INTERNATIONAL MONETARY FUND

Dissecting Medium-Term Growth Prospects for Asia

Natasha Che, Federico J. Diez, Anne Oeking, Weining Xin

WP/25/168

2025
AUG

WORKING PAPER

 WP/25/168

IMF Working Paper
Asia and Pacific Department

Dissecting Medium-Term Growth Prospects for Asia

Natasha Che, Federico J. Diez, Anne Oeking, Weining Xin

Authorized for distribution by Paulo A. Medas
August 2025

ABSTRACT: This paper explores Asia-Pacific's medium-term growth prospects using two approaches. First, growth accounting analysis and machine-learning estimation reveal how demographics, capital deepening, productivity, and human capital shaped Asia's growth. Second, an innovative algorithm forecasts growth by matching countries' current conditions with historically analogous periods using Dynamic Time Warping (DTW). Comparing pattern-based forecasts with traditional projections highlights economic convergence and demographic headwinds. Results show that without ambitious reforms, Asia's growth will likely moderate, though remaining the world's fastest growing region. The paper offers data-driven tools for policymakers to identify growth drivers and generate robust forecasts.

RECOMMENDED CITATION: Che, Natasha, Federico J. Diez, Anne Oeking, and Weining Xin (2025). "Dissecting Medium-Term Growth Prospects for Asia." IMF Working Paper, International Monetary Fund, Washington, DC.

JEL Classification Numbers:	E17, O11, O40, O47, O53
Keywords:	Asia-Pacific; growth forecasting; demographics; productivity; Dynamic Time Warping; machine learning; growth accounting; economic convergence
Author’s E-Mail Address:	NChe@imf.org, FDiez@imf.org, AOeking@imf.org, WXin@imf.org

WORKING PAPERS

Dissecting Medium-Term Growth Prospects for Asia

Prepared by Natasha Che, Federico J. Diez, Anne Oeking, Weining Xin[1]

[1] The authors would like to thank Davide Furceri and Paulo Medas for their guidance and support, and colleagues at APD's A1 and Regional Studies Divisions for their helpful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. All remaining errors are our own.

Dissecting Medium-Term Growth Prospects for Asia

Natasha Che, Federico J. Díez, Anne Oeking, Weining Xin

August 3, 2025

Abstract

This paper explores Asia-Pacific's medium-term growth prospects using two complementary approaches. First, we conduct a growth accounting analysis and machine-learning-based growth driver estimation, revealing how demographics, capital deepening, productivity, and human capital have shaped Asia's growth trajectories. Second, we introduce an innovative algorithm that forecasts growth by matching country's current conditions with historically analogous country-periods using Dynamic Time Warping (DTW). By comparing these pattern-based forecasts with traditional production function-based projections, we highlight the power of economic convergence and the demographic headwinds facing the region. Our results show that, without ambitious reforms, Asia's growth is likely to moderate, although it would remain the world's fastest growing region. Overall, the paper offers a flexible, data-driven tool set for policymakers to identify key growth drivers, generate robust growth forecasts, and think about structural changes of the economy.

Keywords: economic growth; Asian economies; machine learning; growth decomposition, dynamic time warping

1 Introduction

Economic growth remains one of the most important issues in economic public policy debates. Understanding the dynamics of economic growth is crucial, as it influences employment rates, income levels, and welfare. An accurate and realistic forecast of medium-term growth trends is essential not only for economic policymaking—serving as foundational inputs for fiscal frameworks, debt sustainability analyses, and structural reform agendas—but also for investment planning and resource allocation decisions of the entire economy. Accurate forecasting enables policymakers and businesses to make informed decisions that can mitigate risks, harness opportunities, and foster sustainable development, ultimately shaping future economic environment.

Over the past two decades, growth trajectories across advanced economies and emerging market and developing economies have shown notable differences. Starting with the Global Financial Crisis, many economies experienced significant growth deceleration, particularly pronounced in advanced economies, contributing to discussions about secular stagnation (International Monetary Fund, 2015, International Monetary Fund, 2016). While many emerging market and developing economies demonstrated greater resilience in the post-crisis years, this performance was heterogeneous across regions and economies, with several major emerging economies experiencing their own slowdowns by the mid-2010s. More recently, growth in emerging economies has broadly decelerated (International Monetary Fund, 2024b). The current decade has been marked by significant destabilizing shocks—including the COVID-19 pandemic, supply chain disruptions, the war in Ukraine, and rising geopolitical tensions—that have collectively created a challenging global economic environment. These developments prompt important questions about growth prospects and economic structural changes in the coming years.

Many analysts produce comprehensive forecasts by combining systematic modeling with detailed country-specific expert analysis. This integration of quantitative frameworks with granular country expertise has proven particularly valuable for short-term forecasting, where recent developments and country-specific factors play a decisive role. However, the inherent complexity of economic growth dynamics creates limitations for any forecasting method when extending beyond short-term to medium-term horizons.

Yet, medium-term planning requires careful consideration of potential growth trajectories to inform policy frameworks and economic decisions. Traditional medium-term growth forecasting often relies on structured assumptions about the shape of the production function and its associated parameter values. Such assumptions, however, may not fully reflect a country's development stage or other country-specific fundamentals. Moreover, forecasting accuracy tends to decrease as the time horizon extends, reflecting the fundamental uncertainty in predicting economic outcomes over longer periods during which major structural changes could happen and alter growth trajectories. These challenges are particularly acute for rapidly changing economies such as those in the Asia-Pacific region, where structural transformation has altered growth trajectories in ways that are difficult to capture based on current short-term outlooks. In lower-income economies, the challenge is compounded by data constraints that may restrict the application of more data-intensive forecasting approaches.

Against this background, this paper aims to answer the following questions:

- Historically, what have been the main drivers of economic growth of economies in the Asia-Pacific region?

- Going forward, what are Asia-Pacific's medium-term growth prospects?

We follow two complementary strategies to address these questions. First, we examine growth drivers using both growth accounting and machine learning tree-based approaches. By decomposing historical growth into contributions from traditional production factors (e.g., capital, labor, technology), the growth accounting analysis quantifies their relative importance over time, which also informs growth forecasts in the same framework. The tree-based model deploys a predictive framework by using variables in year t to predict average GDP growth over the subsequent five years (t+1 to t+5). This structure naturally addresses potential endogeneity concerns by ensuring all predictors are predetermined relative to the outcome variables. It complements the growth accounting by considering a broader set of variables beyond traditional production factors—such as macroeconomic conditions, institutional quality, and global factors—thus validating the evolution of growth drivers identified in the growth accounting analysis and highlighting the importance of macroeconomic stability, institutional quality, and global spillovers.

Both backward-looking approaches—the growth decomposition and the tree-based approach—yield insights into the main drivers behind the growth performance in the Asia-Pacific region in recent decades. Demographic factors have become a drag on growth in advanced economies (AEs) since the late 1990s/early 2000s, while increases in human capital and labor force participation have emerged as important growth drivers. In emerging market and middle-income economies (EMMIEs), productivity growth and gains in human capital have been the most significant drivers since the 2000s, likely linked to technology transfer and knowledge spillovers from rising trade openness. Low-income developing countries (LIDCs) have seen capital deepening as the key growth driver since the 2000s—a period that also saw improvements in governance, attracting foreign investments and boosting domestic investments.

Building on these insights, the second part of the paper develops medium-term economic growth forecasts for Asia-Pacific through two complementary methods: a growth accounting framework and a Dynamic Time Warping (DTW)-based pattern-matching method. The growth accounting projections—analogous to the historical decomposition—forecast individual factors (labor, capital, productivity) into the medium term based on the current trends. By contrast, the DTW approach—a pattern-recognition approach originally developed for audio signal processing—identifies and learns from similar historical growth episodes across a wide cross-section of economies, effectively allowing for non-linear and state-dependent growth trajectories. By systematically identifying historically similar growth episodes, the DTW method complements standard forecasting frameworks that rely more heavily on parametric relationships. This novel approach offers a way to capture economic transitions—such as structural transformations or demographic shifts—that may not fit neatly into traditional production function-based or short-term models. By combining these theoretical and pattern-based approaches, we can develop richer forecasts that account both for fundamental economic structures and for real-world historical analogies.

Our findings can be summarized as follows. First, our analysis indicates that, without ambitious reforms, growth is expected to moderate across the Asia-Pacific region. Further, while demographics shifts play an increasingly important role, slowing total factor productivity (TFP) and less capital deepening are also forecast to be significant contributors to the growth slowdown. Methodologically, our approaches yield different quantitative results yet point toward the same broad trends: economic convergence among income groups but an overall moderation in growth over time. It is worth noting that, both the growth accounting approach and the DTW method project growth based on historical data—either the country's own observed ongoing trends or the historical patterns of other economies. Therefore, these methods lack the ability to forecast structural changes, for example, significant technological advances, or ambitious reforms which have not been observed in the past and thus to incorporate these into the projections, despite that these changes could substantially alter the growth trajectories. As a result, our projections could be considered as *baseline* projections if the current trends persist and economies follow historical patterns of their own and others.

This paper makes several contributions to the literature on Asia-Pacific's growth drivers, growth forecasting, and economic development. First, we identify detailed historical drivers behind the strong growth in the Asia-Pacific region over the last three decades, using complementary approaches that combine growth decomposition and tree-based analysis. Second, we introduce a novel DTW-based methodology for growth forecasting, systematically employing pattern-matching algorithms to identify analogous historical trajectories—a perspective that complements parametric models more heavily reliant on current factor dynamics. Third, we provide a comprehensive empirical assessment of this new methodology in an Asia-Pacific context, demonstrating how insights gleaned from historical parallels can refine medium-term growth projections. Finally, our forecasts from these two approaches highlight important takeaways for policymakers tasked with navigating demographic transitions and promoting sustainable long-run development.

Overall, our focus on the Asia-Pacific region provides specific insights into the dynamics of one of the world's most vibrant economic areas. By combining multiple approaches, we shed light on how structural trends—such as demographic shifts, capital deepening, and productivity convergence—are likely to shape future growth. The new DTW-based forecasting method, in particular, offers a transparent and adaptable framework for economists and policy institutions to incorporate country-specific knowledge while maintaining methodological consistency, making it a valuable tool for medium-term policy making and economic planning.

2 Literature Review

The literature on economic growth is vast—spanning from the seminal work by Harrod (1939) and Domar (1946), the canonical neoclassical growth model by Solow (1956) and Swan (1956), to the endogenous growth models by Romer (1990) and Aghion and Howitt (1992). The empirical literature on economic growth encompasses two key strands relevant to our analysis: understanding growth determinants and forecasting future growth trajectories. Motivated by theoretical models, as the ones mentioned above, and following the seminal work of Barro (1991), researchers have identified numerous variables correlated with economic growth, including the accumulation of physical and human capital, research and development, macroeconomic policies, financial development, and international trade. Sala-i-Martin (1997) found that 22 out of 59 variables appear to be "significantly" correlated with growth, and Sala-i-Martin et al. (2004) further refined this analysis using Bayesian approaches.

The growth forecasting literature spans several methodological traditions. Traditional approaches typically combine structural macroeconomic models with expert judgment, as reviewed by Hendry and Clements (2003). These methods rely on identifying key growth determinants and their interrelationships, incorporating both country-specific factors and global conditions. Growth econometrics approaches, comprehensively surveyed by Durlauf et al. (2005), have developed sophisticated methods for analyzing growth patterns across economies. Some of these methods help draw forward-looking inferences, while others, such as Aromí (2019), explicitly focus on medium-term growth forecasts using simple models estimated from historical data.

More recently, machine learning (ML) techniques have been applied to economic analysis and forecasting. Despite increasing applications to studying business cycles and predicting economic crises[1] and nowcasting GDP (e.g., Richardson et al., 2021; Babii et al., 2022), there has been limited work applying ML to medium-term growth forecasting. While these methods excel at identifying complex patterns, they often face challenges in economic applications due to limited data availability and structural breaks, as highlighted by Fernald (2015).

Pattern recognition algorithms offer a promising new direction for economic forecasting. Petris and Petrone (2019) demonstrate the utility of time series matching, though primarily for short-term prediction. Dynamic Time Warping (DTW), originally developed for speech recognition by Sakoe and Chiba (1978), has proven particularly effective at finding optimal alignments between time series that may vary in speed or timing. Berndt and Clifford (1994) extended DTW to time series clustering and classification, showing its ability to capture similarities between series of varying lengths and phases. Bagnall et al. (2017) found DTW compared favorably against newer methods in extensive empirical testing.

The literature on growth convergence and development patterns provides important

[1]For example, see Basu et al., 2019; Fouliard et al., 2021; International Monetary Fund, 2021; Hellwig, 2021; Bluwstein et al., 2023; Hacibedel and Qu, 2023; Cebotari et al., 2024.

context for pattern-based forecasting approaches. Pritchett (1997) highlights the diversity of growth experiences while identifying common development trajectories. Acemoglu and Robinson (2012) and Rodrik (2013) emphasize how institutions and structural transformation shape growth paths, suggesting that economies at similar development stages may follow comparable trajectories.

3 Identifying Asia's Growth Drivers

Over the last three decades, economies in the Asia-Pacific region have outperformed the rest of the world in terms of economic growth. Based on IMF World Economic Outlook estimates, average annual GDP growth in the region since the 1990s has been 2 percentage points higher than in the world as whole, 5.6 and 3.6 percent, respectively. GDP per capita has also increased significantly during this period, with the region's per capita income as a share of the world's increasing by over 20 percentage points—from 40 percent in the early 1990s to almost two-thirds of the current world per capita GDP. The following sections explore the main drivers behind this growth performance based on two complementary approaches: a growth decomposition and a tree-based machine learning approach.

3.1 Growth Decomposition

3.1.1 Methodology

First, we explore the main drivers underlying Asia-Pacific's multi-decade high growth performance through the lens of a semi-endogenous growth model. The analysis is carried out via a growth accounting decomposition.[2] That is, breaking down the observed growth in GDP per capita into the contributions made by different production factors—namely, labor, capital, human capital, total factor productivity (TFP, measuring the efficiency of input usage), and demographic factors affecting the number of workers within a country's population.

