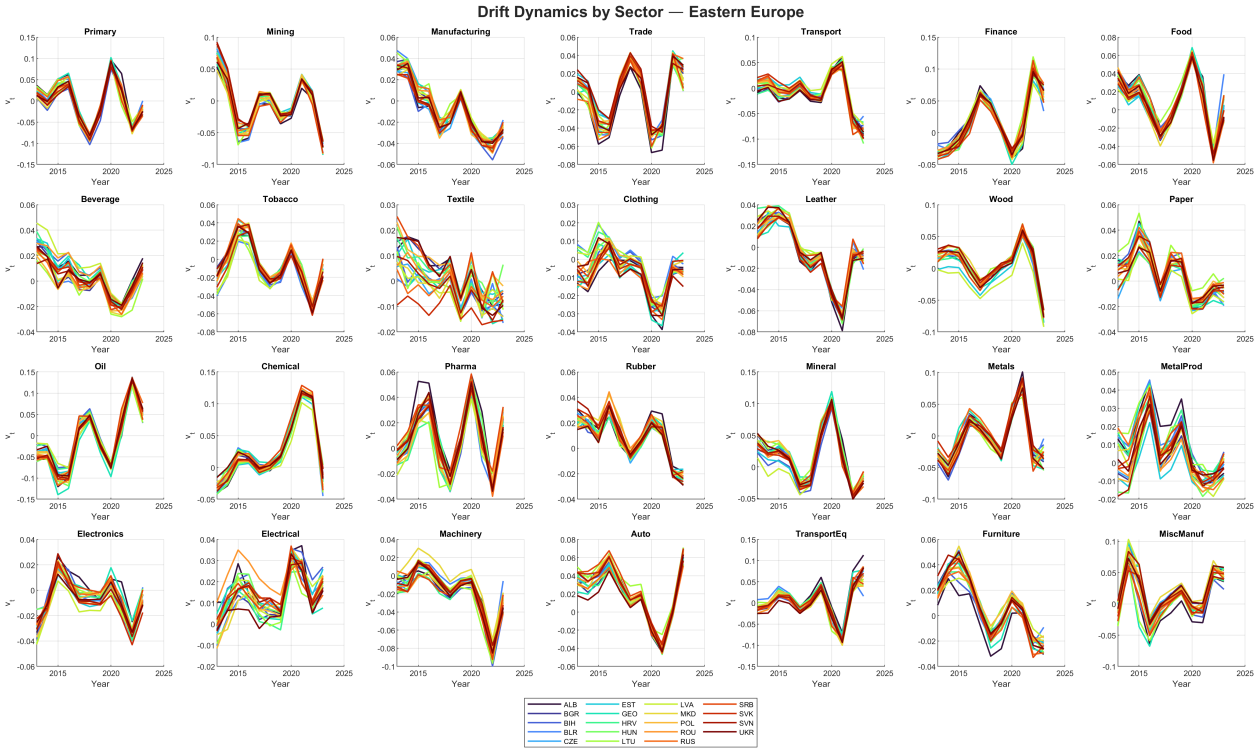


Figure B3: **Temporal evolution of sectoral drift for Eastern Europe countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

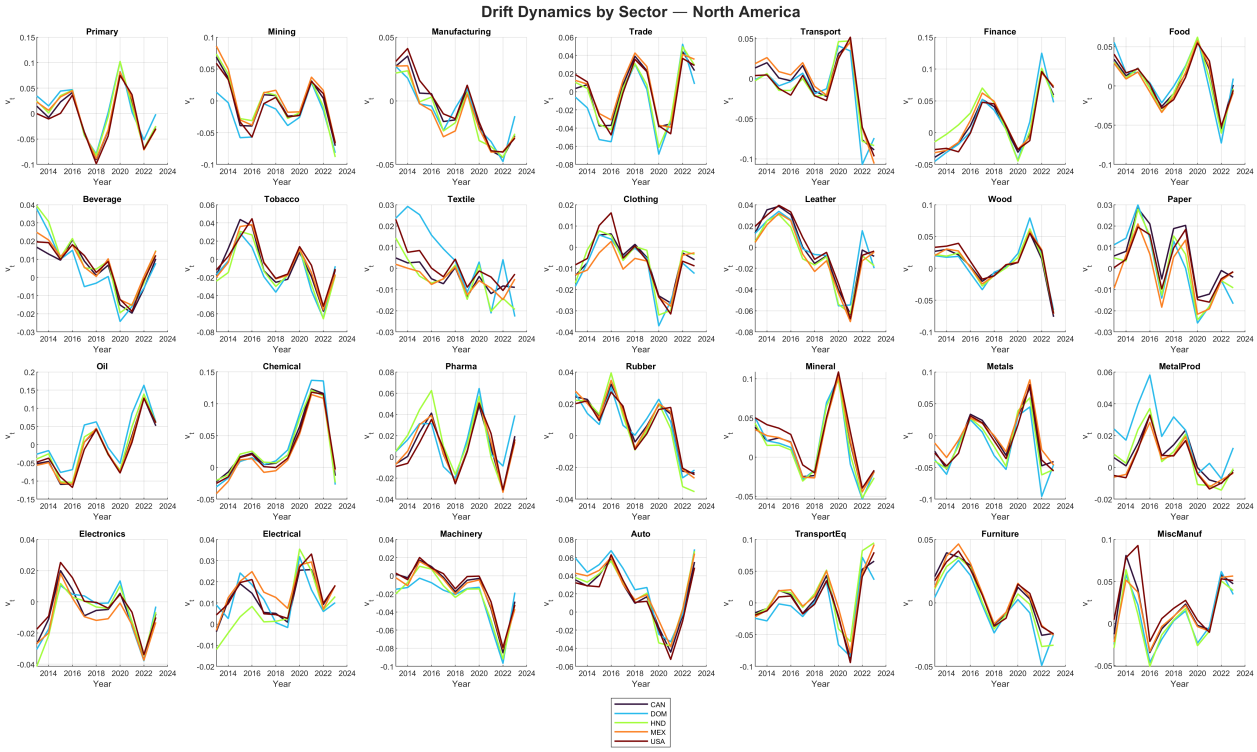


Figure B4: **Temporal evolution of sectoral drift for North America countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

