Dynamic Journal of Arts and Social Science Research
Volume 2, Issue 2 (August 2026), ISSN: 3135-4939

Time-Varying Effects of Trade Openness and FDI on Economic Growth in Bangladesh: Evidence from a Bayesian TVP-VAR-SV Model

Afif Ibn Munir [1]
Anita Azad Riya [2]

[1] *Department of Economics, European University of Bangladesh, Dhaka, Bangladesh.*
[2] *Department of Economics, Varendra University, Dhaka, Bangladesh.*

ARTICLE INFO

Article history:
Received: 18/06/26
Revised: 21/07/26
Accepted: 19/08/26
Published: 26/08/26

How to Cite:
Munir, A. I. & Riya, A. A. (2026). Time-Varying Effects of Trade Openness and FDI on Economic Growth in Bangladesh: Evidence from a Bayesian TVP-VAR-SV Model. *Dynamic Journal of Arts and Social Science Research*. 2(2), 103-127.

Corresponding Author:
Name: Afif Ibn Munir

Email: afif.edu77@gmail.com

Journal Home: https://djassr.com/

Check Updates

ABSTRACT

Bangladesh's growth model is widely described as trade- and FDI-led, yet decades of empirical work assume the underlying growth elasticities of trade openness and foreign direct investment (FDI) are fixed across five decades spanning trade liberalisation, an exchange-rate float, and a compliance-cost shock to the ready-made-garment sector. This study tests that assumption using a Bayesian time-varying-parameter vector autoregression with stochastic volatility, estimated via Hamiltonian Monte Carlo with Minnesota-style shrinkage priors on annual World Bank data for Bangladesh, 1972-2024, the first application of this framework to Bangladesh's trade-FDI-growth nexus. Neither elasticity shows statistically detectable time variation anywhere in the sample, a qualified-null result that survives a three-fold loosening of the coefficient shrinkage prior. The FDI coefficient stays positive and economically meaningful throughout but is never distinguishable from zero, while the trade-openness coefficient hugs a narrow band around zero. Stochastic volatility, by contrast, moves substantially and in economically interpretable ways, with a single data-driven structural break in 2002 rather than at the calendar dates conventionally used to date Bangladesh's policy regimes. These findings suggest that the regime shifts long treated as watersheds for the trade-FDI-growth relationship are better understood as episodes of shifting shock variance than as changes in the elasticities themselves. For policy, this implies that expanding trade volume alone is unlikely to shift Bangladesh's growth-transmission mechanism; diversification toward higher-value-added, more backward-linked production, alongside strong shock-absorption buffers around periods of rising trade and FDI variance, is the more promising lever.

Keywords: Trade openness, foreign direct investment, TVP-VAR-SV, stochastic volatility, Bayesian econometrics, Bangladesh.
JEL Classification: O47, F14, F21, C11

1. INTRODUCTION

Bangladesh's growth record since independence is conventionally told as a trade- and FDI-led story. A country that moved from a heavily protected, aid-dependent economy in the 1970s to one of the world's largest ready-made-garment (RMG) exporters. Trade rose from a roughly a fifth of GDP in the early 1970s to a peak of almost half by the mid-2010s. That narrative matters well beyond academic interest. As Bangladesh approaches graduation from Least Developed Country (LDC) status and the preferential market access that has underwritten much of its export growth, policymakers need to know not just that trade and FDI have historically accompanied growth, but whether the strength of that relationship has itself been stable or shifting. A policy lever calibrated to yesterday's elasticity has limited use if that elasticity has since moved, which is the question this paper takes up.

The empirical literature underpinning the trade- and FDI-led growth narrative, however, rests almost entirely on the assumption that it has not moved. Studies applying the Johansen-Juselius procedure, vector error correction models, and autoregressive distributed lag (ARDL) bounds testing to Bangladesh have each estimated a single, fixed elasticity linking trade openness and FDI net inflows to GDP growth. Some research adopted a step-change dummy in 1991 or a rebuilt exchange-rate series around the 2003 float. But never a coefficient allowed to evolve continuously across the sample. This is a striking omission given that the same five decades span trade liberalisation, an exchange-rate regime change, a compliance-cost shock to the RMG sector following Rana Plaza, and the run-up to LDC graduation, any one of which is a plausible candidate for a shift in the underlying transmission mechanism. No existing Bangladesh-specific study of the FDI-trade-growth nexus has applied a time-varying-parameter framework to test this directly. Even though such methods have been used productively elsewhere to separate genuine coefficient drift from shifts in shock volatility.

This study estimates a time-varying-parameter vector autoregression with stochastic volatility (TVP-VAR-SV), following Primiceri (2005), for GDP growth, trade openness, and FDI net inflows in Bangladesh over 1972-2024. The approach lets the coefficients linking trade and FDI to growth, and separately the volatility of the shocks hitting the system, evolve as latent stochastic processes rather than assuming is fixed. This is well suited to the question at hand: where earlier work could only test for a single break at a pre-specified date, a TVP-VAR-SV traces out the entire coefficient path and lets the data indicate, rather than assume, when and whether the relationship moved. Set against the Johansen/VECM and ARDL bounds-testing tradition summarized above, the advance is structural rather than incremental. Those methods can test whether a relationship holds at one or two pre-chosen dates. Whereas the coefficient path estimated here is continuous over the full sample, so a shift can be located wherever the data puts it rather than wherever the research looked.

The model is estimated by Bayesian Markov chain Monte Carlo using Hamiltonian Monte Carlo with the No-U-Turn Sampler in PyMC, on World Bank World Development Indicators data for Bangladesh. Because a three-variable VAR (1) with a fully time-varying covariance structure implies far more evolving states than the roughly fifty annual observations available can comfortably support, the design incorporates a pre-registered convergence gate and a documented fallback. It freezes the contemporaneous covariance loadings to constants where necessary so that substantive conclusion rests only on parameter that demonstrably converged.

The central finding is a qualified null rather than the regime-shifting pattern the design was built to detect. The growth elasticities of trade openness and FDI show no statistically detectable time variation anywhere in the sample. The trade coefficient stays within a narrow band close to zero throughout, and the FDI coefficient, while consistently positive and economically non-trivial at around 0.6, is never distinguishable from zero at conventional credibility levels. This result holds under a three-fold loosening of the coefficient shrinkage prior, which strengthens rather than weakens confidence that it reflects genuine information content rather than an artefact of prior tightness. Volatility, by contrast, moves substantially and in economically legible ways, peaking around the 1974 post-independence period and tracing a distinct rise in trade-openness volatility from 1990 into the mid-2010s. Altogether these results suggest that the regime shifts long treated as watersheds for Bangladesh's trade-FDI-growth relationship are better understood as episodes of shifting shock variance than as changes in the elasticities themselves. This is a reframing, rather than a rejection of the trade- and FDI-led growth account.

The analysis is confined to a single-country, annual frequency setting: Bangladesh over 1972-2024, with GDP growth, trade openness, and FDI net inflows as the endogenous variables of interest and gross capital formation and exchange-rate movements entering as exogenous controls. The remainder of this paper proceeds as follows. The next section reviews the relevant literature on the trade-FDI-growth nexus in Bangladesh and on time-varying-parameter methods more broadly, followed by a description of the data and methodology, the empirical results, a discussion of their implications, and concluding remarks.

2. Literature Review

On the basis of decades of empirical research that assumes underlying elasticities as fixed, the growthy model of Bangladesh is often described as trade and FDI led. This review traces that literature, the alternative offered by time-varying parameter methods, and the case for the Bayesian TVP-VAR-SV specification used in this study.