The decomposition follows Jones (2022), where real GDP (Y) is given by $Y = K^{\alpha}(AhL)^{(1-\alpha)}$, where K stands for the capital stock, A represents total factor productivity, h is human capital,

[2] Growth accounting provides a straightforward approach to quantify the contributions of different production factors to overall growth, after positing an aggregate production function. Growth accounting goes back to the seminal contribution by Solow (1957), while Aghion and Howitt (2009)) and Jones (2016) provide more recent cross-country examples.

L is employment, and α is the capital share.[3] After some simple algebraic manipulation, this expression can be rewritten in per capita income terms as follows:

$$y \equiv Y/P = (K/Y)^{(\alpha/(1-\alpha))} Ah \frac{L}{P} \tag{1}$$

with P representing the population. Note that introducing labor force LF allows to express GDP per capita as

$$y = (K/Y)^{(\alpha/(1-\alpha))} Ah(L/LF)(LF/P) \tag{2}$$

where L/LF proxies the employment rate and LF/P proxies the inverse of the dependency ratio[4], which can be used to make forward-looking projections (see Section 4.1). Next, we differentiate the expression above with respect to time in order to assess the evolution of growth among Asian economies.

3.1.2 Data

The economies included in the analysis are Australia, Bangladesh, Brunei Darussalam, Cambodia, China, Fiji, Hong Kong SAR, India, Indonesia, Japan, Korea, Lao P.D.R., Malaysia, Maldives, Mongolia, Nepal, New Zealand, Philippines, Singapore, Sri Lanka, Taiwan Province of China, Thailand, and Vietnam. The data come from the Penn World Tables (real GDP in 2017 prices, capital stock at 2017 prices, human capital index, employment, population), International Labor Organisation (average labor income shares), and the United Nations (working-age labor force).

3.1.3 Findings

The data reveal stark differences across country groups, as average economic growth has decreased with income. Indeed, AEs have shown the slowest growth rates, followed by EMMIEs, while LIDCs grew the fastest (Figure 1). The growth drivers also vary substantially across groups:

[3] In the semi-endogenous growth model, technological progress is influenced by both endogenous factors like research and development (R&D) and exogenous factors such as population growth. Unlike purely endogenous models, it assumes that technological change is partly dependent on external influences.

[4] The dependency ratio measures the number of people of non-working age as a ratio to those of working age.

Figure 1: Growth Rates across Income Groups in Asia

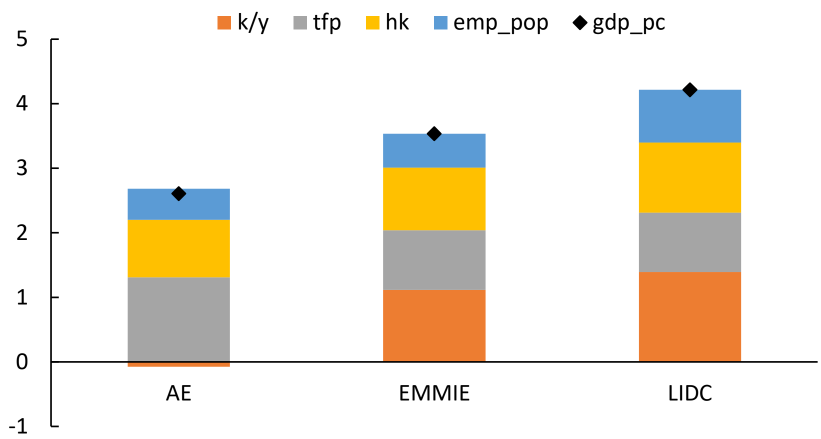

Source: authors' calculations based on data from PWT, ILO, and UN. Notes: The figure shows the average annual (1991-2019) growth rate for each income group, following the decomposition from equation (1) with k/y the capital-output ratio, *tfp* TFP growth, *hk* human capital, *emp_pop* the employment-population ratio and *gdp_pc* GDP per capita growth.

- In AEs, growth has mainly been driven by increases in TFP and human capital and, to a lesser extent, demographic factors. Instead, capital-to-output ratios (i.e., the share of capital in GDP or capital deepening) remained essentially unchanged throughout the period considered and, in fact, marginally detracted from growth.

- Among EMMIEs and LIDCs, in sharp contrast to AEs, capital deepening was the relatively most important factor, accounting for roughly one-third of total growth. Increases in TFP, human capital, and demographic factors (especially in LIDCs) also contributed positively.

Breaking down the growth decomposition by decade allows to assess the dynamic performance of each factor (Figure 2). Once again, notable contrasts are revealed between AEs and EMMIEs/LIDCs. Among the former group, the 1990s showed the strongest performance and growth has declined since then. In contrast, growth accelerated over time for the latter two groups, especially LIDCs.

The decomposition also shows that capital-to-output ratios have accounted for most of the decline in AE growth, perhaps reflecting the increasingly important role of sectors with relatively low capital intensity. Instead, for EMMIEs the data show an uptick in contributions of TFP and a softening of those of capital deepening—whereas capital deepening among LIDCs was the main factor accounting for the acceleration in growth in the last two decades.

Figure 2: Growth Decomposition, by Income Group and Decade

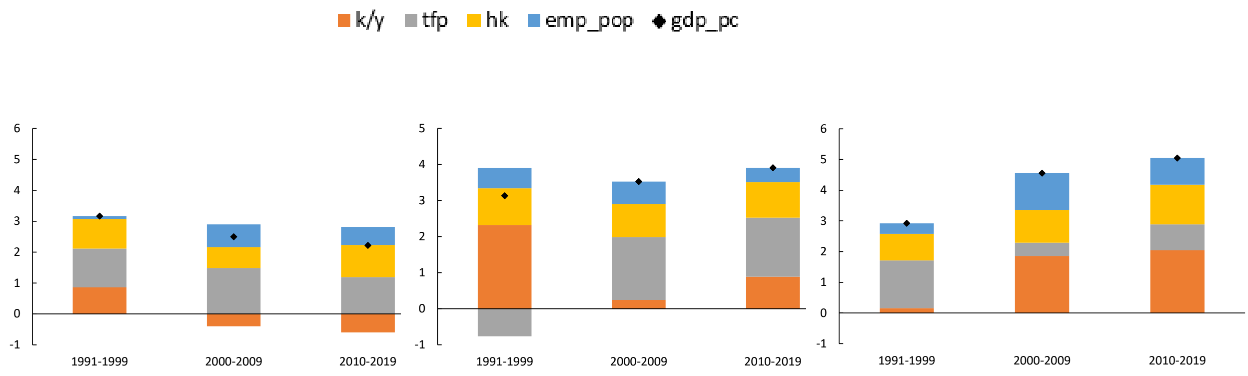

(a) AEs (b) EMMIEs (c) LIDCs

Source: authors' calculations based on data from PWT, ILO, and UN. Notes: Each panel plots the average annual growth rate by country income group and decade. The components follow the decomposition from equation (1) with k/y the capital-output ratio, *tfp* TFP growth, *hk* human capital, *emp_pop* the employment-population ratio and *gdp_pc* GDP per capita growth.

From these results it is also apparent that all AP economies have benefited from demographic tailwinds, as the share of active workers in the total population increased over time. This was done either by (i) increasing the labor force participation rate (i.e., the number of workers within the working-age population), for instance, by increasing female labor force participation; and or (ii) via demographic dividends derived from starting off with a relatively young population, with the number of working-age individuals in the total population increasing over time (i.e., the inverse of the so-called dependency ratio). However, these growth-enhancing factors are unlikely to be sustained in the future—the increase in participation rate is bounded and, more importantly from a quantitative standpoint, by the already-ongoing population ageing.

Figure 3 zooms into the contributions of these demographic factors to GDP per capita growth. While over the last three decades, the overall demographic contribution was similar across country groups, the relative importance of each subcomponent varies widely across groups. For AEs, the average annual 0.5 percentage point increase in the employment-to-population ratio is accounted almost exclusively by higher labor force participation. Instead, for EMMIEs and LIDCs, the increases result from improved inverse dependency ratios, as more younger individuals entered their working-age years. However, in recent years, and especially going forward, these effects are likely to be reversed as AP populations age and a greater mass of retirees become dependent on a relatively smaller mass of working-age individuals, as shown in section 4.1.

Figure 3: Contribution of Employment-Population Ratio to Growth

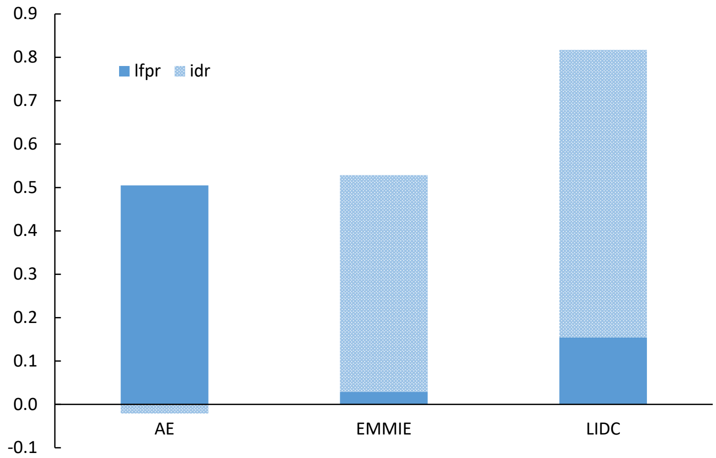

Source: authors' calculations based on data from PWT, ILO, and UN. Notes: The figure shows, for each country group, the average annual (1991-2019) contribution of the employment-population ratio to overall growth, broken down by labor force participation rate (lfpr) and inverse dependency ratio (idr) , following the decomposition from equation (2).

3.2 Tree-Based Approach

Complementing the growth accounting exercise, we now turn to a non-parametric machine learning model to identify growth drivers—beyond traditional production factors—in a predictive framework. Considering a broader set of variables, this approach not only validates the evolution of growth drivers identified above in the growth accounting analysis but also casts light on the importance of macroeconomic stability, institutional quality, and global spillovers.

3.2.1 Methodology

We consider widely-used machine learning tree-based ensemble models for the prediction tasks, namely random forests and XGBoost. Simply put, these tree-based ensemble models consist of multiple regression trees which are built by recursively partitioning the sample into subsamples based on the values of predictors. Such nonparametric model structure lends itself well to prediction task where the target is affected by many predictors in a complex way, including nonlinearities and interactions. To unpack the so-called "black box" of machine learning models and understand the contributions of predictors, we report the results in terms of Shapley values (Strumbelj and Kononenko, 2010 and Lundberg, 2017)—built on the concept of Shapley values from cooperative game theory (Shapley, 1953 and Young, 1985)—which essentially measure the additive marginal contribution of each predictor to the target relative to the sample-average. More technical details on the model training and testing, including hyperparameter tuning, and model explanation, can be found in Appendix A.

3.2.2 Data

The sample consists of 22 Asia-Pacific economies over the period of 1970-2019. Each observation is a country-year pair. For each observation of country c in year t, the target variable is the average real GDP growth in the next five year, i.e., from year $t+1$ to $t+5$, and the predictor values are chosen from year t. As such, the empirical framework is formulated as a prediction framework which addresses the concern that the variables could be affected by real GDP growth if they were included from the same period. We can interpret it as showing how each variable is associated with growth over the medium term. Bearing in mind that machine learning models only tell us about predictive power of the variable rather than causal relationships, many of the results we find are consistent with the empirical and theoretical literature which points to causal interpretation.

We have a total of 40 predictors, which are chosen based on the criteria that they should be associated with interpretable channels described in the literature.[5] We group these predictors into four broader categories: development and demographics, structural conditions, macroeconomics, and global factors and openness (see Appendix B).

3.2.3 Findings

Figure 4 shows the top growth drivers in Asia. First of all, the income level (nominal GDP per capita) has been the most important variable predicting average growth in the next five years which is in line with the principle of economic convergence. Demographics have also played a critical role, as seen in the impact of the working-age population share. Next, trade openness emerges as an important growth engine for Asia: in line with the literature, we find that a higher degree of trade openness and an increase in trade openness over time have been associated with higher growth in the next five years. In addition, higher growth in human capital and an improvement in governance have also been associated with higher growth over the medium term. As for global factors and the external environment, import partners' growth and FDI inflows have been found to be important drivers of a country's growth over the medium

[5]Some of the potential determinants of medium-term growth are excluded because of collinearity with the chosen variables. For example, the level of governance is highly correlated with nominal GDP per capita and thus is not included, while the percentage change of governance is included.

term, adding another layer to Asia's trade-driven growth story.

Figure 4: The Top Drivers of Growth in Asia

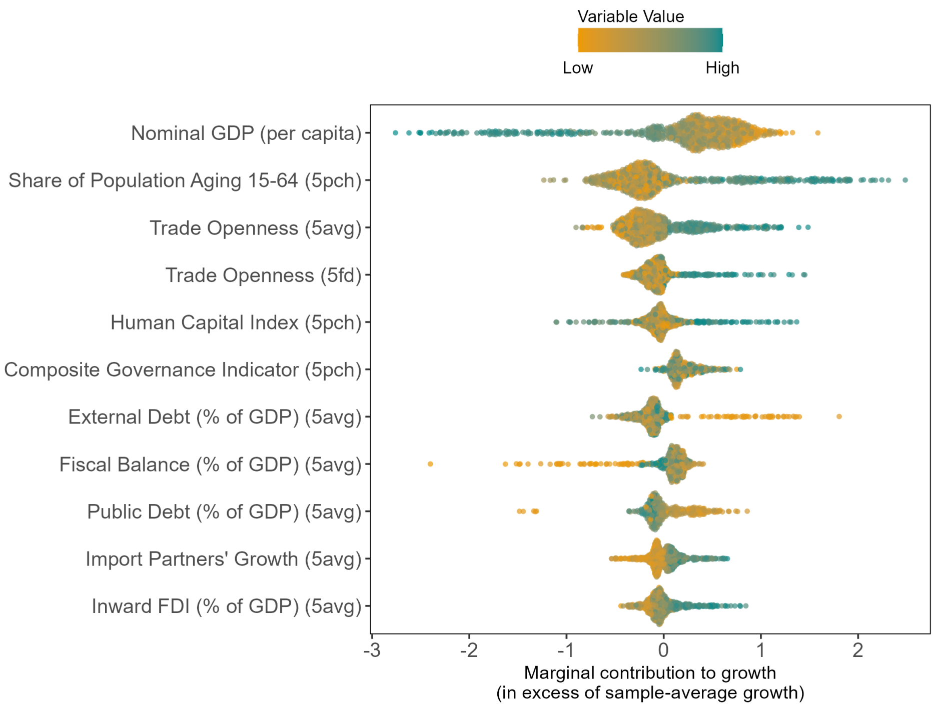

Notes: The figure shows the contribution to five-year average growth in Asia, in excess of the sample average (as measured by the Shapley value) for those most important drivers—ranked by individual average contribution—that together account for 70 percent of the total contributions to five-year average growth for the sample. Each dot represents one observation (one country-year pair), and its color indicates the variable value, with orange being low and green high. Therefore, the value of the contribution (and its negative or positive sign) and color together illustrate how the variable value affects the five-year average growth. For example, for nominal GDP per capita, the top driver with the largest average contribution, most green dots (observations with higher nominal GDP per capita) are to the left (i.e., reduce five-year average growth, relative to the sample average) and orange dots (observations with lower nominal GDP per capita) to the right (i.e., increase five-year average growth, relative to the sample average). Variable transformations are indicated in the parenthesis, with 5avg, 5pch, and 5fd indicating five-year average, five-year average of annual percentage change, and five-year average of annual level change.
Sources: authors' calculations.