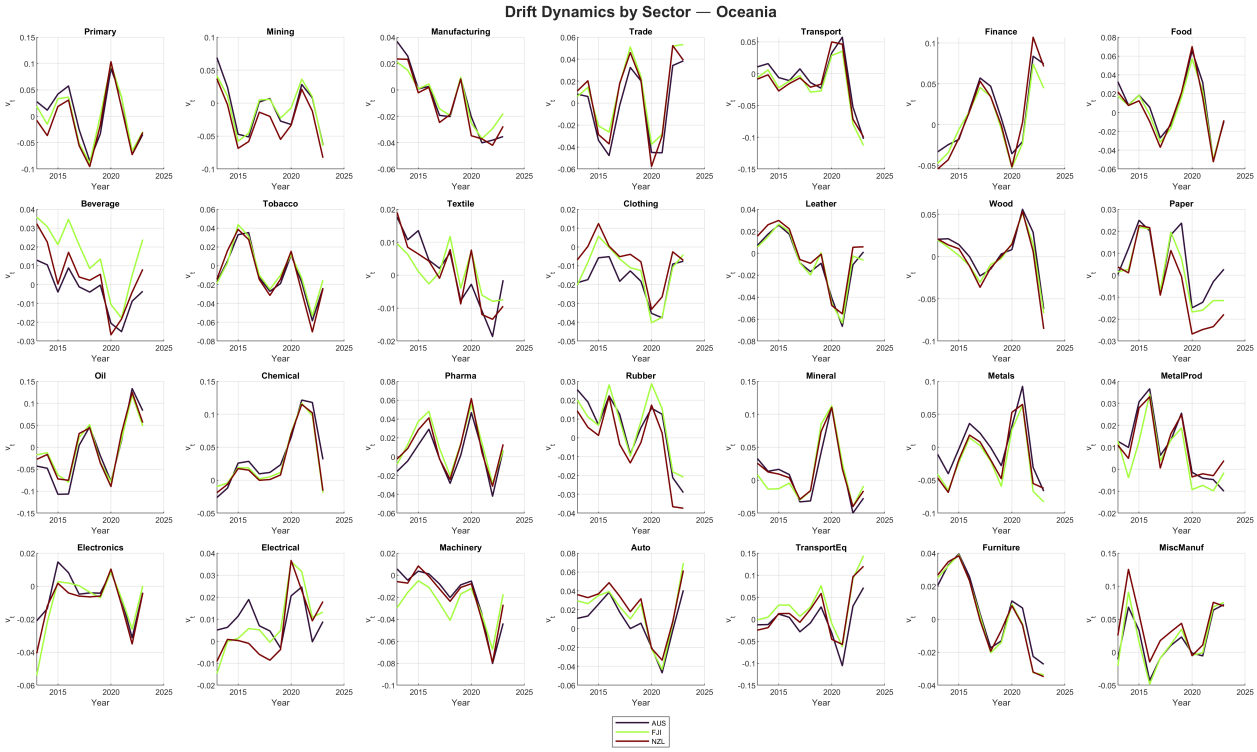


Figure B5: **Temporal evolution of sectoral drift for Oceanian countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

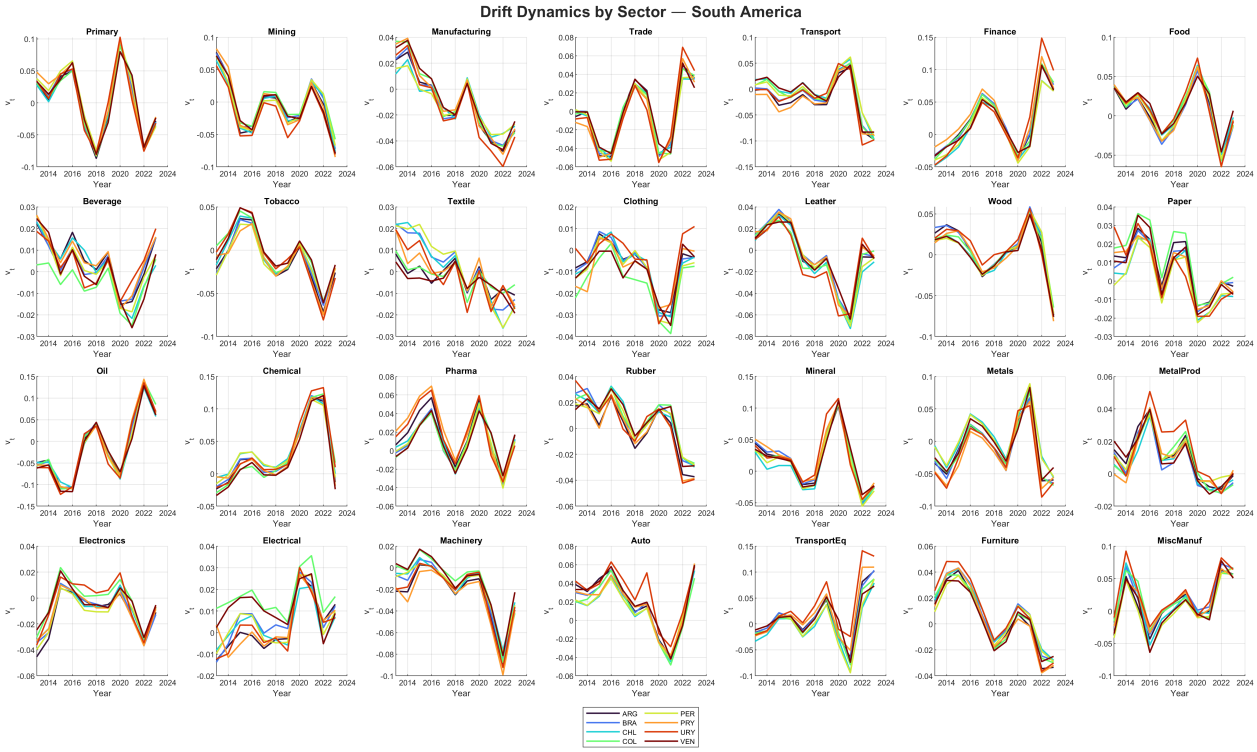


Figure B6: **Temporal evolution of sectoral drift for South America countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