2.1 Fixed-parameter models

Without changing the notion of a single, fixed elasticity, studies on the relationship between FDI, trade and growth in Bangladesh have moved forward through multiple generations built on the fixed-coefficient VAR of Sims (1980). Once researchers found that the underlying series were non-stationary, early regression gave way to cointegration, and the Johansen-Juselius method and corresponding vector error correction model (VECM) became the standard tool. Along term equilibrium was found by Adhikary (2010) where FDI and capital formation increased GDP while trade openness had a negative, if diminishing impact. A bidirectional causal relationship FDI and growth as well as a one-way effect from trade openness and growth was reported by Haque and Amin (2018). The autoregressive distributed lag (ARDL) bounds testing approached has been more favored in recent studies because of its tolerance of mixed orders of integration along with its reliability in short samples. Considering FDI as the dependent variable, a bootstrap ARDL test found no long-run cointegration, which was used to argue that trade openness drives growth in Bangladesh (Saleem et al., 2020). ARDL and dynamic OLS were used by Ridwan et al. (2025), who identify a negative long-run relationship between trade openness and growth, attributable to the country’s dependence on imported technology.

A second part of this literature reports either weak or null results. According to Tabassum and Ahmed (2014), between 1972 and 2011 neither trade openness nor FDI significantly boosted growth, and no long-term relationship between FDI and GDP was established at all. Some inconsistencies also exist in short-run results: some research found bidirectional causality, whereas other research found growth leading FDI rather than the reverse. This is at least compatible with a booming economy drawing foreign investors rather than FDI causing the boom. Instead of direct estimation, pre-established regime splits have been used to manage this instability. Majumder (2016) compared pre- and post-1991 sub-periods and showed that trade and FDI ratios rose sharply after liberalisation, while Razzaque et al. (2017) built a real exchange rate series around the 2003 float but retained a VECM framework with stable long-run coefficients throughout. These are step-change assumptions rather than tests of whether the relationship moved continuously.

Islam (2016) found that GSP suspension and the Rana Plaza collapse had an adverse effect on FDI inflows, a conclusion drawn from the same post-2013 period that other single-elasticity studies read differently, which is difficult to reconcile in a model that assumes one constant elasticity. FDI data in developing countries is also known to underreport informal investment and treat reinvested earnings inconsistently across data vintages, and the standard proxies, trade as a share of GDP and FDI inflows as a share of GDP, are blunt instruments that say nothing about the makeup of the export basket (Huchet-Bourdon et al., 2018). Adding a dummy variable for 1991 or 2003 does not address this; it merely signals a link whose strength may be changing.

The same limitation shows up outside Bangladesh. Cieślik and Tarsalewska (2011) estimate trade and FDI openness jointly for 97 developing countries over 1974- 2006 and find that both channels contribute positively to growth on average but only by pooling across countries with very different trade and FDI histories. It is exactly the kind of heterogeneity a single country, single elasticity estimate cannot separate from genuine time-variation. Using an instrumental-variable panel smooth-transition model for African economies Kinfack and Bonga-Bonga (2023) go further and show that the trade-openness elasticity of growth switches sign with the level of development and it is evidence that a supposedly "fixed" elasticity in these constant-parameter frameworks is often a within sample average hiding real underlying movement. When combined, the Bangladesh-specific literature above and this broader evidence make the same point from opposite direction. That is, the constant-parameter estimators and whichever generation of technique they use, are simply not built to tell a genuinely stable elasticity apart from one that has been averaged over a shifting one.

2.2 Structural break studies

A recurring response to the instability noted above has been to test for breaks rather than to let coefficient evolve continuously. The Zivot-Andrew's test tries to address this by searching for the most likely break date on its own, without requiring the researcher to specify one in advance. Even so, it remains sensitive to how lag length is chosen, also it can perform poorly when a series actually has more than one break (Ahmad et al., 2024). ADF and KPSS frequently give opposite results in short annual macroeconomic samples like this one. When that happens, the more defensible approach is to rely on economic judgement about the series rather than simply following whichever test statistics looks more convincing. A variable that is bounded by regulatory or capacity limits, for example, has good reason to be treated as stationary in levels even if a formal test doesn't clearly reject a unit root (Adeleye et al., 2021).

A similar issue comes up when trying to date structural change across all three series together rather than one at a time. The Bai-Perron framework is still the standard method for detecting multiple breaks, but it has no actively maintained Python implementation, and it assumes linear behaviour within each segment because it relies on segmented ordinary least squares (Bai & Perron, 2003). A more practical alternative in this software environment is the Pruned Exact Linear Time (PELT) algorithm that is applied here with a radial basis function cost. PELT solves essentially the same penalized optimisation problem as Bai-Perron but does so at linear rather than quadratic computational cost, at the cost of being sensitive to the penalty parameter choses (Killick & Eckley, 2014). It is also worth noting that the breaks detected by these methods do not always match the dates a researcher expects: the literature on purchasing power parity breaks, for instance, consistently finds that statistical breaks show up earlier or later than the policy event that prompted the search. For that reason, the more defensible way to use a detected break is as an independent cross-check on calendar markers such as 1991, 2003 and 2013 but not as a replacement for them.

This distinction matters for how the calendar-dummy studies discussed above should be read. Majumder (2016) and Razzaque et al. (2017) locate their breaks in 1991 and 2003 because those are the policy dates, not because a data-driven procedure independently pointed there and whether the joint second-moment structure of the series actually moves at those points rather than at some other year is an empirical question the calendar-dummy literature has not asked and this study returns to directly.

Formal Bayesian model comparison was used by (Chan & Eisenstat, 2015) which showed that stochastic volatility is often the component doing most of the work while coefficient of variation adding relatively little depending on the dataset. In this regard, Chan and Eisenstat (2015) as well as Terrell (2026) are useful as both demonstrate that a well-specified time-varying model can legitimately produce a flat result. The order of integration for GDP growth, trade openness and FDI net inflows need to be established before estimating any of the above models. The Zivot-Andrew's test tries to fix this by searching for the most likely break date on its own without requiring the researcher to specify one in advance.

2.3 Time-varying parameter models

Cogley and Sargent (2005) were first to allow coefficient themselves to drift, applying a random-walk specification to post-war US inflation while keeping the contemporaneous covariance structure fixed. Bayesian VARs used Minnesota-style shrinkage prior to relax the dimensionality problem this causes in small systems. Primiceri (2005) expanded this to provide simultaneous evolution of the entire parameter set, coefficients, volatility, and off-diagonal loadings of the covariance matrix. This specification continues to serve as the model class's reference point. Later research has debated the extent to which that flexibility is actually needed. The framework has been pushed toward large systems that require sparsification to remain tractable (Chan, 2022).

As for time variation, the cross-country evidence is inconsistent rather than consistent. Umar and Aldio (2026) found significant parameter movement in Indonesia around the global financial crisis and COVID-19, formally confirmed with Sup- Wald tests against parameter stability. Terrel et al. (2026) report, by contrast, that a constant-coefficient VAR with stochastic volatility fit six small open economies better than the fully time-varying alternative, with the transmission mechanism stable while shock volatility changed.

Malik and Sah (2024) observed that FDI’s impact on growth across BRICS economies remained nearly flat over the medium term, which they read as evidence that institutional quality changes only slowly. Given that only 52 annual observations are available, Bangladesh's data situation resembles the small-sample settings found in CEMAC-country BVAR studies more than it does the quarterly or monthly datasets typically used for larger emerging economies. To date, no study on Bangladesh has examined the FDI-trade-growth relationship using a time-varying framework; that is the gap this study addresses, rather than any prior assumption that the elasticities have shifted.

Wang et al. (2025) add a methodological caution relevant here: their randomized test for distinguishing whether apparent coefficient movement reflects a genuine stochastic process or a deterministic function of time finds, across several macroeconomic and financial applications, that a deterministic-time specification is often preferred once the test is applied formally. That does not itself decide the Bangladesh case, but it is a reminder that a flat or barely moving coefficient path, of the kind this study eventually reports, is at least as consistent with the wider time-varying-parameter literature as a dramatically drifting one would be.