To better understand how these variables have been associated with growth over the medium term, panel figure 5 plots the marginal contribution to medium-term growth for each variable versus the variable value. The aforementioned variables are associated with growth in a more or less linear way, with economies with lower nominal GDP per capita, a higher growth of the working-age population share, a higher degree of trade openness and an increase in trade openness, higher growth in human capital, an improvement in governance, higher import partners' growth, and more FDI inflows would see higher growth over the medium term.

However, as for domestic macroeconomic variables (Figure 6), their association with

Figure 5: Selected Top Growth Drivers in Asia

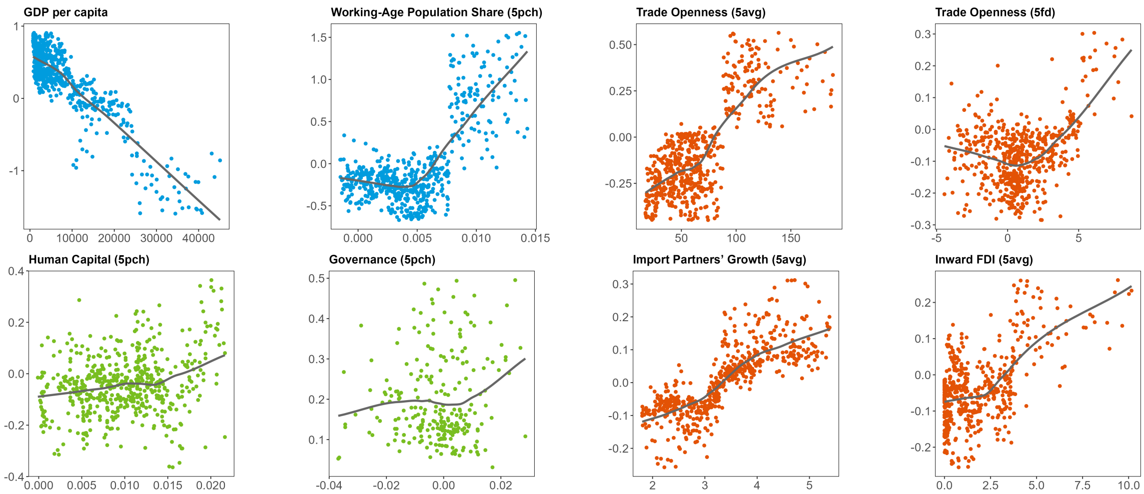

Notes: In each figure, the y-axis shows the variable's marginal contribution to the five-year average growth (in excess of the sample average) as measured by the Shapley value, and the x-axis shows the value of the variable. To better illustrate the trend, outliers (observations with variable value or Shapley value below the 5th percentile or above the 95th percentile of the sample) are not shown, and a local regression with degree up to 1 is used to fit a smooth curve based on the data, as shown by the blue line.
Sources: authors' calculations.

growth could be nonlinear. While economies with higher external debt (i.e., more than 50 percent of GDP) are found to have (slightly) higher growth in the medium term—possibly reflecting the growth-enhancing effects of external borrowing—there are some examples which have seen high growth with much lower external debt. Similarly, when public debt is lower than about 60 percent of GDP, lower public debt is associated with higher growth over the medium term, indicating that strong public finance positions would be important for economic growth. But this relationship is less strong when public debt is above 60 percent, possibly reflecting economies' different debt carrying capacity. We also observe cases with very low public debt contributing both positively and negatively to growth. Relatedly, a weak fiscal balance, for example with a fiscal deficit larger than 5 percent of GDP, is associated with low growth over the medium term, as fiscal sustainability concerns weigh on growth. On the other hand, some weak evidence suggests that larger fiscal balances could be associated with lower growth in the future, likely reflecting the negative growth effects of fiscal consolidation.

How have these growth drivers evolved over time? Figure 7 plots the contribution of the four categories of drivers to each income group's average growth in a five-year window during 1990-2019, in excess of the average growth of the entire Asia-Pacific sample during 1970-2019. First of all, advanced Asia has seen lower growth than the regional average since the 1990s, with

Figure 6: Macroeconomic Growth Drivers in Asia

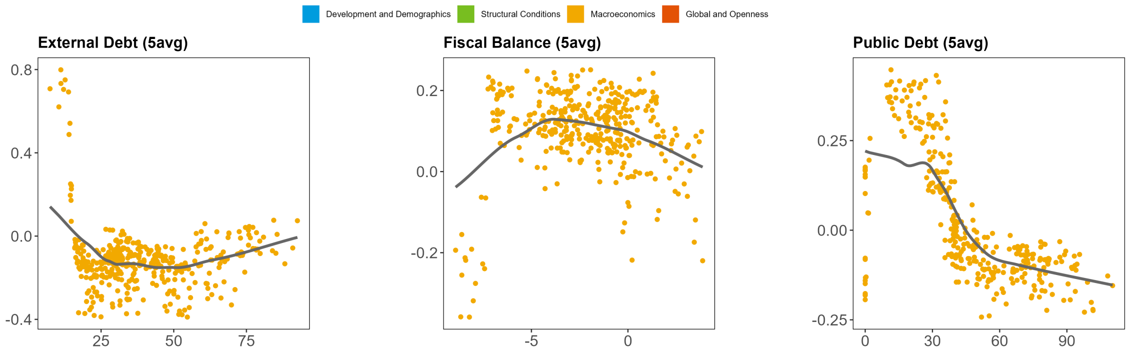

Notes: In each figure, the y-axis shows the variable's marginal contribution to the five-year average growth (in excess of the sample average) as measured by the Shapley value, and the x-axis shows the value of the variable. To better illustrate the trend, outliers (observations with variable value or Shapley value below the 5th percentile or above the 95th percentile of the sample) are not shown, and a local regression with degree up to 1 is used to fit a smooth curve based on the data, as shown by the blue line.

an average growth of 3.3 percent, while low-income economies in Asia have seen the highest growth among the three income groups, with an average growth of 5.5 percent, and emerging Asia in between with an average growth of 4.2 percent.

For Asian AEs, development level and demographics—namely, higher GDP per capita and lower growth rate of working-age population share—have been the main growth drag since the late 1990s. The negative contribution from these factors has also been increasing over time, lowering growth of advanced Asia by as much as 2 percentage points below the regional average in the last decade. Macroeconomic conditions have contributed positively to advanced Asia's growth historically, especially in the early 2000s. This was a period with much stronger fiscal positions in AEs, with more balanced fiscal accounts and declining public debt. However, since the Global Financial Crisis, macroeconomic conditions no longer contributed positively to medium-term growth in advanced Asia (except for the recovery from the Global Financial Crisis). Moreover, global factors have become a growth drag in the last decade, among which importers' growth had the largest negative contributions in the last five years, likely associated with China's growth slowdown. At the same time, structural conditions have continued to improve in advanced Asia in the last decade, including on governance, human capital, and labor force participation, which have contributed positively to medium-term growth.

For Asia's EMMIEs, their gaps from the development frontier and growing working-age populations contributed positively to medium-term growth but the positive contributions have been diminishing with GDP per capita increasing but population aging at the same time. Since early 2010s, these factors started to become a growth drag, lowering emerging economies' growth from the regional average. Since the 2000s, structural conditions have improved significantly, such as enhanced human capital and reduced informality, which have contributed positively to emerging Asia's medium-term growth during the last two decades. While increased trade openness helped promote growth in many of the emerging economies in Asia, lower global growth after the Global Financial Crisis, including in China, has contributed less to the medium-term growth in emerging Asia in the last decade. On the macroeconomic conditions front, higher public debt since the Global Financial Crisis has been one main factor contributing negatively to emerging Asia's growth.

Figure 7: Medium-Term Growth Drivers in Asia, by Income, 1990-2019

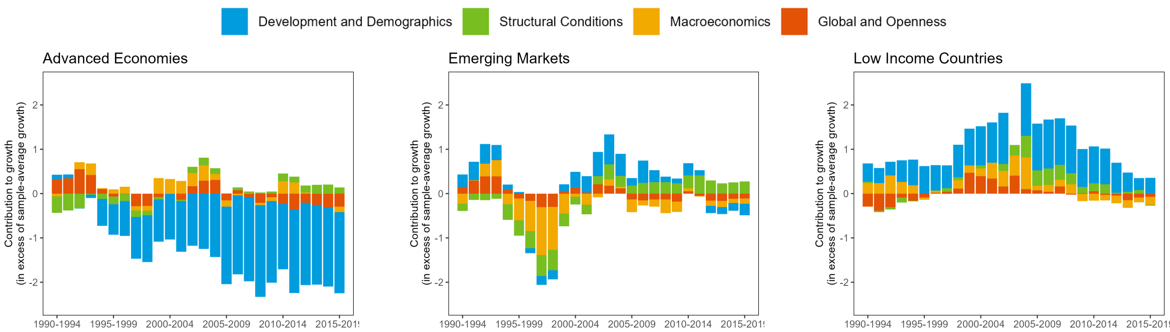

Notes: Each figure shows the contribution of various categories of drivers to the income group's average growth in a five-year window, in excess of the average growth of the entire Asia sample during 1970-2019. Drivers are grouped into four broader categories including development and demographics, structural conditions, macroeconomics, and global factors and openness.
Sources: authors' calculations.

For LIDCs in Asia, development level and demographics have been the most important growth driver (relative to the regional average), amid lower level of development levels and faster growing populations and working-age population shares. It is worth noting that, similar to emerging Asia, these contributions have decreased in the last decade though remaining positive, as GDP per capita increased and the share of working-age population grew slower. Structural improvements have also contributed positively to medium-term growth, especially in the 2000s, a transformative period for governance in low-income Asian economies during which many of them have accelerated economic reforms aimed at liberalizing the economies, improving regulatory frameworks, and attracting foreign investments. As a result, governance in Asian LIDCs improved significantly, with the average governance index (encompassing control of corruption, government effectiveness, regulatory quality, rule of law, and voice and accountability) increasing by about 50 percent in the second half of 2000s. Macroeconomic conditions contributed positively to low-income Asian economies' growth before the last decade, but have become a grow drag since 2010s, when many low-income economies faced fiscal challenges with large fiscal deficits.

3.3 Comparing the Two Approaches

Both approaches—the growth decomposition and the tree-based approach—yield insights into the main drivers behind the growth performance in the Asia-Pacific region in recent decades. Demographic factors became a drag on growth in AEs since late-1990s/early-2000s, while increases in human capital and labor force participation became important growth drivers since then. In EMMIEs, productivity growth and an increase in human capital have been the most important drivers for growth since 2000s, with the former likely linked to technology transfer and knowledge spillovers from increasing trade openness. LIDCs have seen capital deepening as the most important growth driver since 2000s, during which period governance improved significantly which likely helped boost domestic investments and attract foreign investments, accelerating capital deepening.

The two approaches (see Figures 2 and 7) also show how trends have evolved over time. Factors such as demographics and structural transformation (International Monetary Fund, 2024a) have played an important role and are likely to do so going forward. Any forward looking approach will thus have to take these developments into account.

Having identified the key growth drivers, we now move to forecast the evolution of growth through two complementary frameworks: a production-function-based approach and a novel pattern-matching approach (DTW).

4 Forecasting Asia's Medium-Term Growth

We employ two complementary approaches to forecast medium-term growth in Asia: a traditional growth accounting framework and a novel pattern-matching methodology based on Dynamic Time Warping (DTW). The growth accounting approach provides theoretically-grounded projections based on fundamental factors, while the pattern matching approach leverages historical growth experiences to identify likely development paths.

4.1 Growth Accounting Forecasts

4.1.1 Methodology

Our growth accounting approach builds on the semi-endogenous growth framework presented in section 3.1. We project individual factors labor, human capital, physical capital, and TFP, on a country-by-country basis. Aggregate growth is derived by summarizing over these projected contributing factors.

The individual factors are assumed to follow their pre-pandemic trends. This simplistic assumption primarily stems from data considerations, as the availability of data poses a significant constraint on more extensive models with a diverse range of economies. This implies at least four important caveats beyond the significant uncertainty inherent in any such forecast. First, the results do not explicitly take into account convergence, i.e., we do not specify a specific relationship between initial conditions and long-run outcomes. Second, the projections assume the 'status quo' to broadly continue, abstracting from assuming major technological disruptions or structural transformations. Third, as argued in Kaldor (1961) or Jones (2022), the capital-output ratio is constant in the long-run. Here, however, we assume changes in the capital-to-output ratio, in line with the idea that transition dynamics to a long-run equilibrium are very slow and can take decades. Fourth, linear trends do not necessarily capture slowly progressing structural transformation.