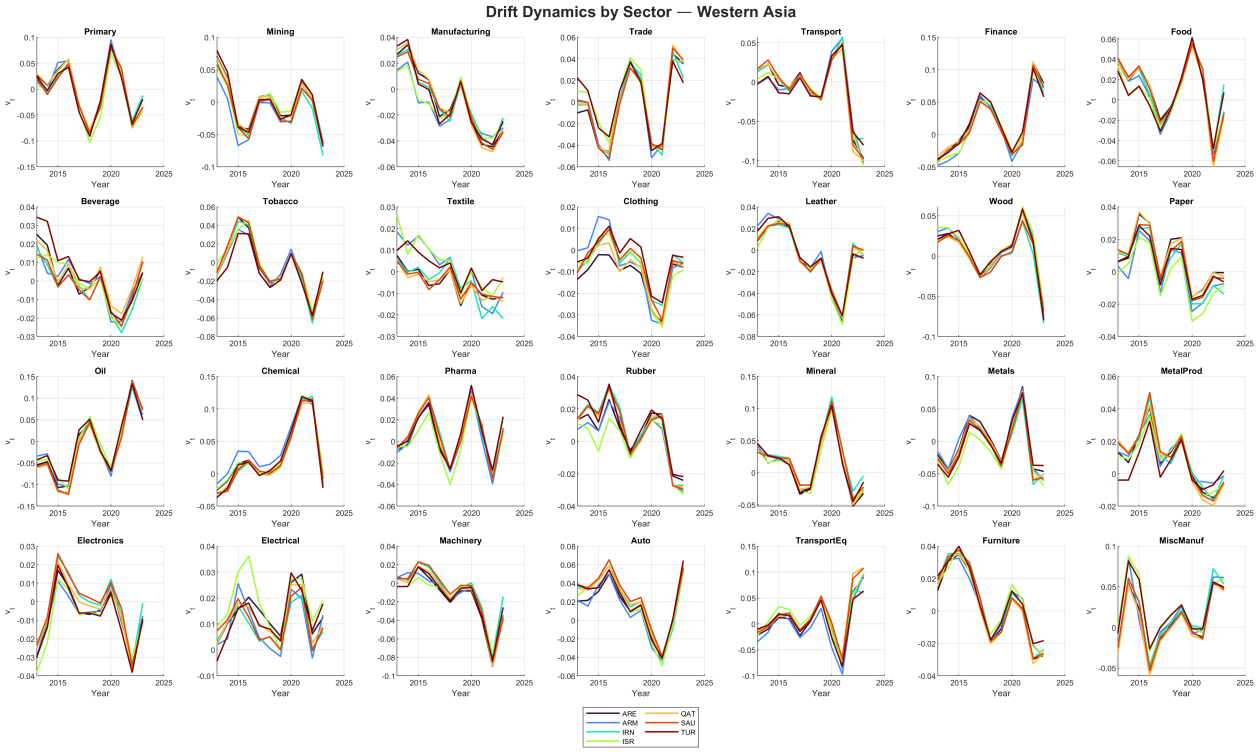


Figure B7: **Temporal evolution of sectoral drift for Western Asia countries.** Each panel displays the evolution of the learned drift vector, $\nu^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

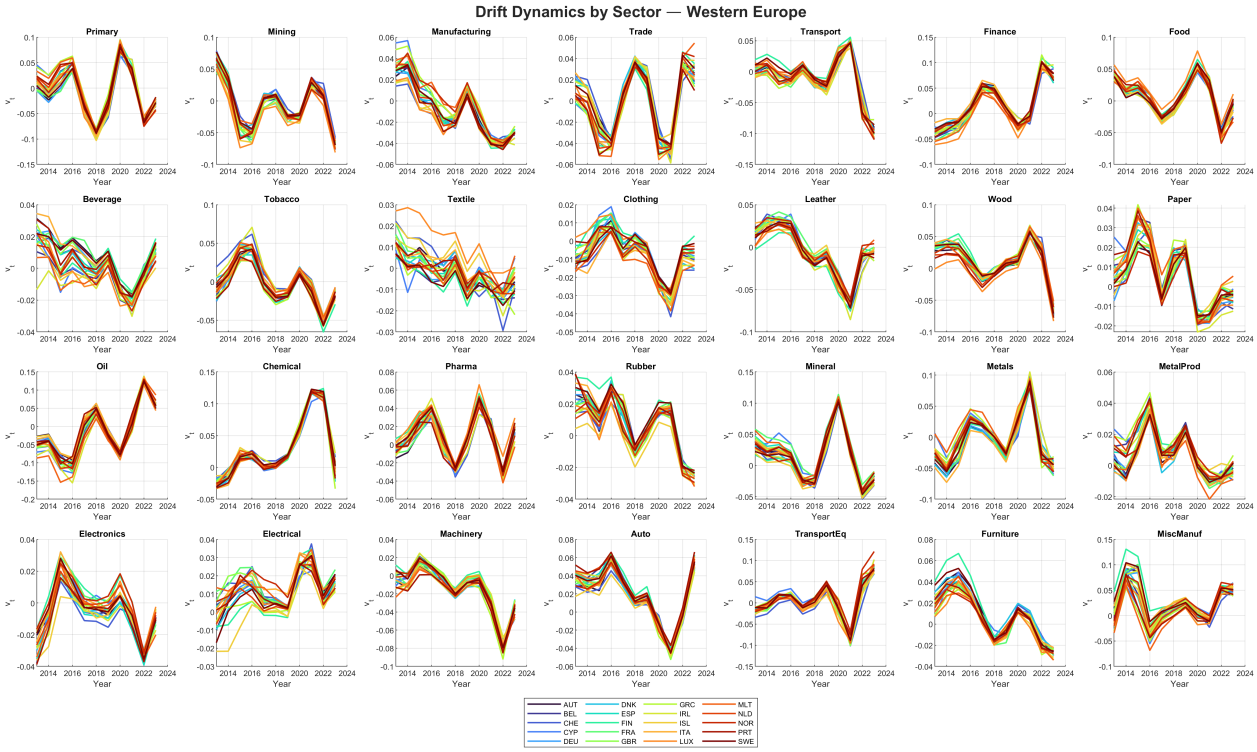


Figure B8: **Temporal evolution of sectoral drift for Western Europe countries.** Each panel displays the evolution of the learned drift vector, $v^i_{t,p}$, for a specific economic sector. The drift (y-axis) represents the expected velocity of change in a country's productive capacity, learned from historical data using the proposed stochastic process model. Each colored line corresponds to an individual country.