2.4 Bayesian Estimation

Primiceri's original estimation strategy combined a Carter-Kohn forward-filtering-backward-sampling routine with the Kim-Shephard-Nelson mixture approximation to linearise the stochastic volatility block. This approach scales poorly as the number of latent states grows and tends to produce highly autocorrelated drawings. Since then, practice has shifted toward Hamiltonian Monte Carlo (HMC) and its self-tuning variant which is the No-U-Turn Sampler (NUTS), which treats the entire sequence of latent states as parameters to be sampled jointly, rather than marginalising them out through a filter. This is now the default approach in probabilistic programming environments such as PyMC and it comes with a diagnostic that filtering methods do not offer directly the divergent transition. This flags regions of the posterior where the sampler's leapfrog integrator has failed to conserve energy, signalling that part of the state space has not been properly explored.

Convergence in this setting is usually assessed through the potential scale reduction factor (R-hat) and effective sample size (ESS). Current practice treats an R-hat below 1.01 as the working threshold which is tighter than the 1.1 once considered acceptable and recommends an ESS of at least 400 for reliable posterior summaries (Roy, 2019). Neither statistic is sufficient on its own. Chains can appear to agree with each other while all being trapped in the same local region of a covariance block which is a particular risk in short annual series where the number of states can run into the hundreds against only around fifty observations (Vats & Knudson, 2021).

The choice of shrinkage prior underlying all of this is itself a live methodological question, not a settled default, and it deserves comparison against the wider Bayesian VAR literature rather than being taken for granted. Bańbura et al. (2010) show that Minnesota-style shrinkage, with the degree of shrinkage set in relation to the cross-sectional dimension of the system, lets a VAR absorb far more variables than its sample size would otherwise support while still producing credible impulse responses, which is the same logic, scaled down, that justifies the Minnesota-style prior used in this Bangladesh application at T = 52. When the covariance-loading block fails this check, a documented response in the literature is to freeze those coefficients at a constant value while leaving the coefficients and volatilities free to evolve. This is treated as a genuine, if reduced, time-varying estimation rather than an abandonment of the approach. The reasoning is that a posterior concentrating near zero is itself informative as it tells us there is an absence of time-variation in that block.

2.5 Small-sample estimation

A recurring assumption in applied work is that a time-varying model only remains if it finds movement. That assumption does not hold up against the theoretical literature on parameter stability. Canova et al. (2015) show that when parameter variation is exogenous to the system, the structural dynamics implied by a time-varying model, its impulse responses in particular, are identical to those from a time-invariant one. A flat coefficient path is therefore not evidence that the modelling choice was wasted but a legitimate finding about the underlying transmission mechanism. Mendoza et al. (2026) makes the same point from the estimation side, showing that loosening the priors governing coefficient and volatility innovation by an order of magnitude and still finding the paths remain flat, actually strengthens a stability finding rather than undermining it, since it rules out the possibility that a tight prior alone was responsible.

Small-sample settings sharpen the case for taking such findings seriously rather than treating them as a modelling failure. He et al. (2008) show that the asymptotic theory behind standard parameter-constancy tests becomes an unreliable guide once the sample shrinks toward the size typical of annual macroeconomic series, and Clements and Mizon (1991) note that the high dimensionality of unrestricted VARs relative to short samples is precisely what motivates informative priors in the first place, rather than being a defect specific to any one dataset.

Comparisons between constant-parameter models with stochastic volatility (CVAR-SV) and the full Primiceri specification have repeatedly found the constant-coefficient version difficult to beat once volatility alone is allowed to move. Chan and Eisenstat (2018) performed a formal Bayesian model comparison for the US and Australia and found that most of the gain a full time-varying VAR offers over a constant-parameter one comes from stochastic volatility rather than coefficient drift, with Bayes factors often favouring a constant-coefficient VAR with stochastic volatility over the fuller specification.

Clark and Ravazzolo (2015) reach a compatible conclusion from a pure forecasting angle: comparing random-walk stochastic volatility against stationary-AR volatility, fat-tailed errors, GARCH, and mixture-of-innovations specifications of US macroeconomic forecasting, they find the conventional stochastic-volatility specification dominates the alternatives for density forecasting, which is one further reason to expect that, in a short annual sample such as this one. The stochastic-volatility block is likely to be where the informative time variation actually resides, rather than the coefficient block. This is consistent with the broader stochastic volatility literature's central claim that separating shifts in the size of shocks from shifts in the underlying transmission of those shocks is itself the analytical contribution regardless of which specific coefficients turn out to move.

Stochastic volatility modelling of this kind traces back to the recognition that macroeconomic shock variance is not stable over time, and that treating it as constant biases standard errors even where the coefficient estimates themselves remain consistent. The shift from GARCH-family, observation driven variance models toward the parameter-driven, log-volatility random walk used in Primiceri (2005) reflects a judgement that macroeconomic shock variance tends to move in persistent, near-permanent shifts, rather than the mean-reverting clustering typical of high frequency financial data.

For a Bayesian TVP-VAR-SV applied to Bangladesh, this body of work is what makes it possible to report a stable coefficient path (which is checked against a looser prior) alongside a clearly evolving volatility profile as two separate and equally reportable findings, rather than treating the first as a disappointment relative to the second.

2.6 The gap this study addresses

Taken together, the Bangladesh specific literature has not moved past the constant parameter assumption despite five decades that include trade liberalisation, an exchange rate float, a compliance cost shock to the RMG sector, and the approach to LDC graduation. Any one of which is a plausible candidate for a shift in trade or FDI elasticity of growth. The broader TVP-VAR-SV literature offers both the tools to test that possibility directly and a theoretical basis for treating a stable coefficient path as a real answer rather than a null result to be explained away, provided the estimation is properly diagnosed and the finding is checked against a looser prior. This study applies that framework via NUTS estimation in PyMC with a pre-registered convergence gate and a documented fallback for the covariance-loading block to the FDI-trade-growth nexus in Bangladesh over 1972-2024, closing a gap that spans both the country-specific and the methodological literature.

3. METHODOLOGY

3.1 Research Design

This study uses time-series Bayesian econometric design to test whether the growth elasticities of trade openness and foreign direct investment (FDI) in Bangladesh have varied over 1972-2024. Existing literature on this topic assumes a constant parameter, that these elasticities are fixed across the time frame. The introduction sets out why that assumption is questionable given the five decades of structural change the sample spans; it is not repeated here. Our design allows the coefficients, and separately the volatility of shocks to evolve as latent stochastic processes estimated jointly with the model. The approach is descriptive and inferential rather than causal. Thus, there has been no exclusion restriction or instrument is invoked. At the same time, the object of interest is the time path of an estimated conditional association, not a treatment effect. Thus, there is no need for identification-strategy subsections are required for this design.

Why a Bayesian rather than a frequentist time-varying-parameter model. A frequentist TVP-VAR would typically estimate the coefficient path via Kalman filtering with maximum-likelihood estimates of the state-innovation variances Q, which at T = 52 and roughly 15-18 evolving states is genuinely difficult likelihood surface to optimize reliably. It also offers no natural mechanism for encoding the researcher's prior belief that coefficients should move gradually rather than erratically. The Bayesian approach used here instead treats the innovation variances as fixed, informative priors (see the priors described just below) and estimates the full joint posterior over the coefficient path via simulation. This has two practical advantages for a short annual sample. Firstly, it makes the degree of coefficient smoothness an explicit statable modelling choice rather than an estimated nuisance parameter. Secondly, it yields credible intervals that properly reflect the compounded uncertainty of a state space with more latent states than observations, which a plug-in frequentist standard error would understate.

Why Bayesian shrinkage compensates for small-sample bias. With T = 52 and around 15-18 time-varying states, an unrestricted estimate of the coefficient path is not identified in any meaningful sense. There are more parameters evolving over time than there are data points to estimate them from. Minnesota-style shrinkage addresses this by centring the prior on economically sensible defaults (own-lag coefficients near their OLS estimate, cross-variable and higher-lag coefficients shrunk toward zero). By keeping the innovation variance Q small relative to the training-sample coefficient variance, so that the coefficient path is only pulled away from that prior centre where the likelihood provides strong evidence to do so. In effect, the prior substitutes for the additional data a longer or higher-frequency sample would otherwise provide. This is why the paper treats the three-fold prior-loosening exercise in the Diagnostic and Robustness Checks below as informative about whether the null finding is a genuine feature of the data or an artefact of how tightly the prior is set.