The decomposition of real GDP growth follows the same expression as above with $Y = K^{\alpha}(AhL)^{(1-\alpha)}$. This expression can be written in growth rates by log-linearization and first differences, with $\hat{x}$ denoting these log differences:

$$\hat{Y} = \alpha\hat{K} + (1-\alpha)(\hat{A} + \hat{h} + \hat{L}) \tag{3}$$

We forecast $\hat{K}, \hat{A}, \hat{h}, \hat{L}$ mostly via linear regressions. The sample includes 20 regional economies: Australia, Bangladesh, Cambodia, China, Fiji, Hong Kong SAR, India, Indonesia, Japan, Korea, Lao P.D.R., Malaysia, Mongolia, Nepal, New Zealand, Philippines, Singapore, Sri Lanka, Thailand, and Vietnam.[6]

[6] Due to our simplifying assumption of a continuation of linear trends, several economies drop out of our sample as they exhibited continuous negative TFP growth rates in past years which would imply negative TFP levels at one point in the future. These economies include Brunei Darussalam, Macao SAR, Maldives, and Myanmar.

4.1.2 Data and Assumptions

We use the following assumptions:

- Capital stock K is projected using the perpetual inventory method. Past capital stock data, depreciation rates, and implicit investment data are taken from the Penn World Table (PWT) version 10.01. Linear regressions are based on the period 2009-19. This implies increasing depreciation rates for most economies, in line with trends towards more intangible capital which usually requires more input to sustain.

- TFP growth is projected based on country-specific linear regressions of TFP levels during the post-GFC period 2009-19. Historical levels are derived as the residual from the growth accounting decomposition in the semi-endogenous growth framework as described in section 3.1 above.

- For employment L, we focus on age cohort-gender groups $L_{i,j} = P_{i,j} \times LFPR_{i,j}$, with $P_{i,j}$ the projected population size of an age cohort (i) and gender (j), as projected by the UN's World Population Prospects medium-fertility scenario. For relevant age groups, we consider 15-24, 25-54, 55-64, and 65+. Future labor force participation rates by age and gender, $LFPR_{i,j}$, are forecast by country-specific linear regressions based on the time period 2000-22 for each cohort-gender group. Historical LFPRs are sourced from the ILO. The projected linear trend is bound by regional maxima or minima, respectively. In case any projected cohort-gender LFPR hits the regional maxima or minima during the projection period, a logistic growth regression with this regional maximum or minimum is used instead of the linear regression. Multiplying these LFPRs with the demographic projections and summing up total labor over all groups yields our employment forecast: $L = \sum_i \sum_j (P_{i,j} \times LFPR_{i,j})$. This implies we assume economies operate at full employment, arriving at a natural rate of unemployment that includes frictional and structural unemployment. Assuming constant natural rates of unemployment, we can abstract from including them in our growth rates calculations.

- We assume human capital continues its linear trend. Human capital in PWT is measured by the average years of schooling and an assumed rate of return to education, estimated via Mincer equations. For our projections, we assume human capital follows its linear trend over 2000-19, improving across all economies, irrespective of income group. This is in line with Jones and Romer (2010) who show that human capital per worker has been rising dramatically throughout the world.

- Finally, the share of labor income in GDP, $1-\alpha$, based on country-specific data from ILO, is assumed to remain constant, hence the same assumption holds for the capital income share α.

4.1.3 Findings

The projections show that potential growth across the region is assumed to decline over the medium- and long-term, in the simple aggregate to around 3.1 percent by 2030 and 2.2 percent by 2040, compared to 3.7 percent pre-pandemic growth. Emerging market and developing economies (EMDEs) will be the key driver behind this decline, with growth projected to fall to 3.8 percent in 2030 and 2.8 percent in 2040, compared to pre-pandemic growth of 4.8 percent. But a slight downward trend is also observed for AEs.

Figure 8: Projected Average Growth in the Asia-Pacific Region

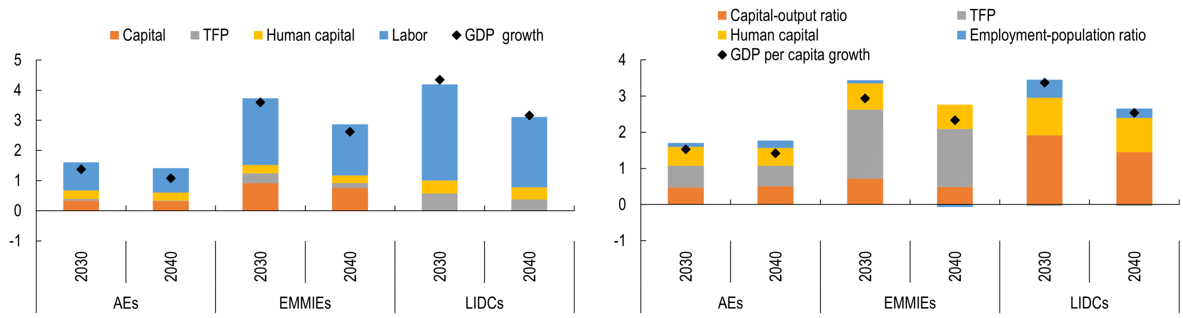

(a) Aggregate growth (b) Per capita growth

Source: Penn World Table version 10.01, International Labour Organization, United Nations; and authors' calculations.
Note: In percent. Aggregate groups based on medians across economies. Median contributions do not always add up to median GDP growth.

For both AEs and EMDEs, a significant factor behind the long-term slowdown will be a lower contribution from labor, i.e., the effect of rapidly ageing societies and the reversal of Asia's demographic dividend (Figure 8). Some of the demographic effects will be dampened by positive net migration and increases in labor force participation rates, especially among females and the elderly. Yet, comparing old populations (i.e., those with a shrinking working-age population) to currently still young populations, lower labor inputs will continuously subtract from growth in the former group throughout the projection period, but labor inputs in the latter group will fall more rapidly, even while still positively contributing overall (Figure 9).

Figure 9: Projected Change in Labor Supply

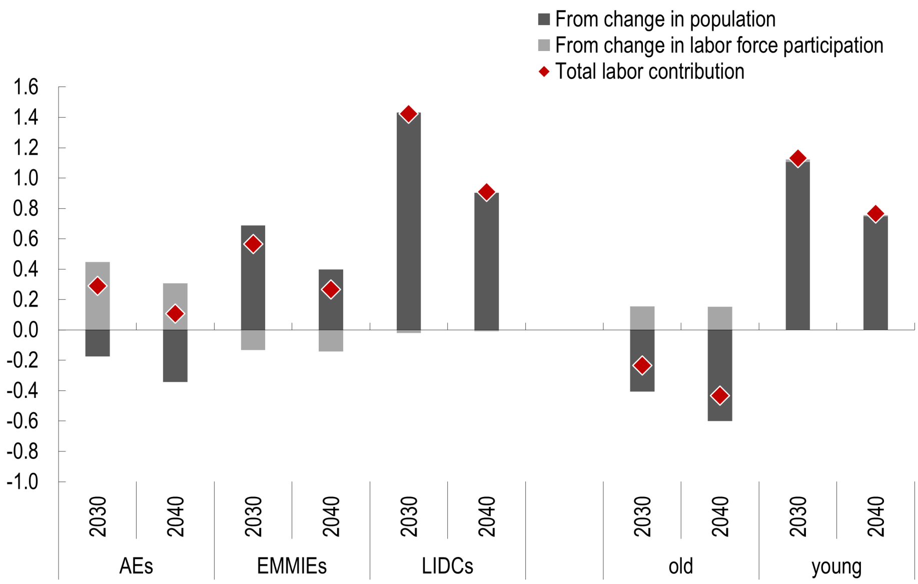

Source: United Nations World Population Prospects 2022; ILO; PWT; WEO; and authors' calculations. Note: In percentage points. Old populations defined as those with shrinking working age population (15-64) over the time horizon 2024 to 2040. Projections of labor supply based on projections of working age population and cohort-gender specific labor force participation rates.

Other factors add to the slowdown. A declining contribution from capital is projected to lower potential growth, both for AEs, and more so for EMMIEs. Finally, the contributions from TFP will slow over time, while the positive contributions from human capital improvements will continue.

While labor contributes significantly to aggregate growth, improvements in per capita metrics—such as human capital and employment ratios—can lead to sustainable long-term growth. In per capita terms, slowing capital-output ratios—especially for LIDCs—and falling TFP contributions—in particular for EMMIEs—explain the downward growth trend.

Providing this starting point by shedding light on some possible drivers of regional growth in the medium- to long-term, we now move to a novel approach for forecasting economic growth.

4.2 Dynamic Time Warping Forecasts

4.2.1 The DTW Algorithm

The Dynamic Time Warping (DTW) algorithm measures similarity between time series that may vary in speed or timing. For two time series $X = \{x_1, x_2, ..., x_n\}$ and $Y = \{y_1, y_2, ..., y_m\}$, the DTW distance is calculated through the following steps:

First, we initialize the cost matrix D of size $n \times m$, where n and m are the lengths of the respective time series (in our application, $n = m = 5$ years). We set $D[0][0]$ to zero and the first row and column to infinity to handle edge cases in the recursion.

Next, we calculate the pair-wise distance $D[i][j]$ between each element x_i from sequence X and y_j from sequence Y. For multivariate data, distances are computed vector-wise using the Euclidean distance:

$$D[i][j] = \sqrt{\sum_{p=1}^{q} (x_{ip} - y_{jp})^2} \tag{4}$$

where q is the number of features in the multivariate time series. Variables are normalized to minimize the impact of scaling.

Finally, we construct the cumulative cost matrix C recursively. The first elements are initialized as:

$$C[1][1] = D[1][1] \tag{5}$$

$$C[i][1] = C[i-1][1] + D[i][1], \text{ for } i = 2, \ldots, n \tag{6}$$

$$C[1][j] = C[1][j-1] + D[1][j], \text{ for } j = 2, \ldots, m \tag{7}$$

For the remaining cells, we update each $C[i][j]$ based on the pair-wise distance and the minimum of the cumulative distances from the three adjacent cells:

$$C[i][j] = D[i][j] + \min(C[i-1][j], C[i][j-1], C[i-1][j-1]) \tag{8}$$

The final DTW distance is given by $C[n][m]$, representing the minimum cumulative distance needed to optimally align the sequences.

In simpler terms, DTW 'warps' the time dimension to align the shape of two time series, such as the economies' growth paths. If the growth spurt of country A was slightly earlier than that of country B, DTW can still consider them 'similar episodes' by stretching one time series to match the peaks and troughs of the other. This is particularly helpful when analyzing economies whose transitions occur at different speeds but follow a shared trajectory.

4.2.2 Pattern Matching Data

The Dynamic Time Warping approach focuses on six fundamental indicators from Penn World Table 10.01 that capture distinct aspects of economic development. These include GDP per capita as a measure of overall development, capital-output ratio reflecting productive capacity, consumption-output ratio indicating demand structure, labor productivity capturing technological progress, real GDP growth measuring economic momentum, and population growth representing demographic dynamics.[7] This harmonized data set spans 183 economies from 1950 to 2019, enabling identification of similar growth episodes across time and economies. These six indicators—GDP per capita, capital-output ratio, consumption-output ratio, labor productivity, real GDP growth, and population growth—were chosen because they collectively capture key aspects of the structural and demographic trajectory of an economy, which is crucial for the matching of medium-term similarity.[8]

4.2.3 The PatternSync Forecasting Procedure

Using the DTW algorithm as our similarity measure, we implement a systematic forecasting procedure which we call PatternSync. The procedure begins by defining a target period $[t-k+1, t]$ for the country of interest, where k is the window size for comparison. In our implementation, we use $k = 5$ years with t as the most recent year available in the data.

With the target period defined, we compute DTW distances between the target country's base period and all possible historical periods of other economies. This computation considers multiple economic variables simultaneously, including measures of economic develop-

[7] Robustness tests dropping one variable at a time show that while all six indicators contribute to forecast accuracy, the method remains functional with five variables, though with degraded performance as measured by higher prediction errors in backtesting. This suggests that the multi-variate approach is robust to minor specification changes but benefits from the full set of macro features for better performance.

[8] GDP per capital and labor productivity are measured using chained PPPs in 2017 USD. Capital-output and consumption-output ratios use current price PPPs.

ment, productive capacity, and demographic dynamics. The DTW algorithm's ability to handle multivariate time series allows us to capture similarity across multiple dimensions of economic development.

After computing distances, we identify the most relevant historical patterns by selecting the top x percent of country-periods with the shortest DTW distances to our target period. In our baseline implementation, we set $x = 2$, effectively selecting the most similar two percents of all available country-periods. For each selected similar country-period, we examine its subsequent growth trajectory over the next y years, where y is our desired forecast horizon, set to 10 years in our backtesting exercise for medium-term projections.

The final step involves aggregating these historical growth trajectories to generate forecasts. We employ multiple aggregation methods to provide a comprehensive view of potential growth paths. The primary methods include taking the median of subsequent growth rates across all similar periods, computing a simple average, and calculating a weighted average where weights are proportional to the inverse of DTW distances. This last approach gives greater importance to more similar country-periods. We also compute various percentile values of the growth distribution, which can help identify potential upside and downside scenarios based on historical patterns.

4.2.4 Backtesting Results

The backtesting procedure was conducted on 26 Asia-Pacific economies[9] to evaluate the forecasting performance of the DTW PatternSync method. For each economy with available data, the algorithm was run using rolling 5-year base periods, beginning with 1995-1999 and ending with 2015-2019, with a forecast horizon of 10 years for each base period. The target economies included a diverse mix of advanced economies, emerging markets, and low-income developing economies.