C Computation of Proximity and Product Complexity Index (PCI)

The structure of global production can be represented as a network of related products, where links capture how likely it is that two products require similar capabilities. This relatedness is quantified through product proximity matrices, which measure the empirical probability that two products are co-exported by the same countries. The topology of this network, known as the product space (Hidalgo et al., 2007), captures the underlying structure of productive knowledge in the world economy.

Building on this network, the Product Complexity Index (PCI) quantifies the sophistication of each product based on the diversity and knowledge intensity of the countries that export it. Following the methodology proposed by Hidalgo and Hausmann (2009), we compute both the product proximities and the PCI using export data. The procedure proceeds in four steps: (i) estimation of revealed comparative advantage, (ii) construction of the binary specialization matrix, (iii) computation of the product proximity network, and (iv) derivation of the PCI using the Method of Reflections.

C.1 Revealed Comparative Advantage (RCA)

For each country $i \in \{1,\dots,N\}$ (where $N = 95$), product $p \in \{1,\dots,28\}$, and year $t \in \{2012,\dots,2023\}$, we compute the index of Revealed Comparative Advantage (RCA) as:

$$\mathrm{RCA}^{i}_{p,t} = \frac{\tilde{x}^{i}_{p,t} \big/ \sum_{p'=1}^{28} \tilde{x}^{i}_{p',t}}{\sum_{j=1}^{N} \tilde{x}^{j}_{p,t} \big/ \sum_{j=1}^{N} \sum_{p'=1}^{28} \tilde{x}^{j}_{p',t}}, \tag{5}$$

where $\tilde{x}^{i}_{p,t}$ is the export value of product p by country i in year t. A country is said to have a revealed comparative advantage in product p at year t when $\mathrm{RCA}^{i}_{p,t} \geq 1$.

We define a binary country–product–year specialization matrix $M^{i}_{p,t}$ such that:

$$M^{i}_{p,t} = \begin{cases} 1, & \text{if } \mathrm{RCA}^{i}_{p,t} \geq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

C.2 Product Proximity Network

We construct time-varying Product Spaces by calculating a sequence of proximity matrices $\{\Phi_t\}_{t=1}^{T}$. For each year t, we compute the proximity matrix Φ_t using the specialization matrix $M^{i}_{p,t}$ from year t. The proximity between any two products p and p' is defined as:

$$\Phi_{t,pp'} = \min\Big(P(M^{i}_{p,t} = 1 \mid M^{i}_{p',t} = 1),\ P(M^{i}_{p',t} = 1 \mid M^{i}_{p,t} = 1)\Big),$$

where the conditional probabilities are estimated empirically over the set of countries $i = 1,\dots,N$. This yields a time-specific map of technological relatedness that reflects the product space structure as it existed in the year immediately preceding the forecast.

C.3 The Method of Reflections for Time-Varying PCI

The Product Complexity Index (PCI) is computed for each year using the same year-by-year approach to maintain temporal consistency. For each year t, we compute $\mathrm{PCI}_{p,t}$ using the specialization matrix $M^i_{p,t}$.

Initial conditions (order 0).

$$k^i_{0,t} = \sum_{p=1}^{28} M^i_{p,t}, \qquad k_{p,0,t} = \sum_{i=1}^{N} M^i_{p,t}. \tag{7}$$

Iterative refinement (order n). For each iteration $n \geq 1$:

$$k^i_{n,t} = \frac{1}{k^i_{0,t}} \sum_{p=1}^{28} M^i_{p,t}\, k_{p,n-1,t}, \tag{8}$$

$$k_{p,n,t} = \frac{1}{k_{p,0,t}} \sum_{i=1}^{N} M^i_{p,t}\, k^i_{n-1,t}. \tag{9}$$

This recursive process propagates information between countries and products using only data from year t.

In matrix form, the relationships can be expressed as:

$$\tilde{M}_{ii',t} = \sum_{p=1}^{28} \frac{M^i_{p,t} M^{i'}_{p,t}}{k_{p,0,t}}, \tag{10}$$

$$\tilde{M}_{pp',t} = \sum_{i=1}^{N} \frac{M^i_{p,t} M^i_{p',t}}{k^i_{0,t}}. \tag{11}$$

The eigenvector of $\tilde{M}_{pp',t}$ associated with the second-largest eigenvalue defines the relative sophistication of each product at time t, denoted by $K_{p,t}$.

Finally, the time-varying Product Complexity Index is obtained by standardizing $K_{p,t}$ within each year:

$$\mathrm{PCI}_{p,t} = \frac{K_{p,t} - \langle K_t \rangle}{\sigma(K_t)}, \tag{12}$$

where $\langle K_t \rangle$ is the mean and $\sigma(K_t)$ the standard deviation of $\{K_{p,t}\}_{p=1}^{28}$ for year t. This yields a sequence of PCI vectors $\{\mathrm{PCI}_{p,t}\}$ that reflect the evolving complexity landscape of the product space.