3.2 Data and Sample

All series are drawn from the World Bank World Development Indicators (WDI) databased for Bangladesh. The raw annual dataset spans 1972-2024 (T = 53). It features a complete dataset for GDP growth, trade openness and gross capital formation. As for the FDI net inflows, it is populated from the mid-1990s onward. Three variables are treated endogenous for the vector autoregression (VAR): GDP growth, trade openness, and FDI net inflows. Then there are two variables which are exogenous controls: gross capital formation (% of GDP) and the annual percentage change in the official exchange rate (LCU per US$). Three calendar dummies (1991, 2003, 2013) are retained only as interpretive markers against which the estimated coefficient and volatility path read. After first-differencing trade openness and using one lag, the estimation sample is T = 52, with the constant-parameter VAR benchmark further losing one observation to lag construction (T = 51).

Why annual data. World Bank WDI coverage for Bangladesh's trade, FDI and GDP series is only reliably available at annual frequency across the full 1972-2024 window; sub-annual (quarterly) national accounts and balance-of-payments data for Bangladesh do not extend back to early post-independence and famine-era volatility of the mid-1970s is one of the episodes the analysis is built to speak to. The trade-off is explicit: annual frequency caps the usable sample at T = 52, which is precisely the small-sample setting the Bayesian shrinkage approach above is designed for, rather than an oversight to be corrected with more data that does not exist for this period.

3.3 Data Cleaning and Treatment

Unit root and stationarity properties were assessed jointly via the augmented Dickey-Fuller (ADF) test, KPSS test, and the Zivot-Andrews (ZA) test. Where, ZA test allowed for a single endogenously determined structural break. ADF and KPSS are known to lose power and can point to opposite conclusions once a series contains a break. GDP growth failed to reject a unit root under ADF (p=0.222) and rejected stationarity under KPSS (p=0.010). But Zivot-Andrew's statistic rejected the unit root null overwhelmingly (−13.637 against a 1% critical value of −5.28). This pattern was read as omitted-break bias in the standard tests rather than genuine non-stationarity. It is consistent with GDP growth being a bounded rate variable and the series was retained in levels (Adeleye et al., 2021). Trade openness failed to reject a unit root under all three tests and was first-differenced before entering the VAR. FDI net inflows produced conflicting evidence (ADF rejected a unit root at the 1% level while KPSS and Zivot-Androws did not). Given the direct ADF evidence and the theoretical prior that a bounded, mean-reverting ratio is unlikely to be integrated, FDI was retained in levels. No further outlier treatment, interpolation, or missing-value imputation was applied. The WDI extract used contains no gaps over the window ultimately modelled.

Table 1: Variable Definitions and Measurement

Role	Variable	WDI code	Treatment
Endogenous	GDP growth (annual %)	NY.GDP.MKTP.KD.ZG	Levels
Endogenous	Trade (% of GDP)	NE.TRD.GNFS.ZS	First-differenced
Endogenous	FDI, net inflows (% of GDP)	BX.KLT.DINV.WD.GD.ZS	Levels
Exogenous control	Gross capital formation (% of GDP)	NE.GDI.TOTL.ZS	Levels
Exogenous control	Exchange rate, annual % change	PA.NUS.FCRF (derived)	% change

3.4 Model Specification

The model follows Primiceri's (2005) time-varying-parameter VAR with stochastic volatility (TVP-VAR-SV). It is an extension of the fixed-coefficient VAR of Sims (1980) and the coefficient-only drift of to let coefficients, contemporaneous covariance loadings and equation-level volatilities all evolve as random walks. For the endogenous vector $y_t = (GDP\ growth_t, \Delta trade_t, FDI_t)'$ and exogenous controls x_t, the reduced-from VAR(1) at each t is

$$y_t = X_t' \beta_t + \Gamma x_t + A_t^{-1} H_t^{0.5} \epsilon_t, \ \ \epsilon \sim N(0, I), \tag{1}$$

with time-varying coefficients evolving as

$$\beta_t = \beta_{t-1} + u_t, \ \ u_t \sim N(0, Q), \tag{2}$$

time-varying log-volatilities $H_t = diag(h_{1t}, h_{2t}, h_{3t})$ evolving as

$$logh_{i,t} = logh_{i,t-1} + \eta_{i,t}, \ \ \eta_{i,t} \sim N(0, 0.15^2), \tag{3}$$

and the lower-triangular contemporaneous loadings A_t evolving as $A_t = A_{t-1} + \zeta_t$. The exogenous coefficients Γ are held constant across t, a deliberate dimension-reduction given T = 52, since a full VAR (1) in three endogenous variables already implies roughly 15-18 evolving states. Lag order was fixed at $p = 1$ by unanimous agreement of AIC, BIC, FPE and HQIC over a search capped at three lags (Table 4).

Why VAR (1) rather than VAR (2). Beyond the information-criteria result itself, two further considerations rule against a longer lag length here. First, each additional lag multiplies the number of time-varying states that the T = 52 sample has to identify. Moving from one to two lags would roughly double the size of the coefficient block. This will worsen an already heavily over-parameterised state space (the Discussion returns to this). Second, annual frequency data of this kind rarely carries genuine second-lag dynamics once a first lag and the exogenous controls are included. Since most of the transmission from trade of FDI shocks to growth plausibly operates within the same year or with a one-year lag at annual frequency; a second lag would mostly re-absorb residual first-lag persistence rather than add distinct information. VAR(2) was not estimated as robustness check in this round of revisions. It is flagged here as a natural extension for a future, higher-frequency version of this study rather than attempted at T = 52.

3.5 Estimation Technique

The model was estimated by Bayesian Markov chain Monte Carlo using Hamiltonian Monte Carlo with the No-U-Turn Sampler (NUTS) in PyMC. NUTS samples the full sequence of latent states jointly rather than marginalizing them through a Kalman filter. Priors follow a Minnesota-style shrinkage scheme calibrated on a 1972-1981 training sample: $\beta_0 \sim N\left(\hat{\beta}_{OLS}, 4\,\widehat{VAR}\left(\hat{\beta}_{OLS}\right)\right)$, with innovation variance Q fixed tight relative to this training-sample variance given the short series; $\log h_{i0} \sim N\left(\log \hat{\sigma}_i^2, 4\right)$;; and A_0off-diagonal elements $N(0, 1)$. Posterior inference is reported via medians and 90% credible intervals rather than point estimates with classical standard errors, the natural counterpart to Bayesian estimation.

The Minnesota prior in plain terms. A Minnesota prior is a way of telling the model, before it sees the data, what a "sensible" coefficient looks like in a VAR. A variable's own lag should matter more than another variable's lag, and coefficients on more distant lags should matter less than coefficients on recent ones. Concretely here, the prior centres each coefficient's starting value on its ordinary-least-squares estimates from a short 1972-1981 training window. Then only allows the coefficient to drift away from that starting point slowly, by constraining the innovation variance to be small. This does two things a reader unfamiliar with Bayesian methods may not expect. It does not force the coefficient to stay perfectly constant (the model is still free to detect real movement if the likelihood strongly supports it). But it does mean that, in a short sample like this one, an essentially flat estimated path is the honest default outcome whenever the data does not contain strong evidence against it, rather than evidence that estimation has failed to find something that is there.