While forecasts were generated specifically for these 26 Asia-Pacific economies, the algorithm utilized data from all 183 economies in the database to identify similar country-periods, maximizing the potential for finding relevant historical patterns. To assess forecast accuracy, the algorithm's predictions were compared against trend growth rates calculated using HP-filtered

[9] The economies included in the backtest were: Australia, Bangladesh, Brunei Darussalam, Bhutan, China, Fiji, Hong Kong SAR, Indonesia, India, Japan, Cambodia, Korea, Lao P.D.R., Sri Lanka, Macao SAR, Maldives, Myanmar, Mongolia, Malaysia, Nepal, New Zealand, Philippines, Singapore, Thailand, Taiwan Province of China, and Vietnam.

historical real GDP data for each target country during the backtesting timeframe. The evaluation metrics included average error (AE), root mean squared error (RMSE), and hit rate (HR)[10], calculated for each year of the forecast horizon. In identifying the most similar country-periods, we chose those episodes that are within the top 2% shortest DTW distance to a target country's base period.

Figure 10: Median Absolute Error And Hit Rate Over Time

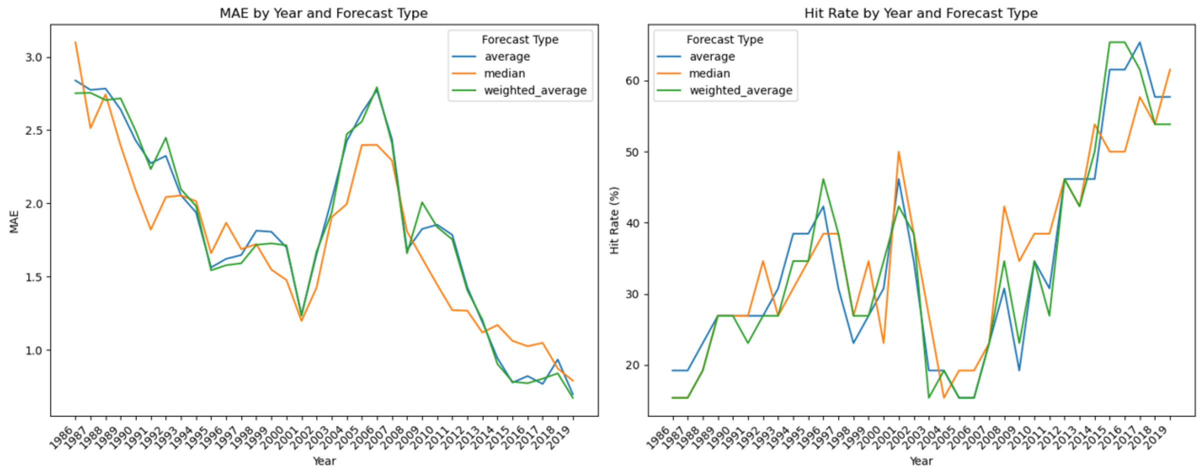

Source: Authors' calculations.

The backtesting results demonstrate several notable patterns in the DTW PatternSync forecaster's performance. First, as shown in Figure 10, which plots MAE and hit rate across all forecast horizons over time, the median absolute error (MAE) exhibits a clear downward trend and the hit rate a clear upward trend as time goes on though with a notable reversal during 2005-10 that likely reflects the challenges of forecasting during the Global Financial Crisis period. By the end of the backtesting sample, the MAE stabilizes at around 0.9 percentage points for 5-year-ahead forecasts. This improvement in forecast accuracy can likely be attributed to the growing pool of historical country-periods available for comparison, which provides the algorithm with a richer set of patterns to learn from. This pattern also suggests that the forecast should continue to become more reliable as more years of data accumulate, barring any significant changes in the fundamental structure of the global economy.

Second, examining the error patterns across different forecast horizons reveals an interesting feature of the DTW PatternSync method (Figure 11). Unlike traditional forecasting

[10]Hit rate is defined as the percentage of time the forecasted value falls within a 25% band of the actual value.

approaches that typically show deteriorating performance as the forecast horizon extends, our method maintains relatively stable accuracy between short-term (Years 1-3) and medium-term (Years 4-7) horizons. This stability suggests that the pattern-matching approach is particularly valuable for medium-term forecasting, where conventional methods often struggle.

Figure 11: Median Absolute Error by Forecast Horizon

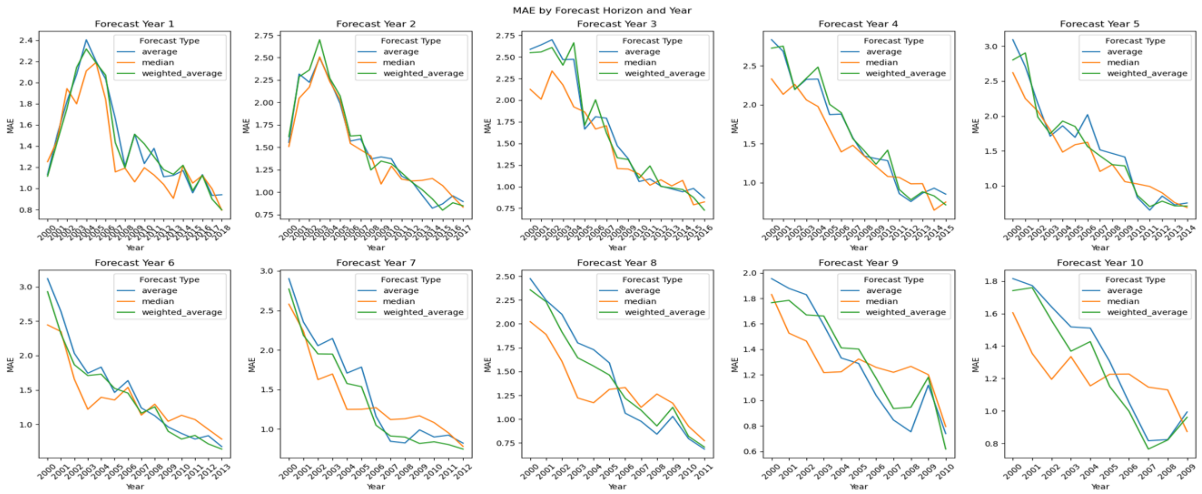

Source: Authors' calculations.

The hit rate analysis provides further demonstration of the method's performance. Figure 12 shows that the hit rate improves consistently over time, reaching approximately 60% by the end of the sample period for 5-year-ahead forecasts. This means that in recent periods, the forecaster successfully predicts trend growth within a 25% band of the actual outcome more than half the time.

Regarding different aggregation methods, we find that weighted average forecasts, where weights are inversely proportional to DTW distances, generally perform slightly better than simple averages or medians, particularly for longer forecast horizons. This suggests that the similarity scores captured by DTW distances contain valuable information for prediction accuracy. However, the performance difference between aggregation methods is modest, indicating that the method's success relies more on effectively identifying similar country-periods than on the specific approach to aggregating their subsequent growth paths.

We compare the DTW PatternSync forecasts with two widely-respected conventional forecasts: the IMF World Economic Outlook (WEO) and the Economist Intelligence Unit (EIU) forecasts, in Table 1. The comparison covers 2011-2019 for Asia-Pacific economies, evaluating

Figure 12: Hit Rate by Forecast Horizon

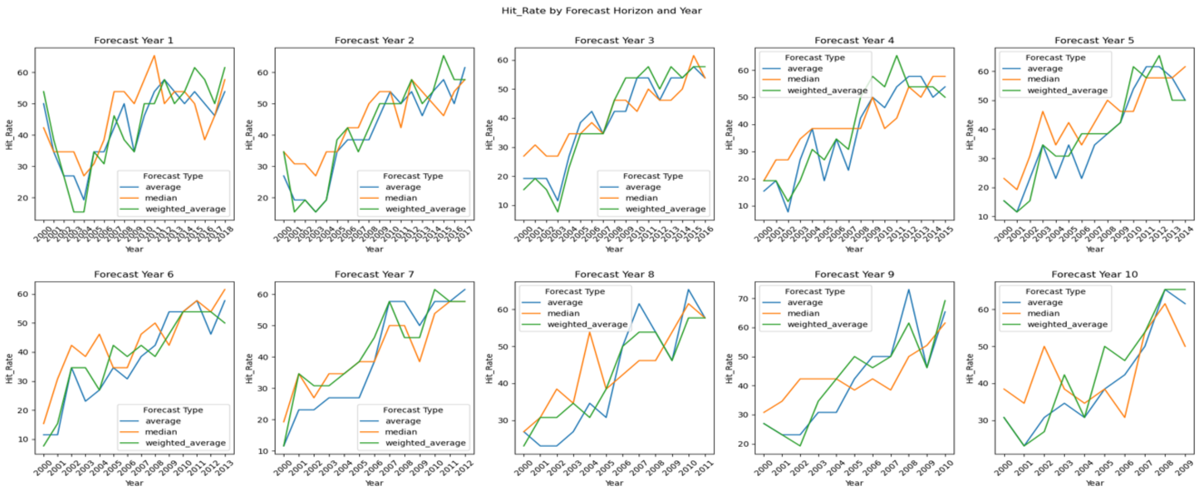

Source: Authors' calculations.

performance across multiple metrics including median absolute error, root mean squared error (RMSE), average error, and hit rate.[11]

For short-term forecasts (Years 1-2), conventional forecasts demonstrate strong performance, particularly in hit rates. The WEO and EIU achieve hit rates of 68-79% in Year 1, compared to 65% for DTW PatternSync, suggesting their incorporation of current economic conditions and expert judgment adds particular value for near-term predictions. The EIU notably achieves the lowest RMSE in Year 1 at 0.80, compared to 1.41 for WEO and 1.49 for DTW.

However, as the forecast horizon extends into the medium term (Years 4-7), the relative performance shifts. While all three methods see some deterioration in accuracy, the decay is notably smaller for DTW PatternSync. Its median absolute error remains relatively stable around 0.7-0.8 percentage points, while both WEO and EIU see larger increases in forecast errors. By Year 5, DTW PatternSync achieves comparable or better performance across most metrics. Its hit rate remains steady at around 62-65% through Year 9, while the conventional forecasts' hit rates decline to the 40-60% range.

Another distinguishing feature is the directional bias in forecasts. The average errors

[11] Errors and hit rates reported are both comparing the forecasted growth rate with the trend growth rate from the HP-filtered realized real GDP. Comparing the forecasted growth with the unfiltered actual growth outturn yields qualitatively similar results, though the magnitudes of errors in the latter case are larger across all forecast methods. For the DTW method, the table used the median forecasts from the algorithm in the calculations.

Table 1: Forecast Error and Hit Rate Comparison for Asia-Pacific economies (2011-2019)

		Forecast horizon (in year)								
		1	2	3	4	5	6	7	8	9
Median absolute error (in pp)	WEO	0.64	0.86	0.97	1.03	1.12				
	EIU	0.56	0.58	0.68	0.76	0.76	0.69	0.63	0.62	0.57
	DTW	0.82	0.81	0.81	0.78	0.79	0.69	0.74	0.78	0.72
Root mean square error (in pp)	WEO	1.41	1.75	2.03	2.06	1.88				
	EIU	0.80	0.91	1.22	1.38	1.39	1.15	1.20	1.21	1.08
	DTW	1.49	1.37	1.29	1.27	1.25	1.17	1.10	1.03	1.00
Average error (in pp)	WEO	0.30	0.44	0.58	0.65	0.69				
	EIU	0.11	0.27	0.40	0.63	0.51	0.35	0.23	0.18	0.17
	DTW	-0.33	-0.32	-0.32	-0.30	-0.31	-0.33	-0.31	-0.33	-0.36
Hit rate	WEO	68%	59%	55%	50%	41%				
	EIU	79%	71%	68%	59%	63%	61%	63%	59%	56%
	DTW	65%	64%	64%	63%	62%	65%	65%	65%	65%

Source: Authors' calculations.

reveal that both WEO and EIU tend to have an optimistic bias, with positive average errors that grow with the forecast horizon. In contrast, DTW PatternSync shows a mild negative bias that remains relatively stable, suggesting it may provide a useful complement to conventional forecasts by offering a different perspective on medium-term growth prospects.

A natural question arises about whether the DTW approach can effectively identify relevant historical patterns for AEs that are near the technological frontier. While AEs have fewer direct development analogues, the method remains applicable for several reasons. First, even frontier economies experience patterns in demographics, productivity growth, and structural transformation that can find historical parallel. Second, our six-dimensional matching approach can identify similarity across multiple economic phases rather than just development levels– an AE may experience demographic patterns simliar to a historical non-AE and the growth trajectory of the latter may shed some light on that of the former, even if the growth level does not match. Third, AEs can match with their own historical periods or with other AEs during comparable structural transitions. Our backtesting results suggest the method maintains reasonable accuracy for AEs, though the smaller pool of relevant historical analogues does represent a limitation compared to emerging economies with more abundant catching-up episodes in the historical record, thus the variance of projections for AEs is likely higher.

4.2.5 Forward-Looking Forecasts

Using the DTW PatternSync method, we generate long-term growth forecasts for Asia-Pacific economies through 2040. The forecasts use 2015-2019 as the base period for pattern matching, leveraging the full historical dataset to identify similar country-periods and their subsequent growth trajectories. For each country, we calculate both the median forecast and a percentile-based forecast, where the percentile is selected based on the country's recent growth performance relative to its historical similar peers.

Several notable patterns emerge from the median forecasts summarized in Figure 14. First, there is a clear divergence between AEs, EMMIEs and LIDCs. The AE group is projected to grow at around 1.8-2.0% on average over 2030-2040, while EMMIEs maintain higher growth rates of 3.2-3.6% and LIDCs of 4.5-5.3%. However, the growth differential between the groups narrows over time, consistent with the economic convergence theory. Among EMMIEs and LIDCs, several fast-growing economies including Cambodia, Vietnam, and Bangladesh are projected to maintain robust growth rates above 4% through 2040, though with some moderation over time. China's growth is expected to moderate more significantly, from around 4.2% in 2030 to 3.3% by 2040 (based on 75th percentile forecasts reflecting its historical outperformance). This is broadly in line with previous work (e.g., Muir et al., 2024). India's growth shows more stability, maintaining around 3.3% through the forecast horizon. Appendix C demonstrates how the DTW method works in practice and its value for country-specific analysis beyond simple point forecasts by examining Cambodia as a detailed case study.