D Additional results for forecasting Sectoral activity composition

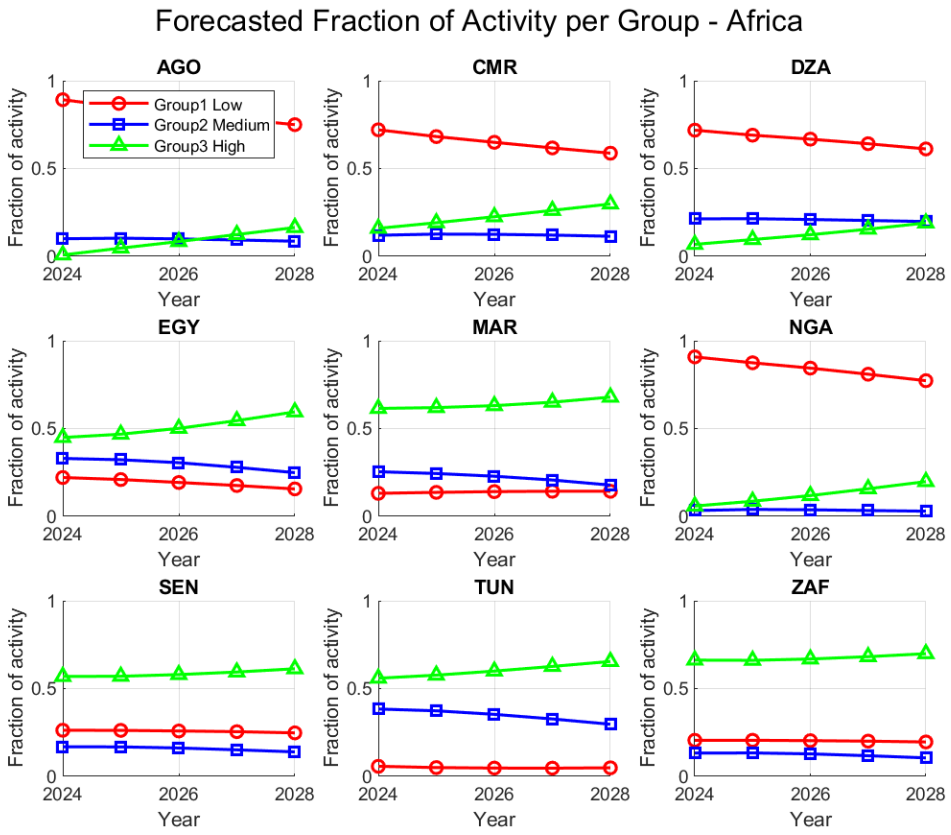


Figure D1: **Forecasted Evolution of Economic Structure for African Countries.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

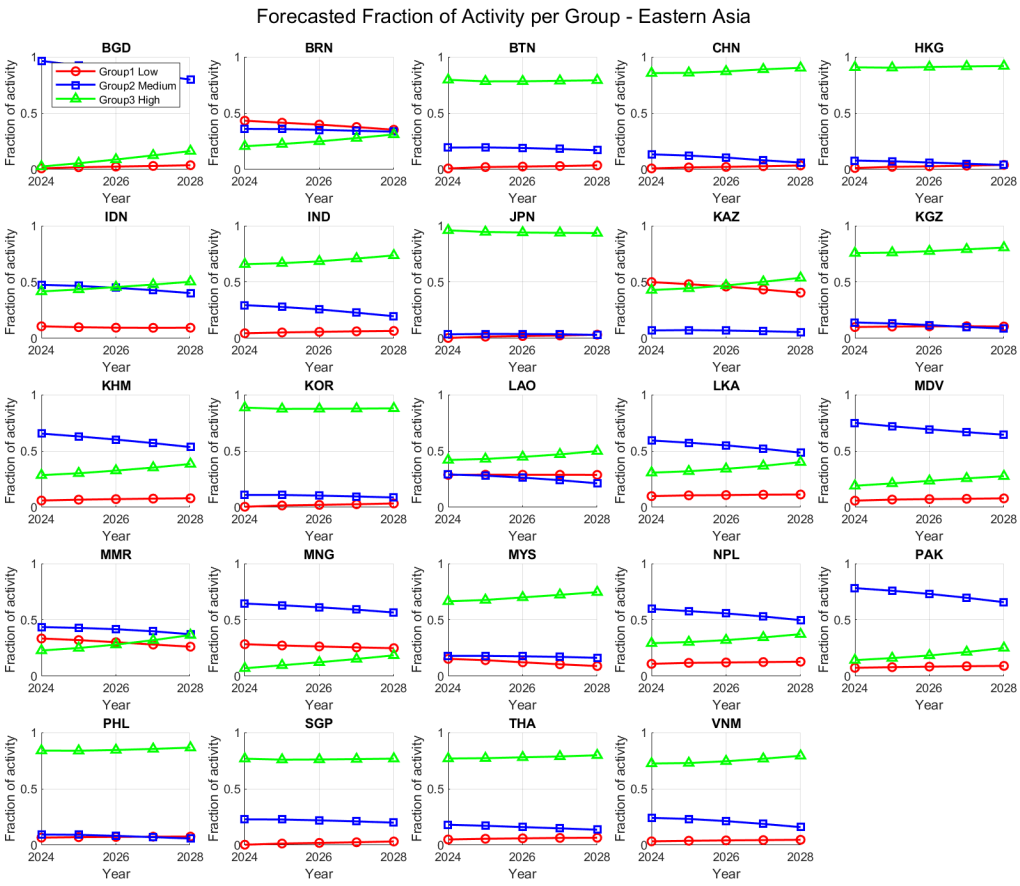


Figure D2: **Forecasted Evolution of Economic Structure for Eastern Asia.** Each subplot shows the expected fraction of total economic activity ($f_{t,g}^{i}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