3.6 Diagnostic and Robustness Checks

Prior to the main model, a constant-parameter VAR (1) benchmark was estimated as an explicit comparator. And its residuals were tested for ARCH effects. The resulting p-values (0.162, 0.496, 0.052 for the GDP growth, trade and FDI equations respectively) gave only modest FDI-concentrated evidence of conditional heteroskedasticity. So, the stochastic-volatility extension is motivated primarily on theoretical grounds rather than an overwhelming ARCH-LM signal. Structural instability was checked independently of the calendar regime markers using the Pruned Exact Linear Time (PELT) algorithm with a radial basis function cost on the standardized three-variable signal. A single break at 2002 was detected stably across penalty values of 2, 3 and 5, one year ahead of the 2003 exchange-rate float marker. Neither 1991 nor 2013 was recovered as a break in the joint second-moment structure. This finding has been reported rather than reconciled with the calendar narrative. As a robustness check on the headline coefficient paths, the model was re-estimated with the coefficient random-walk innovation standard deviation loosened three-fold (from 0.05 to 0.15). It holds the frozen covariance loadings and all other priors fixed. Under both specifications the 90% credible interval for the trade- and FDI-growth coefficients excludes zero in no sample year. So, the qualified-null conclusion is robust to this prior choice even though the precise magnitude of apparent drift is not (Mendoza et al., 2026).

3.8 Computational Details

Convergence was assessed against a pre-registered gate of $\hat{R} < 1.01$ and effective sample size (ESS) comfortably above a few hundred draws as per the current practice (Roy, 2019). Divergent transitions monitored as an additional check not available under classical filtering estimation (Vats & Knudson, 2021). The full Primiceri covariance specification, with time-varying contemporaneous loadings, failed this gate (maximum $\hat{R} = 1.029$, minimum ESS 111 of 4,000 post-warmup draws, despite zero divergences). Per the pre-registered fallback, the contemporaneous covariance loadings a_{21}, a_{31}, a_{32} were re-estimated as time-invariant constants while coefficients and log-volatilities remained fully time-varying. This fallback specification, run over 4 chains of 500 tuning and 400 post-warmup draws each (1,600 total post-warmup draws, target acceptance 0.95, max tree depth 12, zero divergences). It converged comfortably on the parameters of substantive interest (maximum $\hat{R} = 1.009$for the 612 time-varying coefficients, 1.001 for the constant exogenous coefficients, 1.003 for the frozen loadings), with a single isolated soft spot in the FDI equation's volatility path (maximum $\hat{R} = 1.021$, minimum ESS 293). This has been noted as a limitation rather than suppressed.

What freezing the covariance loadings implies for inference. Freezing the contemporaneous loadings to a constant means the model no longer allows the contemporaneous, within-year relationships among GDP growth, trade and FDI shocks to change over time. Those relationships are estimated once, from the whole sample, and held fixed. This has a specific, limited scope: it does not bias or otherwise compromise the time-varying coefficient path or the volatility path reported as the paper's headline results. Because those are estimated conditional on, not through, the covariance structure and both were shown to converge cleanly under the fallback. What it does restrict is any claim about how the transmission of a shock in one equation to the others has itself evolved. For example, whether an FDI shock has more or less tightly linked to a contemporaneous trade shock over time; that specific question is left open by this specification and is flagged as a direction for future work with a longer or higher-frequency series, where the full covariance block is more likely to be identified.

3.9 Software and Tools

All analysis was conducted in Python, using pandas and openpyxl for data preparation, statsmodels and arch for unit-root testing, lag selection and the constant-parameter VAR benchmark, ruptures for PELT structural-break detection, and PyMC with ArviZ for TVP-VAR-SV estimation, posterior summarization and convergence diagnostics.

3.10 Benchmarking Against the Literature

The design extends the Bangladesh-specific FDI-trade-growth literature, which has remained within a constant-parameter Johansen/VECM or ARDL bounds-testing tradition despite documented regime change (Adhikary, 2010; Haque & Amin, 2018; Ridwan et al., 2025; Saleem et al., 2020). By applying the Primiceri (2005) TVP-VAR-SV framework, previously unapplied in context of Bangladesh, estimated via modern NUTS-basesd HMC rather than the original Carter-Kohn/KSN algorithm. The chosen design accepts a parsimony-flexibility trading-off characteristic of this model class in short annual samples: exogenous coefficients and following the convergence gate. The contemporaneous covariance structure is held constant to keep the state space estimable at T = 52, a documented simplification rather than a retreat from the time-varying-parameter approach (Chan & Eisenstat, 2015). Consistent with theoretical work showing that a flat coefficient path under a time-varying specification is itself a legitimate structural finding rather than a failure of the exercise (Canova et al., 2015). The resulting stable coefficient paths, checked against a three-fold looser prior, are reported alongside clearly time-varying stochastic volatilities as two separate and equally reportable results.

4. Results

4.1 Descriptive statistics and unit root properties

Table 2 summarises the three endogenous series over 1972-2024. Average GDP growth over the period is 4.65% (s.d. 3.33). It has been depressed by the severe post-independence contraction of –13.97% recorded in 1972. Trade openness rose from a low base to average 27.43% of GDP (s.d. 9.24). This reflects Bangladesh's gradual liberalisation. While FDI net inflows averaged just 0.39% of GDP (s.d. 0.48). It remained close to zero until the mid-1990s (Figure 1).

Table 3 reports ADF, KPSS and Zivot-Andrews unit root tests. GDP growth technically fails to reject a unit root under ADF ($p = 0.222$) and rejects stationarity under KPSS ($p = 0.010$). The Zivot-Andrews statistic, which allows for a single structural break, rejects the unit root null overwhelmingly instead, a pattern consistent with omitted-break bias in the standard tests rather than genuine non-stationarity, and with GDP growth being a bounded rate variable. It is therefore retained in levels. Trade openness fails to reject a unit root under all three tests including Zivot-Andrews with a break allowed. It is thus first-differenced before entering the VAR. FDI net inflows show conflicting evidence (ADF rejects a unit root at the 1% level, $p = 0.008$, while KPSS and Zivot-Andrews do not). Given the direct ADF evidence and the theoretical prior that FDI ratios are bounded and mean-reverting, FDI is retained in levels. The resulting endogenous vector is (GDP growth, Δtrade openness, FDI ratio), T = 52 (1973-2024).

Table 2: Descriptive statistics, 1972-2024

	mean	std	min	max
GDP growth (%)	4.65	3.33	−13.97	9.59
Trade (% of GDP)	27.43	9.24	11.00	48.11
FDI net inflows (% of GDP)	0.39	0.48	−0.05	1.74

Source: World Bank World Development Indicators, Bangladesh. All series in levels, annual, 1972-2024 (T = 53)

Table 3: Unit root tests

Variable	ADF stat	ADF p	KPSS stat	KPSS p	ZA stat	ZA break year
GDP growth (%)	−2.158	0.222	1.094	0.010	−13.673	2005
Trade (% of GDP)	−1.565	0.501	0.808	0.010	−3.919	2016
FDI net inflows (% of GDP)	−3.486	0.008	0.714	0.012	−3.617	2016

ADF: augmented Dickey-Fuller (lag length by AIC). KPSS: Kwiatkowski-Phillips-Schmidt-Shin (constant, bandwidth). Zivot-Andrews (ZA): unit root test allowing one endogenously-determined structural break in the constant (lag search capped at 4 given T = 53)

4.2 Lag selection and constant-parameter VAR benchmark

Information criteria unanimously select a lag order of $p = 1$ (Table 4) which is consistent with the short annual sample. The constant-parameter VAR(1) benchmark (Table 5) shows GDP growth mean-reverting toward its own recent history, FDI persisting strongly year to year, and trade-openness change responding negatively to lagged growth, the full coefficients and significance levels are in Table 5 rather than repeated her. ARCH-LM tests on the benchmark residuals give $p = 0.162$ (GDP growth) and $p = 0.052$ (FDI). There is a modest, borderline evidence of conditional heteroskedasticity concentrated in the FDI equation, rather than a uniformly strong case for stochastic volatility across all three equations. The stochastic-volatility extension below is accordingly motivated primarily on theoretical grounds (Bangladesh's documented liberalisation-era volatility regimes) rather than on an overwhelming ARCH-LM signal.