The percentile forecasts warrant careful interpretation. A high percentile (e.g., 75th for China, Singapore, and Malaysia) indicates that these economies have recently outperformed their historical peers. The percentile forecast essentially projects what their growth would be if they were to maintain this same level of relative outperformance. However, historical evidence suggests that economies rarely maintain very high or very low percentile rankings among peers over multiple decades. Specifically, the cross-country growth performance ranking exhibits a certain level of mean reversion, i.e., economies that are significant outperformers in a given period tend to see their ranking drop over time, and vice versa. To illustrate, Figure 13 presents a partial regression plot regressing economies' future growth ranking in 10 years against their current growth ranking.[12] The regression coefficient is negative and significant when taking

[12]The sample includes non-overlapping 10 year data on 183 economies from 1959 to 2019. The results

into account country fixed effect, indicating that economies are statistically unlikely to maintain the same performance ranking over extended time.[13] Therefore, while the percentile forecasts provide useful information about growth potential under sustained outperformance, the median forecasts may offer a more realistic baseline for long-term projections.

Figure 13: Real GDP Growth Ranking vs Growth Ranking in 10 Years (1950-2019)

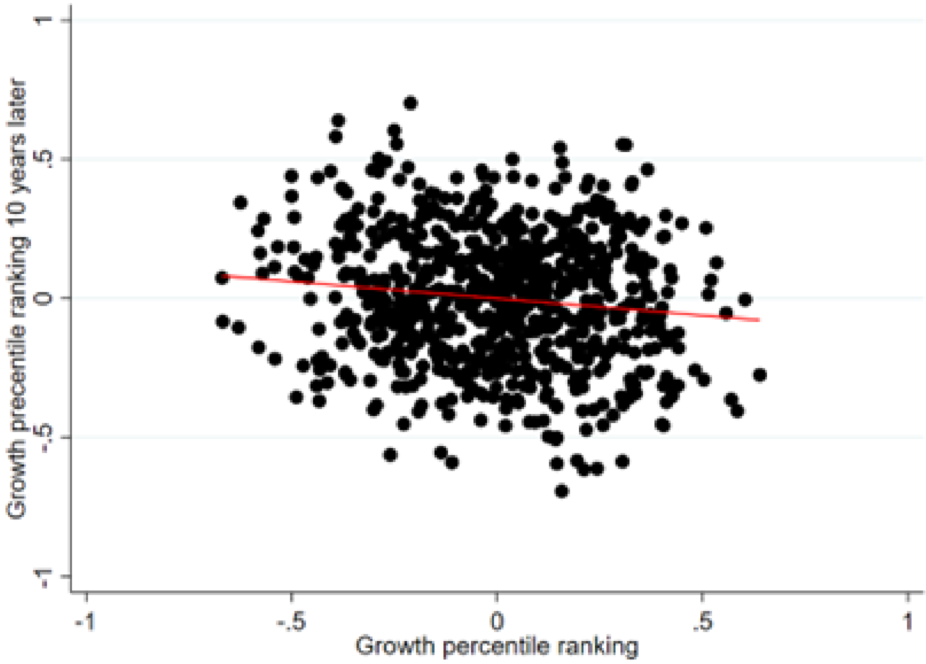

Source: Authors' calculations.

The general pattern of growth moderation over time across most economies merits some qualification. While consistent with the convergence hypothesis, these forecasts implicitly assume the continuation of the relatively stable global technological and productivity growth environment that characterized our post-World War II sample period. This period featured relatively predictable rates of technological change and broad-based peace and prosperity in many regions. However, the forecasts' accuracy could decline significantly if the coming decades were to see more disruptive changes in the pace or distribution of technological progress—for instance, if artificial intelligence or other breakthrough technologies create highly uneven productivity gains across economies. Similarly, major shifts in the global economic order or significant geopoliti-

presented in Figure 13 include country fixed effects.

[13]If run without the country fixed effects, the regression coefficient for the current growth ranking is estimated at around 0.16, with the coefficient for the constant term at 0.42. This implies that a country that is currently ranked at 75th – 80th percentile among peers is expected to see its ranking drop to 54th to 55th percentile in 10 years.

cal disruptions could create growth patterns that deviate substantially from historical reference points. These limitations are inherent to any forecasting method that relies on historical data, but they are particularly relevant for long-term projections in a rapidly evolving global economy.

4.3 Comparing the Two Approaches

Although the two very different forecasting approaches yield different overall findings, they still share some similar broad features. First, growth differs by income group, with the forecast for the lower income group to continue to see higher growth, as shown in Figure 14. Similarly, these growth differentials are forecast to narrow over time, in line with economic convergence. Second, both approaches forecast a moderation of growth over time across all income groups. The growth decomposition suggests that demographics might be an important explanatory factor, reinforced by slower TFP growth and less capital deepening.

Figure 14: Projected Medium- and Long-term Growth

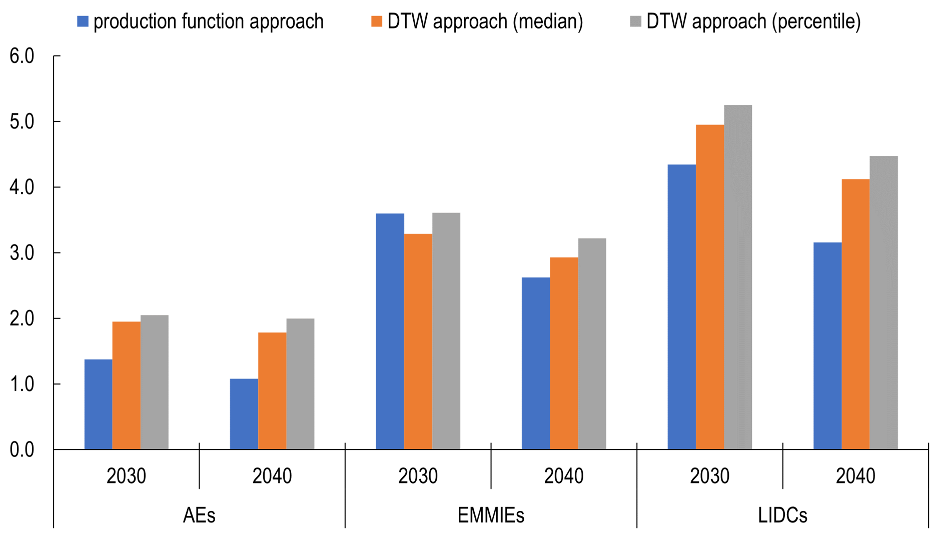

Notes: Only economies with projection results for both methodologies are included.
Source: Authors' calculations.

Yet, the DTW approach almost consistently forecasts higher growth than the growth decomposition. What could explain those differences? The machine learning methodology allows for the inclusion of an extensive pool of historical data for comparison, showing higher forecasting accuracy over longer horizons than the traditional growth decomposition. This also implies that long-term structural factors such as growth convergence and structural transformation would be implicitly taken into account while not incorporated in the growth accounting framework.

In addition, unlike the production-function approach that projects each factor linearly (capital, TFP, etc.), DTW searches for historical analogues that match current structural and macroeconomic patterns. This naturally accounts for complex non-linearities—e.g., rapid structural transformation—that a purely factor-based projection might miss. As a result, the DTW forecasts often imply slightly higher [or lower] growth paths for specific economies, especially those still undergoing substantial structural shifts.

The DTW approach suggests that historical analogies can offer additional insight into the likelihood of sustained structural transformation, which may help policymakers assess whether outlier 'growth miracles' are replicable. By identifying parallels in demographics and sectoral shifts, DTW-based forecasts can inform country-specific strategies to maintain or accelerate convergence.

It is worth noting that, both the growth accounting approach and the DTW method project growth based on historical data—either the country's own observed ongoing trends or the historical patterns of other economies. Therefore, these methods lack the ability to forecast structural changes, for example, significant technological advances, or ambitious reforms which have not been observed in the past and thus to incorporate these into the projections, despite that these changes could substantially alter the growth trajectories. As a result, our projections could be considered as *baseline* projections if the current trends persist and economies follow historical patterns of their own and others.

In addition, our growth projections serve as methodological contributions designed to understand long-term structural patterns through historical analysis, rather than operational forecasts for policy guidance. Unlike IMF country desk forecasts that integrate real-time data and expert judgment on country-specific circumstances, our projections are derived purely from historical patterns and econometric relationships to illustrate forecasting methodologies. Readers should therefore view our results as complementary analytical tools that suggest potential growth paths under historical development patterns, rather than as predictions of most likely outcomes given current policy settings.

5 Conclusion

This paper provides a comprehensive analysis of medium-term growth prospects in Asia through multiple complementary methodologies. Our analysis yields several key insights about both historical growth patterns and future trajectories.

Our examination of historical growth drivers reveals significant variation across Asian economies at different development stages and across time. Through growth accounting decomposition, we find that emerging market and developing economies' outperformance over the past three decades was primarily driven by capital deepening and favorable demographics, accounting for roughly one-third of their total growth. In contrast, advanced economies' growth was mainly driven by TFP and human capital improvements. Our machine learning analysis further identifies economic convergence, demographics, and trade openness as the most important growth determinants, while also highlighting the critical roles of human capital development and institutional quality in sustaining growth momentum.

The paper makes an important methodological contribution by developing an integrated forecasting framework that combines traditional growth accounting with novel pattern recognition techniques. Our application of Dynamic Time Warping to economic forecasting demonstrates particular value for medium-term projections, maintaining more stable accuracy over longer horizons compared to conventional methods. When tested against IMF WEO and EIU forecasts, this pattern-matching methodology shows comparable or superior performance for medium-term horizons while avoiding the systematic optimistic bias often present in traditional forecasts.

Looking ahead, our analysis points to a significant moderation in Asia's growth rates over the next two decades, if current trends persist and economies follow historical patterns of their own and others. Using both growth accounting and pattern-matching approaches, we project aggregate regional growth to decline from pre-pandemic levels of 3.7 percent to around 3.1-3.5 percent by 2030 and 2.2-3.1 percent by 2040. This slowdown reflects several factors: demographic headwinds as populations age, diminishing returns to capital accumulation, and gradual convergence in productivity levels. The decline is expected to be more pronounced in emerging market and developing economies, where growth could fall from pre-pandemic levels of 4.8 percent to 3.8-4.1 percent by 2030 and 2.8-3.6 percent by 2040.

However, significant variation exists across economies, reflecting differences in demographic trajectories, reform momentum, and integration into global value chains. Fast-growing economies like Cambodia, Vietnam, and Bangladesh are projected to maintain robust growth rates mostly above 4 percent through 2040, though with some moderation over time. China's growth is expected to moderate more significantly to around 2.5-3.3 percent by 2040, while India's growth shows more stability at similar levels.

These findings have important implications for the public discussion on economic growth. First, they underscore the importance of structural reforms to offset demographic headwinds and boost productivity growth. To sustain high growth, it is critical to reduce the impact of aging, including by supporting labor force participation and facilitating migration. Reforms to enhance productivity are also critical, and include improving the business and investment conditions, increasing human capital stock, and removing obstacles to efficient resource allocation. Second, they highlight the value of complementing traditional forecasting approaches with methods that learn from historical development patterns. Finally, they suggest the need for careful calibration of medium-term policy frameworks to account for likely growth moderation.

Future research could extend this approach in several directions. The pattern-matching methodology could be enhanced by incorporating additional economic and institutional variables, or by developing more sophisticated ways to weight historical parallels. The approach could also be applied to other regions or to specific sectors within economies. As more data becomes available on how economies navigate various development challenges, the method's ability to identify relevant historical patterns should further improve.

References

Acemoglu, D., & Robinson, J. A. (2012). *Why nations fail: The origins of power, prosperity, and poverty.* Crown Books.

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*(2), 323–351.

Aghion, P., & Howitt, P. (2009). *The Economics of Growth* (Vol. 1). The MIT Press.

Aromí, J. D. (2019). Medium term growth forecasts: Experts vs. simple models. *International Journal of Forecasting*, *35*, 1085–1099.

Babii, A., Ghysels, E., & Striaukas, J. (2022). Machine learning time series regressions with an application to nowcasting. *Journal of Business & Economic Statistics*, *40*(3), 1094–1106.

Bagnall, A., Lines, J., Bostrom, A., Large, J., & Keogh, E. (2017). The great time series classification bake off: A review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery*, *31*, 606–660.

Barro, R. J. (1991). Economic growth in a cross section of countries. *The Quarterly Journal of Economics*, *106*(2), 407–443.

Basu, S. S., Perrelli, R. A., & Xin, W. (2019). External crisis prediction using machine learning: Evidence from three decades of crises around the world. *Computing in economics and finance, Ottawa, Canada.*

Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining*, 359–370.

Birdsall, N., Ross, D., & Sabot, R. (1995). Inequality and growth reconsidered: Lessons from east asia. *The World Bank Economic Review*, *9*(3), 477–508.

Bluwstein, K., Buckmann, M., Joseph, A., Kapadia, S., & Şimşek, Ö. (2023). Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach. *Journal of International Economics*, *145*, 103773.

Breiman, L. (2001). Random forests. *Machine learning*, *45*, 5–32.

Breiman, L., Friedman, J., Stone, C. J., & Olshen, R. (1984). *Classification and regression trees.* CRC Press.

Burman, P., Chow, E., & Nolan, D. (1994). A cross-validatory method for dependent data. *Biometrika*, *81*(2), 351–358.

Cebotari, A., Chueca-Montuenga, E., Diallo, Y., Ma, Y., Turk, M. R. A., Xin, W., & Zavarce, H. (2024). *Political fragility: Coups d'état and their drivers* (Working Paper No. 034). International Monetary Fund.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

Domar, E. D. (1946). Capital expansion, rate of growth, and employment. *Econometrica*, *14*(2), 137–147.

Durlauf, S. N., Johnson, P. A., & Temple, J. R. (2005). Growth econometrics. *Handbook of economic growth*, *1*, 555–677.

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the penn world table. *American economic review*, *105*(10), 3150–3182.

Fernald, J. (2015). Productivity and potential output before, during, and after the great recession. *NBER macroeconomics annual*, *29*(1), 1–51.

Fouliard, J., Howell, M., Rey, H., & Stavrakeva, V. (2021). *Answering the queen: Machine learning and financial crises* (tech. rep.). National Bureau of Economic Research.