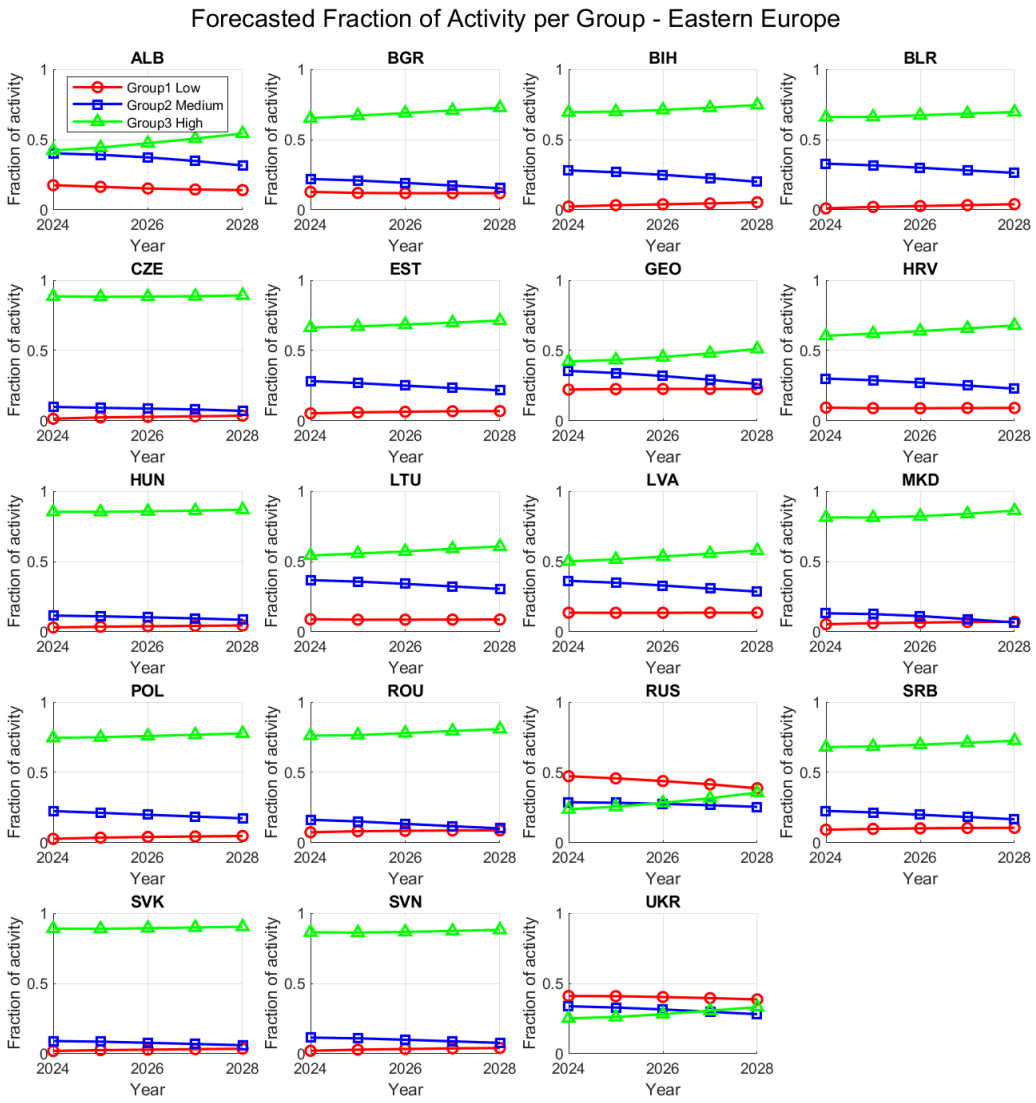


Figure D3: **Forecasted Evolution of Economic Structure for Eastern Europe.** Each subplot shows the expected fraction of total economic activity ($f_{t,g}^{i}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

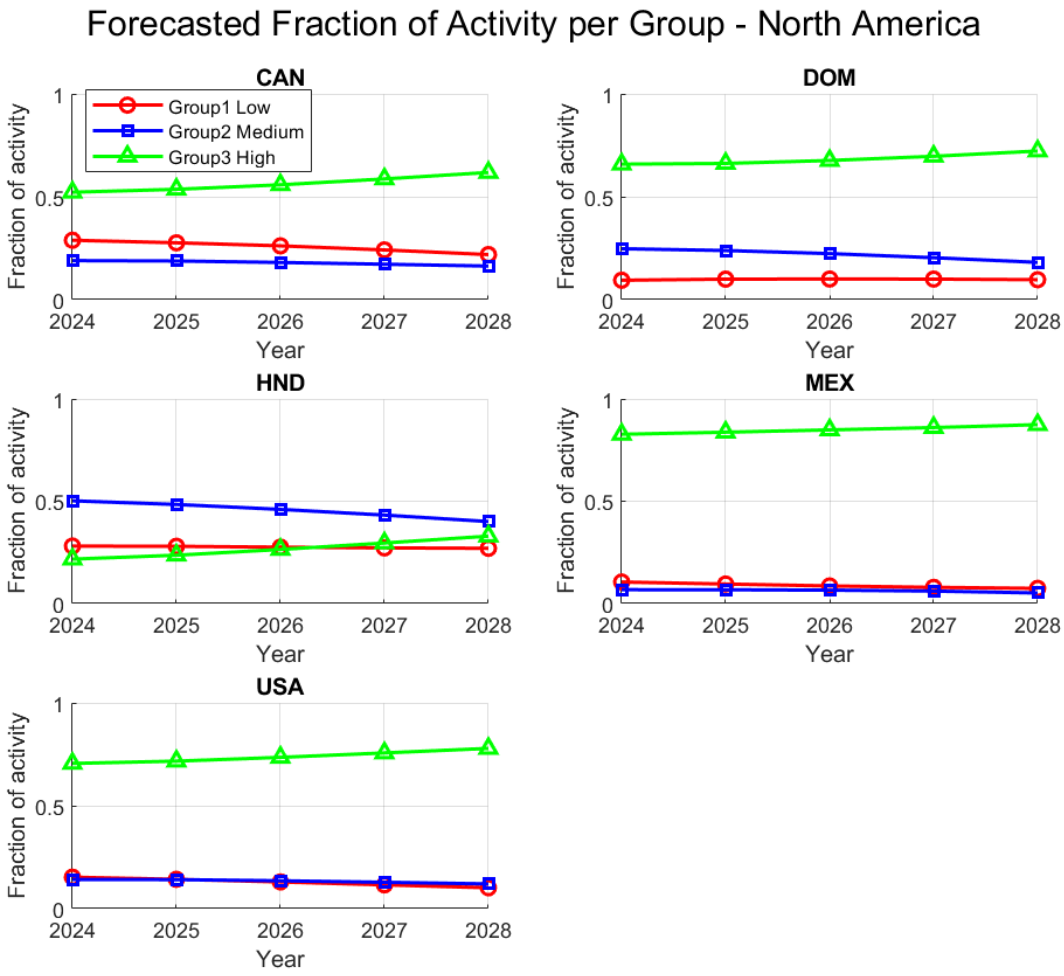


Figure D4: **Forecasted Evolution of Economic Structure for North America.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