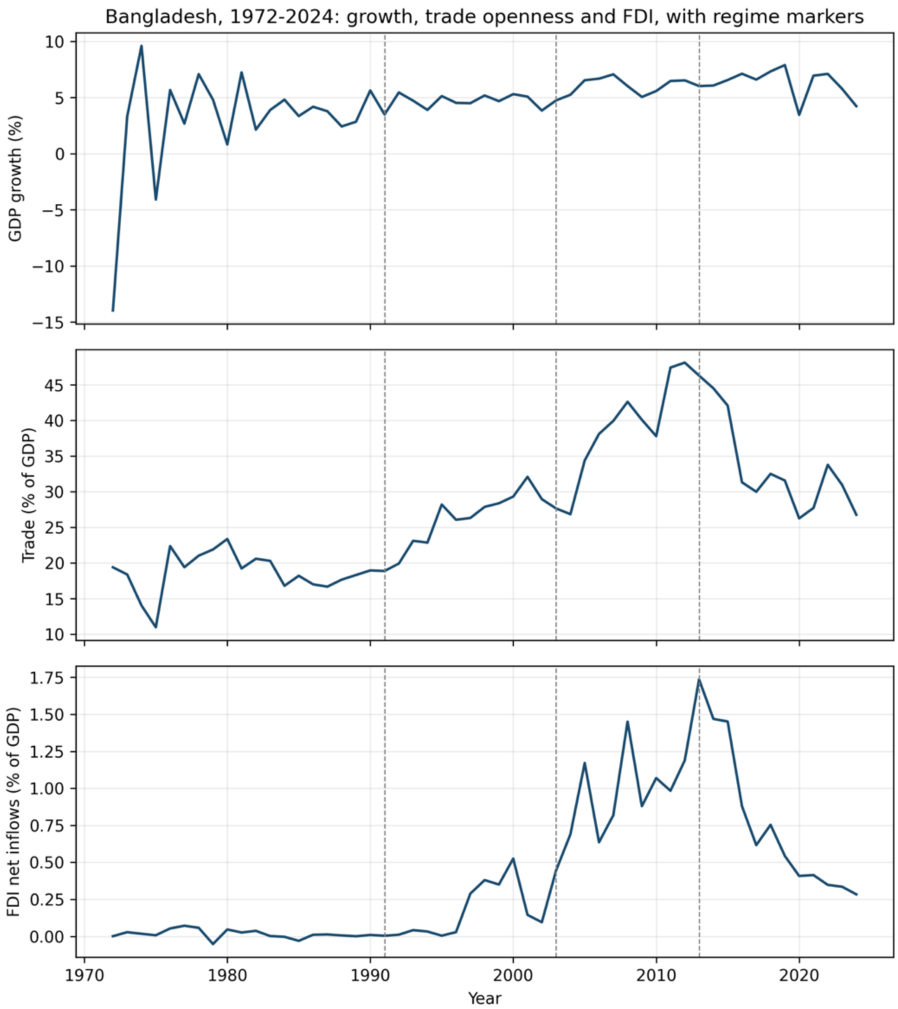

Figure 1: *Bangladesh, 1972-2024: GDP growth, trade openness and FDI net inflows, with 1991/2003/2013 regime markers.*

Table 4: Lag order selection

	AIC	BIC	FPE	HQIC
0	1.021	1.369	2.778	1.153
1	0.2614*	0.9563*	1.304*	0.5250*
2	0.4128	1.455	1.530	0.8083
3	0.6274	2.017	1.931	1.155

Asterisk denotes the order selected by each criterion. Maximum lag searched: 3 (annual data, T = 52)

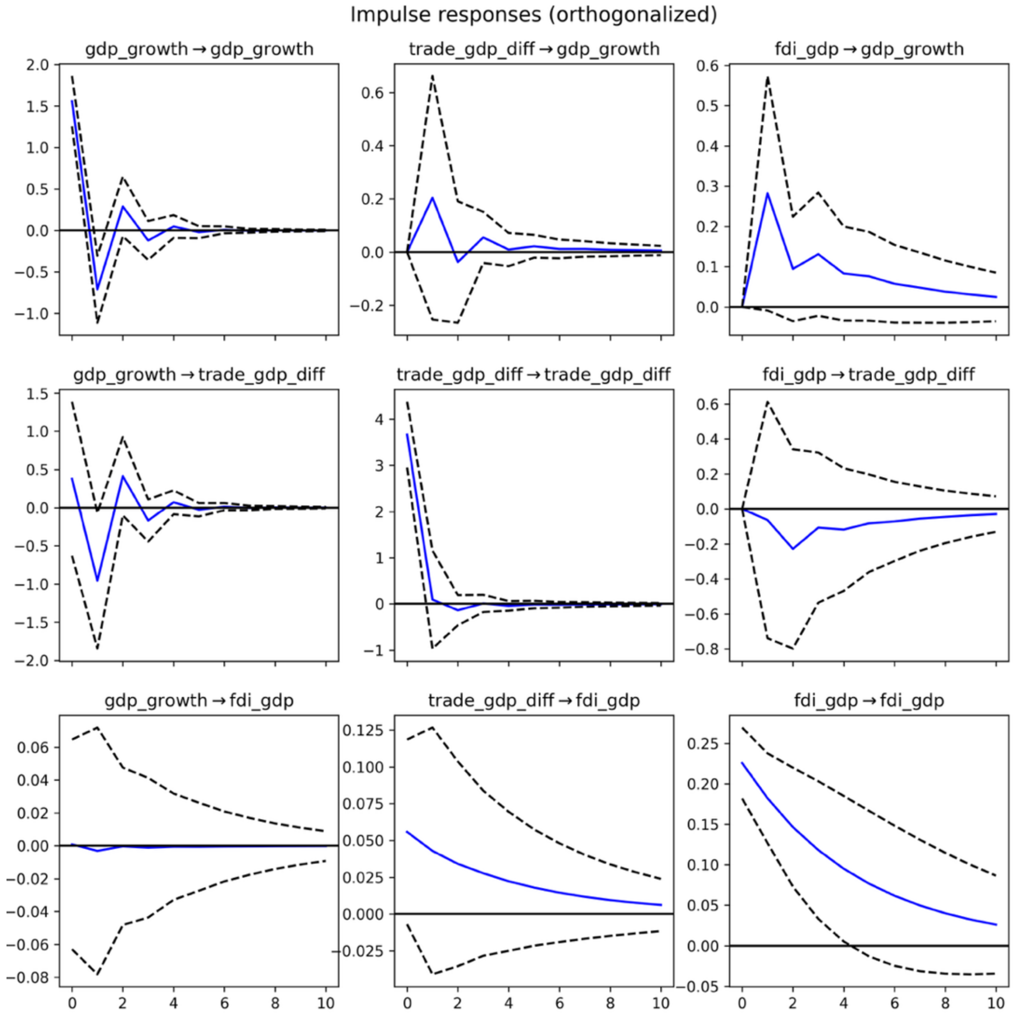

Figure 2: Orthogonalised impulse responses, constant-parameter VAR(1) benchmark

Table 5: Constant-parameter VAR(1) benchmark

Regressor	GDP growth	ΔTrade (% GDP)	FDI (% GDP)
const	4.754***	1.107	−0.036
gcf_gdp	0.114**	0.086	0.007
Fx_pct_change	−6.457***	7.285	−0.225
L1.gdp_growth	−0.464***	−0.621**	−0.002
L1.trade_gdp_diff	0.037	0.030	−0.001
L1.fdi_gdp	1.252*	−0.280	0.807***
Observations		51	
Log likelihood		−220.45	

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Exogenous regressors (gross capital formation %GDP, FX % change) enter with constant coefficients throughout, including in the TVP-VAR-SV below.