Hacibedel, B., & Qu, R. (2023). Understanding and predicting systemic corporate distress: A machine-learning approach. *Journal of Credit Risk*.

Harrod, R. F. (1939). An essay in dynamic theory. *The Economic Journal*, *49*(193), 14–33.

Hausmann, R., Rodrik, D., & Velasco, A. (2005). *Growth diagnostics* (tech. rep.). John F. Kennedy School of Government, Harvard University. Cambridge, MA.

Hellwig, K.-P. (2021). *Predicting fiscal crises: A machine learning approach* (Working Paper No. 150). International Monetary Fund.

Hendry, D. F., & Clements, M. P. (2003). Economic forecasting: Some lessons from recent research. *Economic Modelling*, *20*(2), 301–329.

International Monetary Fund. (2015). Uneven growth: Short- and long-term factors. *World Economic Outlook, April 2015* (Chapter 3).

International Monetary Fund. (2016). Subdued demand: Symptoms and remedies. *World Economic Outlook, October 2016* (Chapter 1).

International Monetary Fund. (2021). *How to assess country risk: The vulnerability exercise approach using machine learning*. Technical Notes; Manuals, Washington DC.

International Monetary Fund. (2024a). Asia-pacific's structural transformation: The past and prospects. *Regional Economic Outlook for Asia and Pacific, November 2024* (Notes).

International Monetary Fund. (2024b). Slowdown in global medium-term growth: What will it take to turn the tide? *World Economic Outlook, April 2024* (chapter 3).

Jones, C. I. (2022). The past and future of economic growth: A semi-endogenous perspective. *Annual Review of Economics*, *14*, 125–152.

Jones, C. I., & Romer, P. M. (2010). The new kaldor facts: Ideas, institutions, population, and human capital. *American Economic Journal: Macroeconomics*, *2*(1), 224–245.

Jones, C. (2016). The facts of economic growth. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of macroeconomics* (pp. 3–69). Elsevier.

Kaldor, N. (1961). Capital accumulation and economic growth. In F. Lutz & D. Hague (Eds.), *The theory of capital* (pp. 177–222). Palgrave Macmillan.

Lundberg, S. (2017). A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*.

Muir, D. V., Novta, N., & Oeking, A. (2024). *China's path to sustainable and balanced growth* (Working Paper No. 238). International Monetary Fund.

Petris, G., & Petrone, S. (2019). Time series matching for economic forecasting. *Journal of Business & Economic Statistics*, *37*(3), 429–441.

Pritchett, L. (1997). Divergence, big time. *Journal of Economic Perspectives*, *11*(3), 3–17.

Racine, J. (2000). Consistent cross-validatory model-selection for dependent data: Hv-block cross-validation. *Journal of econometrics*, *99*(1), 39–61.

Richardson, A., van Florenstein Mulder, T., & Vehbi, T. (2021). Nowcasting gdp using machine-learning algorithms: A real-time assessment. *International Journal of Forecasting*, *37*(2), 941–948.

Rodrik, D. (2013). Unconditional convergence in manufacturing. *The Quarterly Journal of Economics*, *128*(1), 165–204.

Rodrik, D. (2018). An african growth miracle? *Journal of African Economies*, *27*(1), 10–27.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, *98*(5), S71–S102.

Sakoe, H., & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, *26*(1), 43–49.

Sala-i-Martin, X. (1997). *I just ran four million regressions.* National Bureau of Economic Research Cambridge, Mass., USA.

Sala-i-Martin, X., Doppelhofer, G., & Miller, R. I. (2004). Determinants of long-term growth: A bayesian averaging of classical estimates (bace) approach. *American economic review*, *94*(4), 813–835.

Shapley, L. S. (1953). Stochastic games. *Proceedings of the national academy of sciences*, *39*(10), 1095–1100.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, *70*(1), 65–94.

Solow, R. M. (1957). Technical change and the aggregate production function. *The Review of Economics and Statistics*, *39*(3), 312–320.

Stiglitz, J. E. (1996). Some lessons from the east asian miracle. *The World Bank Research Observer*, *11*(2), 151–177.

Strumbelj, E., & Kononenko, I. (2010). An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, *11*, 1–18.

Swan, T. W. (1956). Economic growth and capital accumulation. *Economic Record*, *32*(2), 334–361.

Wade, R. (1990). *Governing the market: Economic theory and the role of government in east asian industrialization.* Princeton University Press.

World Bank. (1993). *The east asian miracle: Economic growth and public policy.* Oxford University Press.

Wu, J. C., & Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, *48*(2-3), 253–291.

Young, H. P. (1985). Monotonic solutions of cooperative games. *International Journal of Game Theory*, *14*(2), 65–72.

Appendix A Identifying Growth Drivers: Algorithm, Estimation, and Explanation

Appendix A.1 Machine Learning Algorithms: Tree-Based Models

In the category of tree-based ensemble methods, this paper considers two different ensemble learning methods for regression tasks, which involves predicting continuous outcomes: Random Forests (Breiman, 2001) and XGBoost (Chen and Guestrin, 2016). Ensemble learning methods for regression tasks are an extension of those for binary classification tasks, in which the base learner is the binary classification tree.

The binary classification tree method (Breiman et al., 1984) uses a decision tree to flag an observation by going from the original complex sample to smaller and purer subsamples. Each decision tree consists of a root node, branches departing from parent nodes and entering child nodes, and multiple terminal nodes which are also called leaves. In the structure of classification tree, leaves represent the flagged classes (determined by the class with the most votes within one leaf) and branches represent the conjunctions of indicators that lead to the classes. Observations in the root node are sent to left or right child nodes according to some splitting rules that identify indicators and corresponding thresholds. Once the whole sample is split into two subsamples, such process is repeated on each child node recursively until each leaf consists of observations in one class, or some stopping criteria are met (e.g., the maximum depth or the minimum leaf size of a tree is reached). In other words, a decision tree is made up by many splits, which consists of a parent node, two child nodes, and branches departing from the parent node and entering child nodes. The indicator and threshold used to split the sample at each node are chosen based on some measures of impurity, such as the Gini impurity index. Because of the recursive algorithm, the binary classification tree structure partitions the classification (or prediction) space into multiple smaller spaces, which allows for a complex relationship between the classification (or prediction) outcome and predictors, such as non-linearities, non-monotonicities, and interactions among indicators.

Binary classification trees are prone to overfitting when a tree grows fully to fit all observations in the training sample, which results in a deep tree with small leaves containing only few observations with strict rules. Such a deep tree will fail to make accurate predictions for new observations because it includes too much noise from the training sample that is irrelevant to new predictions. To reduce overfitting of one single binary classification tree, ensemble models consisting of many binary classification trees were proposed.

Random Forest: Among ensemble models based on binary classification trees, the simplest is Random Forests introduced by Breiman (2001), which applies the general techniques of bootstrap aggregating (bagging). The Random Forest consists of multiple binary classification trees, each of which is grown on a random sample selected with replacement from the training sample, which decreases the variance of the model without substantially increasing the bias. Additionally, it performs random feature sampling such that only a random subset of predictors selected from the entire set are considered at each split, effectively preventing strong correlations among trees. In the end, class predictions for new observations are made by taking the majority vote of classes determined by individual trees, and scores of new observations are calculated by taking the average of scores generated by individual trees. Bootstrap aggregating and feature sampling together help Random Forests prevent overfitting and thus achieve better prediction performance.

XGBoost: In addition to Random Forests, we also consider XGBoost (Chen and Guestrin, 2016), another ensemble learning algorithm with the binary classification tree as the building block. XGBoost employs gradient boosting, in which individual binary trees are trained sequentially. To be specific, each new binary tree is trained to learn the residuals of previous trees, which are the differences between the predicted and actual values of the target variable. Moreover, XGBoost adds a penalty term to the loss function used in training the model, which helps prevent overfitting by discouraging the model from becoming too complex. The complexity is measured by sum of the depth of all trees and the number of trees in the XGBoost model. The depth of a tree in an XGBoost model refers to the number of splits in the tree. A tree with more splits has a greater depth and can capture more complex interactions in the data but may also be more prone to overfitting. Also, the more trees you have in the XGBoost model, the more likely you will capture some "idiosyncratic" pattern only for this sample and suffer over-fitting problem when make prediction on another sample. Thus, the algorithm builds new trees to minimize the errors in the training set while controlling the complexity of the trees to ensure it has good generalization ability.

Instead of using the binary classification tree as the base learner, ensemble learning methods for regression tasks use the regression tree. Different from the binary classification tree, the regression tree uses mean squared error (MSE) as the splitting criteria and each leaf node represents a predicted value which is the average of the target values in that node. In terms of model evaluation, the root mean squared error (RMSE) is usually used in a regression task, while the area under the roc curve (AUC) or some combinations of type I and type II errors are usually used in a classification task.

Appendix A.2 Model Training, Testing, and Hyperparameter Tuning

The sample consists of 22 Asia-Pacific economies over the period of 1970-2019. Each observation is a country-year pair. For each observation of country c in year t, the target variable is the average real GDP growth in the next five year, i.e., from year $t+1$ to $t+5$, and the predictor values are chosen from year t. As such, the empirical framework is formulated as a prediction framework which addresses the concern that the variables could be affected by the real GDP growth if they are included from the same period. The root mean squared error (RMSE) is chosen as the evaluation metric in hyperparameter tuning and model testing.

We employ the block time-series cross-validation method based on Burman et al. (1994) and Racine (2000) to tune hyperparameters and evaluate model performance, which accounts for the cross-sectional dependence in the panel structure of our sample. Exact steps are as described below:

Step 1: Construct 5 training set, validation set pairs based on the entire sample, which consists of observations with years belong to 1970 - 2004, 2005 - 2006; 1970 - 2006, 2007 - 2008; 1970 - 2008, 2009 - 2010; 1970 - 2010, 2011 - 2012; 1970 - 2012, 2013- 2014. Therefore, each training set consists of observations with years since 1970 until the cutoff year (namely 2004, 2006, 2008, 2010, and 2012), and each test set consists of observations in the next two years (namely, 2005-06, 2007-08, 2009-10, 2011-12, and 2013-14).

Step 2: For each ML algorithm, starting with a random set of hyperparameters, we train the model on the training sets and apply the model to the corresponding validation sets to get five RMSEs. In each tree-based algorithm, 1000 trees are constructed, and the hyperparameters we choose to tune for each model are shown below:

Random Forest: maximum tree depth for base learners and subsample ratio of columns when constructing each tree.

XGBoost: L2 regularization term on weights and boosting learning rate.

This step is repeated with different sets of hyperparameters which are chosen by Bayesian optimization over 100 iterations. For each ML algorithm, the set of hyperparameters that yield the lowest average RMSE across the five validation sets are chosen as the optimal set of hyperparameters.

Step 3: Finally, the ML algorithm with its optimal set of hyperparameters that yields the lowest average RMSE is chosen as our final selected ML algorithm and the corresponding hyperparameters.

Appendix A.3 Model Explanation: Shapley Values

To unpack the black box of machine learning models and understand the contributions of predictors, we report the results in terms of Shapley values (Strumbelj and Kononenko, 2010 and Lundberg, 2017)——built on the concept of Shapley values from cooperative game theory (Shapley, 1953 and Young, 1985)—which essentially measure the additive contribution of each predictor to the likelihood of a coup relative to the sample-average predicted probability of a coup.

Beginning with linear models, Shapley value of a predictor for an observation is simply the estimated coefficient multiplied by the observation's value of the predictor. Specifically,

$$\hat{f}(x_i) = \phi_0(\hat{f}) + \sum_{k=1}^{n} \phi_k(x_i; \hat{f}) = \hat{\beta}_0 + \sum_{k=1}^{n} \hat{\beta}_k x_{i,k} \tag{9}$$

where $\hat{f}(x_i)$ is the model prediction for observation x_i, $\hat{\beta}_0$ is the estimated unconditional expected value of $\hat{f}(x_i)$, and $\hat{\beta}_k$ is the estimated slope coefficient for the kth predictor. Shapley value of the kth predictor for observation x_i is calculated as $\hat{\beta}_k x_{i,k}$, and the sum of Shapley values of all predictors for observation x_i is the difference between its model prediction and the average prediction in the training sample.

Shapley values for non-linear model draws on ideas from cooperative game theory, and is implemented in a model-agnostic way, i.e., it offers a way to decompose the model prediction into contributions of predictors for any machine learning algorithm. The idea is to think about a model prediction for an observation as a cooperative game where each predictor value of the observation is a "player" and the prediction is the "payout". In this way, the "gain" of this game is the difference between the payout and the average prediction for all observations, and all predictor values of the observation (all players in the game) collaborate to receive the gain. Then the Shapley value of a predictor for the observation is defined as the average marginal contribution of a predictor value for the observation across all possible coalitions, given the observation values for all other predictors. Specifically, a model prediction can be linearly decomposed as

$$\hat{f}(x_i) = \Phi_0^S(\hat{f}) + \sum_{k=1}^{n} \Phi_k^S(x_i; \hat{f}) \tag{10}$$

$$\Phi_k^S(x_i; \hat{f}) = \sum_{x' \subseteq \{x_1, x_2, \ldots, x_n\} \backslash \{x_k\}} \frac{|x'|!(n - |x'| - 1)!}{n!} \Big(\hat{f}(x_i | x' \cup \{x_k\}) - \hat{f}(x_i | x') \Big) \tag{11}$$

where $x' \subseteq \{x_1, x_2, ..., x_n\} \backslash \{x_k\}$ is the set of coalitions used in the model prediction consisting of all predictors but the kth predictor for which the Shapley value $\Phi_k^S(x_i; \hat{f})$ is calculated, $|x'|$ is the number of predictors included in each coalition except the kth predictor, $\frac{|x'|!(n-|x'|-1)!}{n!}$ is the weighting factor, and $\hat{f}(x_i|x' \cup \{x_k\})$ and $\hat{f}(x_i|x')$ are the model prediction with and without the kth predictor conditional on the set of coalition consisting of all other predictors, which implies that the difference is the pay-off for including the kth predictor, i.e., x_k, in the coalition x'.