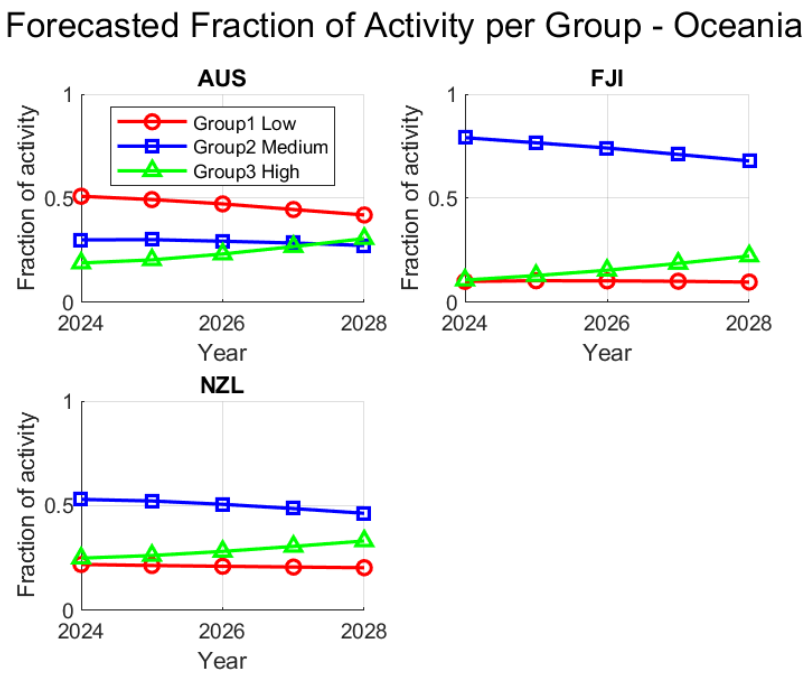


Figure D5: **Forecasted Evolution of Economic Structure for Oceania.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

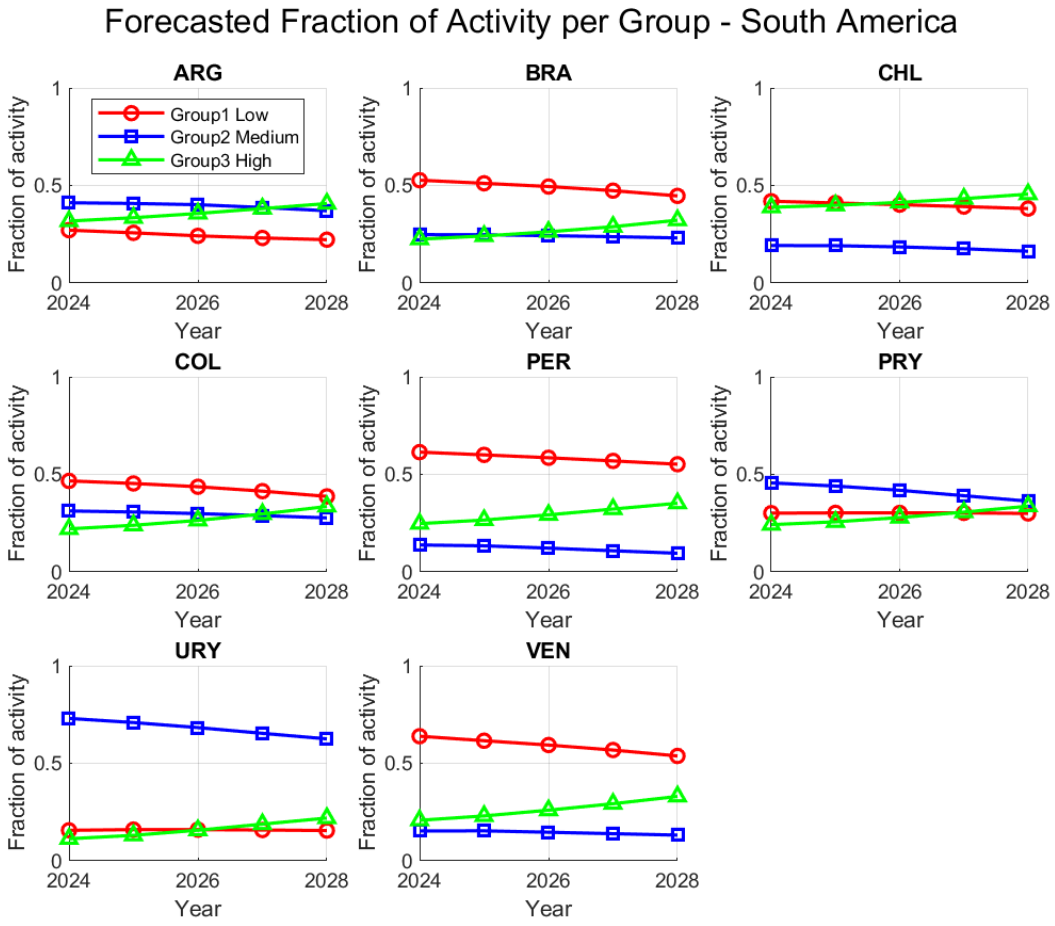


Figure D6: **Forecasted Evolution of Economic Structure for South America.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