4.3 Structural break detection

An empirical, data-driven check on the 1991/2003/2013 regime narrative is provided by ruptures' PELT algorithm with an RBF cost function on the standardized signal used here. A single break at 2002 is detected stably across penalty values 2, 3 and 5 (Table 6, Figure 3), one year ahead of the 2003 exchange-rate float marker used to motivate the regime dummies. Neither 1991 nor 2013 recover as a break in the joint second-moment structure of the three series under this method. This is reported as a finding rather than smoothed over and suggests the 1991 and 2013 dates may matter more as policy-dated thresholds than as detectable shifts in this particular multivariate signal.

Table 6: Detected structural breaks, sensitivity to PELT penalty

Penalty	Detected break years
1	[1977, 2002, 2012, 2017]
2	[2002]
3	[2002]
5	[2002]
8	[]

PELT with RBF cost function on standardized (GDP growth, Δtrade, FDI) signal, minimum segment length 5.

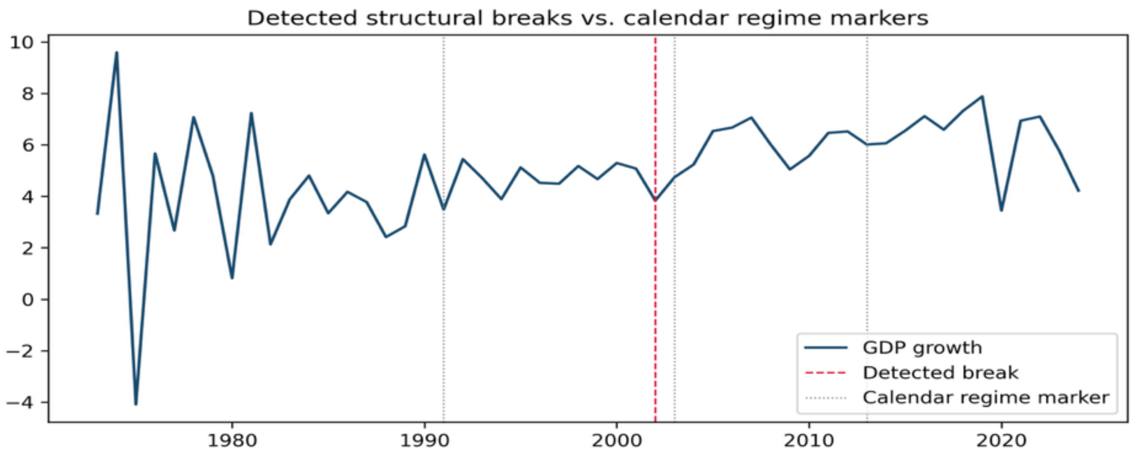

Figure 3: Detected structural break (2002) vs. calendar regime markers (1991, 2003, 2013)

4.4 TVP-VAR-SV estimation and convergence

The full Primiceri (2005) style specification: time-varying VAR(1) coefficients, time-varying stochastic volatility, and time-varying contemporaneous covariance loadings, was estimated via Hamiltonian Monte Carlo (NUTS, PyMC). The full specification's time-varying covariance loadings failed the convergence gate (maximum $\hat{R} = 1.029$, effective sample size as low as 111 out of 4,000 post-warmup draws, despite zero divergences). While the time-varying coefficients and volatilities passed comfortably. Per the pre-registered fallback path, the contemporaneous covariance loadings were re-estimated as time-invariant constants. The coefficients and stochastic volatilities remain fully time-varying. This is a documented simplification of the full Primiceri covariance structure, not a retreat from the time-varying-parameter model itself (see the Methodology section above for what this restricts and does not restrict about the inference that follows).

Table 7: MCMC convergence diagnostics, fallback specification

Parameter group	N params	Max R-hat	Mean R-hat	Min ESS (bulk)	Median ESS (bulk)	Min ESS (tail)
beta (time-varying var coeff.)	612	1.0090	1.0032	532	763	512
gamma (constant exog. coeff.)	6	1.0012	1.0000	930	1608	964

log_h (log stochastic volatility)	153	1.0211	1.0056	293	1403	478
a21_const (covar. loading)	1	1.0012	1.0012	1405	1405	1508
a31_const (covar. loading)	1	1.0005	1.005	1272	1272	1595
a32_const (covar. loading)	1	1.0033	1.0033	1879	1879	1644

4 chains, 500 tuning + 400 post-warmup draws each (1,600 total post-warmup draws), NUTS target acceptance 0.95, max. tree depth 12. Zero divergences across all four chains. Exogenous regressors standardized prior to estimation.

Convergence is comfortable for the parameters the headline results depend on: the time-varying coefficient array (maximum $\hat{R} = 1.009$ across 612 individual parameters), the constant exogenous coefficients (maximum $\hat{R} = 1.001$), and the frozen covariance loadings (maximum $\hat{R} = 1.003$). A single, isolated soft spot remains in the stochastic-volatility path for the FDI equation specifically (maximum $\hat{R} = 1.021$, minimum ESS 293). This plausibly reflects the genuinely large historical swings in FDI's variance (near-zero for two decades, then repeated spikes); this is noted as a limitation rather than suppressed. Trace plots are provided below (Figure 4).

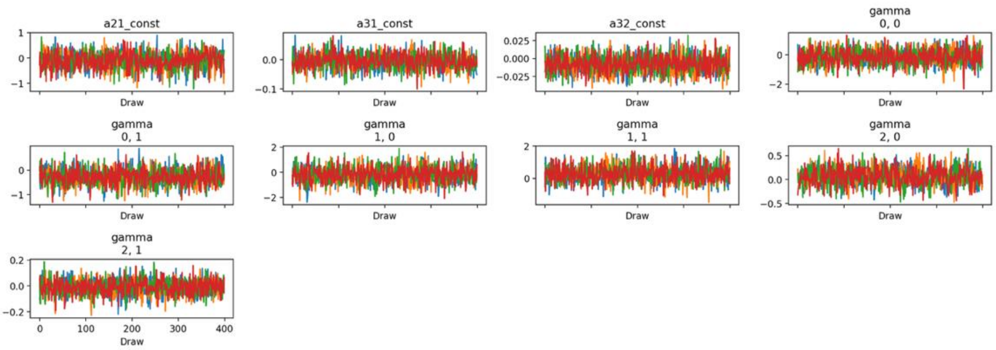

Figure 4: MCMC trace plots, constant exogenous coefficients

4.5 Headline results: time-varying coefficient paths

Figures 5 and 6 show the posterior median and 90% credible interval for the time-varying effect of Δtrade openness and of the FDI ratio on GDP growth, respectively, over 1974-2024. Table 8 reports the full year-by-year posterior summary.

Both paths are, substantively, flat: the trade-openness coefficient ranges only from 0.019 to 0.020 across the entire sample. The FDI coefficient ranges from 0.578 to 0.618, against a shrinkage prior whose cumulative standard deviation over the 51-year sample is only about 0.36. In neither case does the 90% credible interval exclude zero in any year. Which means that in every single year of the sample, a zero elasticity remains among the range of values the posterior treats as plausible for that coefficient. The model cannot rule out a genuinely zero effect of either trade openness or FDI on growth at any point between 1974 and 2024. At T = 52, that is the state space telling us something real about the limits of what an annual, five-decade sample can pin down, not an artefact to be explained away. The FDI coefficient's point estimate is nonetheless consistently positive and economically non-trivial (≈ 0.6: a one-percentage-point rise in the FDI/GDP ratio is associated with roughly 0.6 percentage points of additional GDP growth). Even though it is not statistically distinguishable from zero at conventional levels throughout the sample. The trade-openness coefficient, by contrast, hugs zero throughout, with no evidence of either an economically meaningful effect or of time-variation in it.