Appendix B Identifying Growth Drivers: List of Variables

Table 2: List of Variables

Variable Name	Transformation	Category	Source
GDP per capita	level	Development and Demographics	World Economic Outlook Database
Population Growth	5avg	Development and Demographics	World Development Indicator Database
Working-Age Population Share	5pch	Development and Demographics	World Development Indicator Database
ToT %Change	level	Global and Openness	World Economic Outlook Database
Export Partners' Growth	5avg	Global and Openness	World Economic Outlook Database
Fed Funds Rate (Shadow)	5avg+5fd	Global and Openness	Wu and Xia (2016)
Import Partners' Growth	5avg	Global and Openness	World Economic Outlook Database
Inward FDI	5avg	Global and Openness	IFS and BoP Databases
Trade Openness	5avg+5fd	Global and Openness	World Economic Outlook Database
Country Uncertainty	level+5avg+5pch	Macroeconomics	World Uncertainty Index Database
NEER %Change	level+5avg	Macroeconomics	IFS and BoP Databases
REER %Change	level+5avg	Macroeconomics	IFS and BoP Databases
External Debt	5avg+5fd	Macroeconomics	World Economic Outlook Database
Fiscal Balance	5avg+5fd	Macroeconomics	World Economic Outlook Database
Inflation	5avg+5fd	Macroeconomics	World Economic Outlook Database
M2 Growth	5avg	Macroeconomics	World Economic Outlook Database
Private Credit	5avg+5fd	Macroeconomics	World Economic Outlook Database
Public Debt	5avg+5fd	Macroeconomics	World Economic Outlook Database
Real Interest Rate	5avg+5fd	Macroeconomics	World Economic Outlook Database
Reserves Coverage	5avg+5fd	Macroeconomics	World Economic Outlook Database
Unemployment Rate	5avg+5fd	Macroeconomics	World Economic Outlook Database
Governance	5pch	Structural Conditions	Worldwide Governance Indicator Database
Human Capital	5pch	Structural Conditions	Penn World Table; Feenstra et al. (2015)
Labor Force Participation	5pch	Structural Conditions	World Development Indicator Database
Political Stability	5pch	Structural Conditions	Worldwide Governance Indicator Database
Share of Wage and Salaried Workers	5pch	Structural Conditions	World Development Indicator Database

Notes: The table lists variables used in the machine learning exercise to identify growth drivers. The column *Transformation* indicates how variables are transformed and included as the predictor, with level, 5avg, 5pch, and 5fd indicating the variable itself, five-year average, five-year average percentage change, and five-year average level change. The column *Category* presents the classification of the variable based on the four categories including development and demographics, global and openness, macroeconomics, and structural conditions. The last column *Source* presents the data source with IFS and BoP databases be the International Financial Statistics and Balance of Payments databases.

Appendix C DTW Country Case Study: Cambodia

To better demonstrate how the DTW PatternSync method works in practice and its value for country-specific analysis beyond simple point forecasts, we examine Cambodia as a detailed case study. The method's value lies not just in generating growth projections, but in providing a structured way to analyze a country's growth potential through the lens of historical parallels. By identifying similar country-periods, the method allows us to examine patterns of growth sustainability, study the characteristics of successful peer economies, and understand what factors differentiated sustained high performing peers from those that are less so. This comparative historical perspective can be particularly valuable for policymakers seeking to understand both opportunities and challenges in maintaining high growth trajectories.

Using the DTW method, we identify a set of historically similar country-periods by selecting the top 2 percent most similar to Cambodia's 2015-19 target period. This analysis yields 156 historically similar episodes, with a notable concentration among East and Southeast Asian economies at different stages of development. Table 3 presents some examples. These economies tend to share common economic characteristics with Cambodia in various dimensions, providing relevant benchmarks for growth potential. Through analyzing these similar episodes, we can not only project Cambodia's medium-term growth path but also gain insights into the factors that historically enabled some economies to maintain exceptional growth while others experienced mean reversion.

Figure 15 plots the forecasted medium term growth path by aggregating the subsequent growth rates across these similar country-periods using median, average, and weighted average aggregation. The forecasted potential growth for 2030 ranges from 5.8 percent (average forecast) to 6.3 percent (medium forecast). For the target period, Cambodia's growth outperformed most of the similar country-periods in our sample, with its trend growth situated at 75th to 80th percentile of the top 2 percent most similar cohort. Therefore, as a reference Figure 15 also plots the growth forecasts if Cambodia continues to perform at this percentile range over the medium term. These percentile forecasts point to a potential growth of 7.5 to 7.9 percent for 2030.

The 75^{th} to 80^{th} percentile forecasts may better serve as an aspirational goal rather than the baseline forecast given historical patterns in the data, due to the mean reversing pattern observed the cross-country growth performance ranking, i.e. economies that are significant

Table 3: Examples of Top Similar Country-Periods to Cambodia 2015-2019

Country	Start Year	End Year
India	2001	2005
Indonesia	2002	2006
Lao P.D.R	2008	2012
Sri Lanka	1983	1987
Japan	1951	1955
Bangladesh	2004	2008
Korea	1973	1977
Poland	1973	1977
Philippines	2004	2008
Bolivia	2005	2009
Djibouti	2011	2015
China	1986	1990
Vietnam	2010	2014
Peru	1990	1994
Morocco	2002	2006
Thailand	1985	1989
Uzbekistan	2001	2005

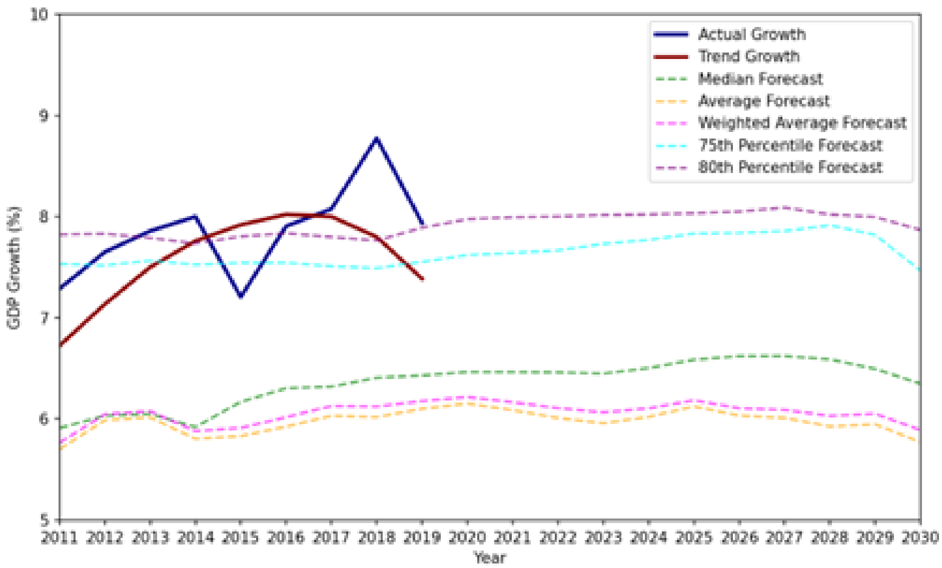

Figure 15: Medium Term Growth Forecast for Cambodia Using Top 2 Percent Most Similar Country-Periods

outperformers in a given period tend to see their ranking drop over time, and vice versa (see Section 4.2.5). The mean-reverting pattern can also be observed by looking at Cambodia's similar peers that ranked at the same growth performance level during their base period of comparison. Figure 16 displays the subsequent growth for economies that are similar to Cambodia's target period (i.e. 2015-2019) during their base period, while having growth performance also similar, i.e. their growth rate in the base period is in the 75th to 80th percentile range among all the country-periods that have a DTW similarity score to Cambodia within the top 2 percent. Among the 7 country-periods that fit these criteria, the average and median growth rates 10 years after the base period are 6.4 percent and 5.9 percent respectively. Notably, in all except one instance (China 1996-2000) among the seven, growth performance has dropped out of the 75th-80th percentile range after a decade. This reversion to the mean in growth rate is not unique to these 7 instances that are in the same growth range as Cambodia. Among all the similar economies that ranked equal to or higher than Cambodia in terms of growth performance during their base period (i.e. their growth rate during base period is in the 75th to 100th percentile range of all similar economies to Cambodia), which include 40 country-period instances, the average and median growth rates a decade after their base period are 5.8 percent and 6.4 percent respectively. That is, even though this sub-group demonstrated very strong growth during their period of comparison with Cambodia 2015-2019, their performance over the medium term are overall indistinguishable from the sample mean or median.

A small group of "outliers" in Cambodia's similar country-period cohort did exhibit sustained high growth over an extended period. The outliners are mainly from 3 economies (their base periods in parentheses): Japan (1956-1967), South Korea (1971-1986), and China (1996-2002). These economies had a growth performance during the base period ranked at 75th percentile or above among Cambodia's top 2 percent most similar peers, while they also managed to remain among the 75th percentile 10 years after the base period. The growth literature identified some commonalities among these outliers that may have contributed to their unusually high level of growth sustainability. Several studies have highlighted the role of export-oriented industrialization policies and strategic integration into global value chains as key drivers of rapid and sustained growth in these East Asia economies (Rodrik, 2018; Wade, 1990; Stiglitz, 1996). Additionally, high investments in human capital through education and health infrastructure have been cited as critical factors that allow these economies to continuously step up into higher

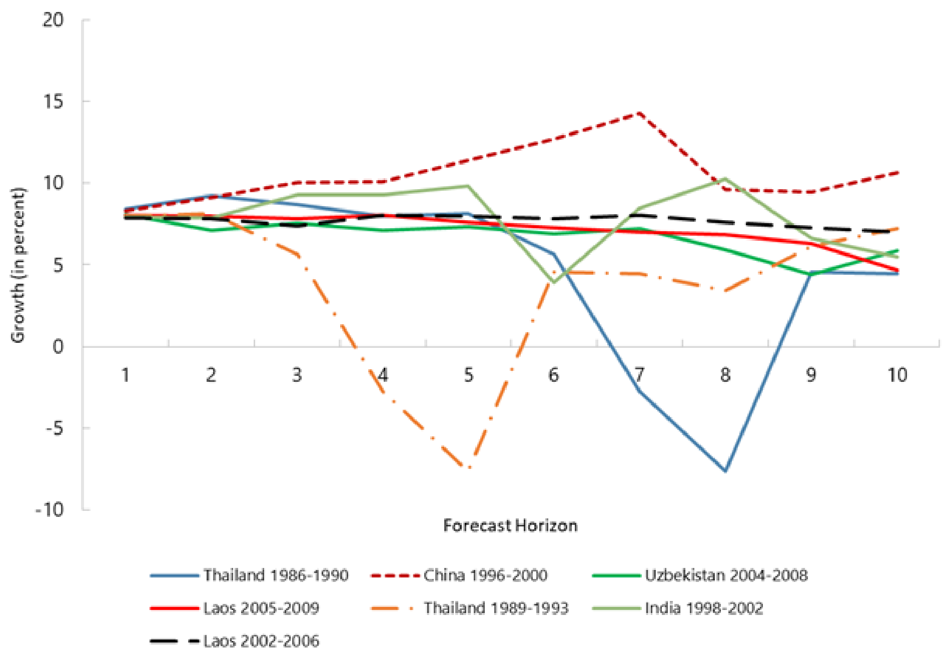

Figure 16: Subsequent Growth for Similar Country-Periods That Are in the Same Growth Performance Range (75th-80th Percentile) as Cambodia 2015-2019

tiers of global value chain (World Bank, 1993; Birdsall et al., 1995). Institutional reforms that improved governance, protected property rights, and fostered an enabling environment for private sector development have also been identified as important enablers of sustained growth (Acemoglu and Robinson, 2012; Hausmann et al., 2005). Data shows that compared to these outliners during their base periods, Cambodia appears to lag behind in terms of human capital and labor productivity.

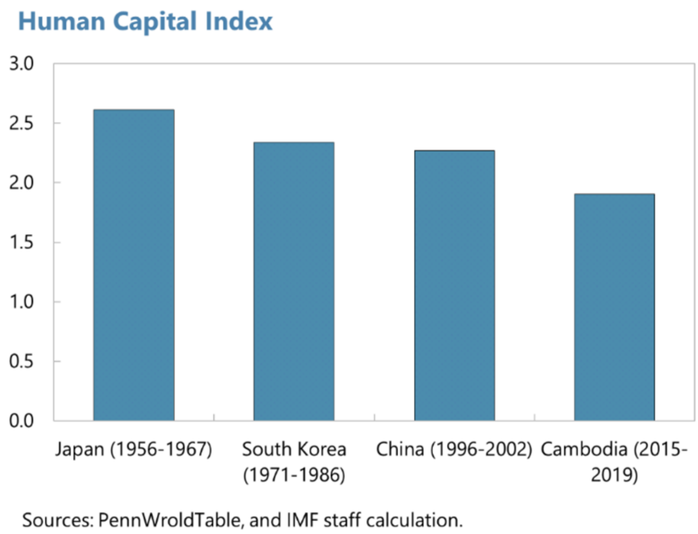

Figure 17: Human Capital Index Comparison

The Cambodia case study demonstrates how the DTW PatternSync framework provides analytical value beyond simple numerical forecasts. First, by systematically identifying historical parallels, the method enables a structured analysis of potential growth trajectories - in Cambodia's case, suggesting baseline growth of 5.8 to 6.3 percent by 2030. More importantly, the framework reveals patterns in growth sustainability through a rich set of comparable episodes, helping policymakers understand the statistical likelihood of sustained outperformance. The method also facilitates systematic identification of critical success factors by allowing detailed examination of historical peers who maintained exceptional growth - in this case, highlighting the role of export-oriented industrialization, human capital investment, and institutional reforms in sustained high growth cases like Japan, South Korea, and China. This combination of quantitative forecasting with tailored historical pattern matching makes DTW PatternSync a useful tool for country-specific growth analysis, especially for developing economies where traditional forecasting approaches may be constrained by data limitations or evolving economic structures. The framework provides a systematic way to learn from relevant historical experiences and identify potential policy priorities for improving growth performance.

PUBLICATIONS

Dissecting Medium-Term Growth Prospects for Asia

Working Paper No. WP/2025/168