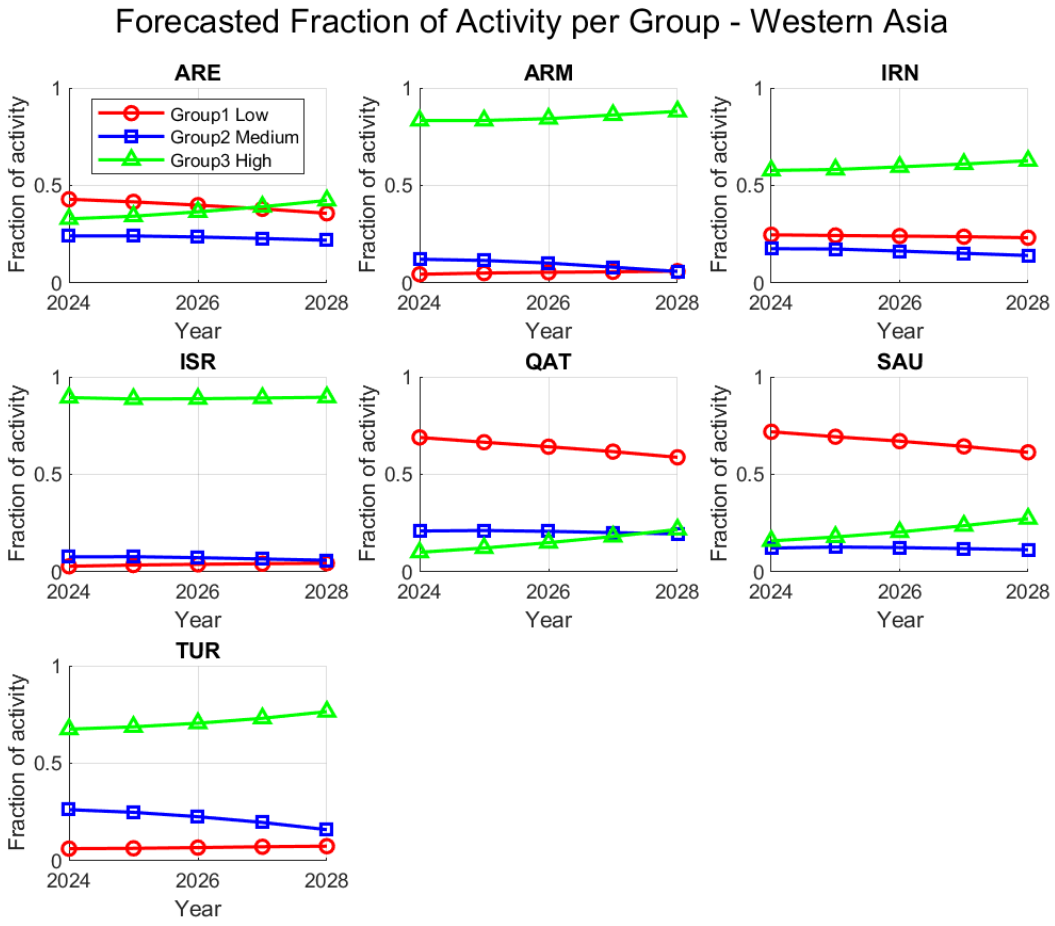


Figure D7: **Forecasted Evolution of Economic Structure for Western Asia.** Each subplot shows the expected fraction of total economic activity ($f^i_{t,g}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.

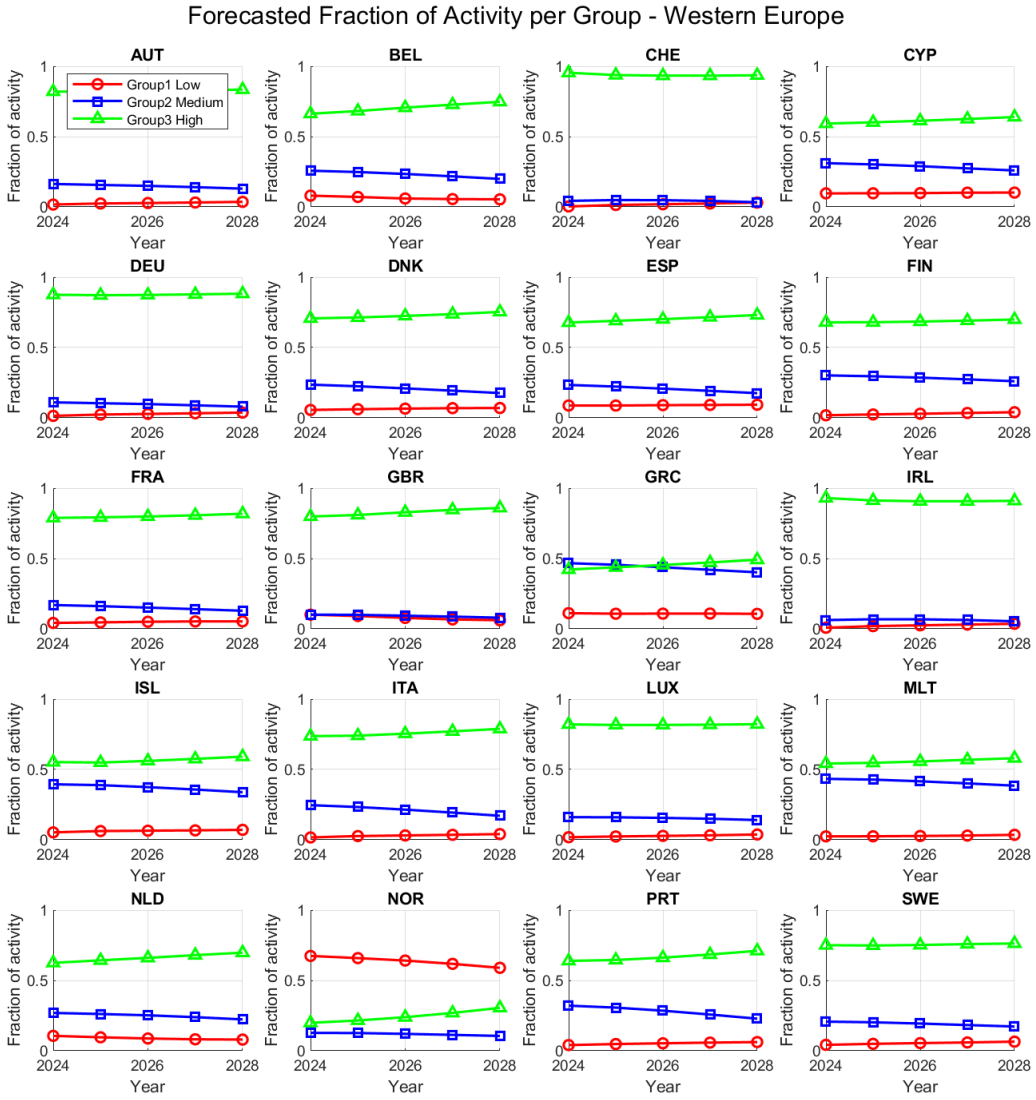


Figure D8: **Forecasted Evolution of Economic Structure for Western Europe.** Each subplot shows the expected fraction of total economic activity ($f_{t,g}^{i}$) allocated to low-complexity (Group 1, red circles), medium-complexity (Group 2, blue squares), and high-complexity (Group 3, green triangles) sectors over the 2024–2028 horizon. The trajectories represent the mean forecast from Monte Carlo simulations of the stochastic model, illustrating the predicted structural transformation dynamics for a diverse set of national economies.