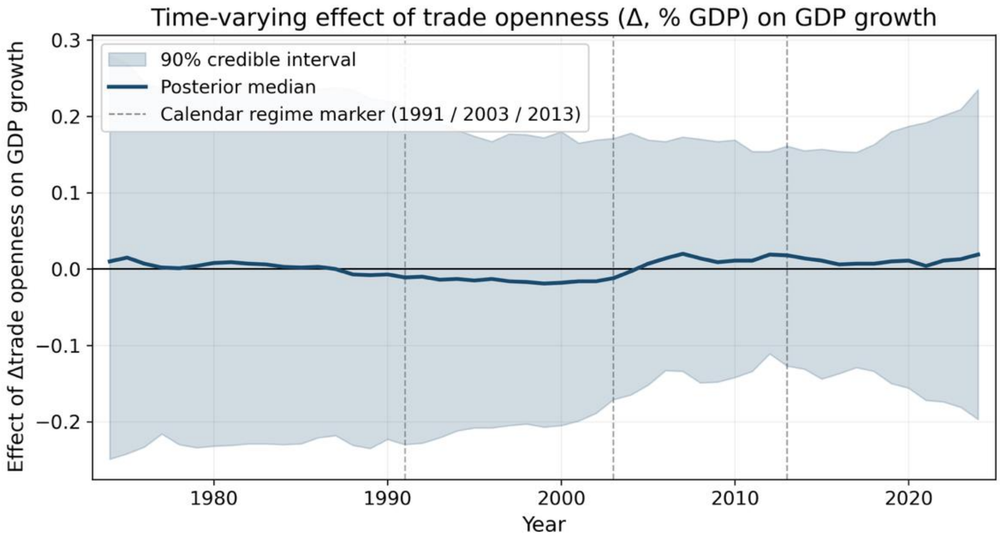

Figure 5: Time-varying effect of trade openness (Δ, % GDP) on GDP growth, posterior median and 90% credible interval.

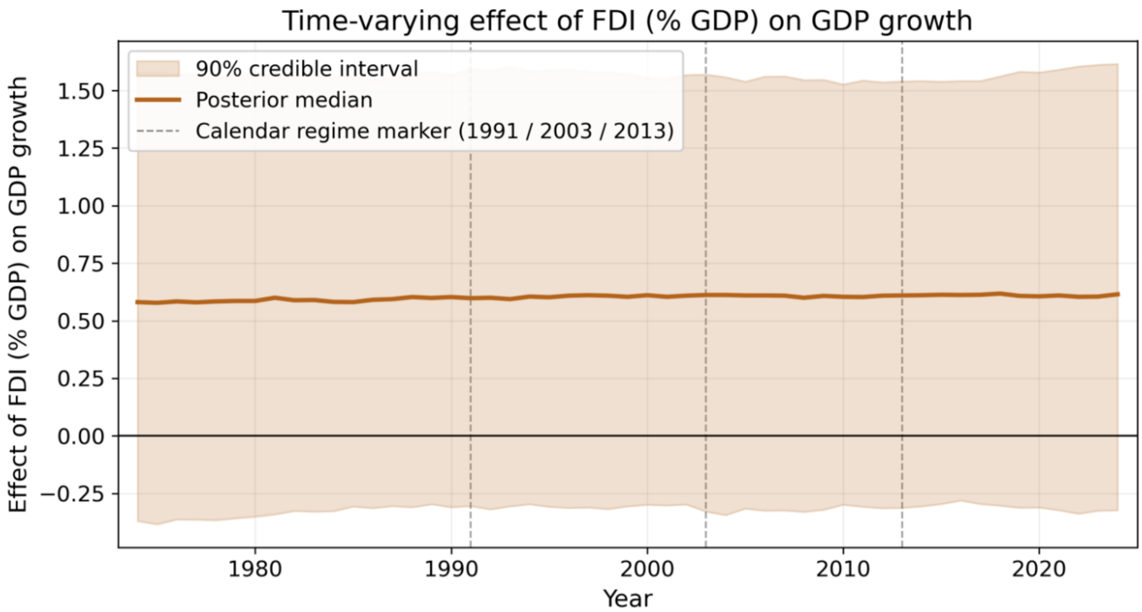

Figure 6: Time-varying effect of FDI (%GDP) on GDP growth, posterior median and 90% credible interval.

Table 8: TVP-VAR-SV coefficient summary, selected years

Year	Trade (med)	Trade (5%)	Trade (95%)	FDI (med)	FDI (5%)	FDI (95%)
1974	0.010	−0.249	0.280	0.581	−0.370	1.577
1975	0.015	−0.242	0.269	0.578	−0.384	1.579
1976	0.007	−0.233	0.244	0.584	−0.363	1.570
1977	0.002	−0.216	0.236	0.580	−0.364	1.570
1978	0.001	−0.230	0.241	0.584	−0.366	1.555
1979	0.004	−0.234	0.238	0.586	−0.358	1.566
1980	0.008	−0.232	0.245	0.586	−0.351	1.577

1981	0.009	−0.231	0.241	0.600	−0.341	1.568
1982	0.007	−0.229	0.243	0.589	−0.326	1.564
1983	0.006	−0.229	0.238	0.590	−0.329	1.567
1984	0.003	−0.230	0.235	0.582	−0.327	1.565
1985	0.002	−0.229	0.238	0.581	−0.307	1.560
1986	0.003	−0.221	0.235	0.591	−0.314	1.546
1987	0.000	−0.218	0.238	0.594	−0.305	1.553
1988	−0.007	−0.231	0.234	0.603	−0.310	1.574
1989	−0.008	−0.235	0.223	0.599	−0.297	1.580
1990	−0.007	−0.223	0.220	0.603	−0.310	1.567
1991	−0.011	−0.230	0.213	0.598	−0.305	1.593
1992	−0.010	−0.228	0.199	0.600	−0.319	1.585
1993	−0.014	−0.221	0.191	0.594	−0.306	1.602
1994	−0.013	−0.212	0.182	0.605	−0.297	1.583
1995	−0.015	−0.208	0.174	0.602	−0.308	1.587
1996	−0.013	−0.208	0.167	0.609	−0.313	1.592
1997	−0.016	−0.205	0.177	0.611	−0.311	1.580
1998	−0.017	−0.203	0.176	0.609	−0.318	1.578
1999	−0.019	−0.207	0.172	0.604	−0.306	1.571
2000	−0.018	−0.205	0.180	0.611	−0.299	1.555
2001	−0.016	−0.199	0.165	0.604	−0.302	1.549
2002	−0.016	−0.189	0.169	0.609	−0.298	1.568
2003	−0.012	−0.171	0.171	0.612	−0.329	1.570
2004	−0.003	−0.165	0.178	0.612	−0.344	1.558
2005	0.007	−0.152	0.169	0.610	−0.316	1.539
2006	0.014	−0.133	0.167	0.610	−0.325	1.561
2007	0.020	−0.134	0.173	0.609	−0.324	1.562
2008	0.014	−0.149	0.170	0.600	−0.330	1.546
2009	0.009	−0.148	0.167	0.608	−0.321	1.547
2010	0.011	−0.142	0.169	0.604	−0.299	1.527
2011	0.011	−0.134	0.154	0.603	−0.308	1.544
2012	0.019	−0.111	0.154	0.609	−0.314	1.536
2013	0.018	−0.127	0.161	0.610	−0.313	1.540
2014	0.014	−0.131	0.155	0.611	−0.306	1.542
2015	0.011	−0.144	0.157	0.613	−0.296	1.539
2016	0.006	−0.137	0.154	0.612	−0.281	1.542
2017	0.007	−0.129	0.153	0.613	−0.296	1.542
2018	0.007	−0.134	0.163	0.618	−0.303	1.562
2019	0.010	−0.150	0.180	0.608	−0.312	1.582
2020	0.011	−0.156	0.187	0.606	−0.311	1.579
2021	0.004	−0.172	0.192	0.610	−0.323	1.591
2022	0.011	−0.174	0.201	0.604	−0.338	1.605
2023	0.013	−0.181	0.209	0.605	−0.325	1.612
2024	0.019	−0.197	0.235	0.615	−0.323	1.616

Posterior median and 90% credible interval (5th/95th percentile), fall-back specification.

4.6 Robustness: sensitivity to the coefficient innovation-variance prior

The flat coefficient paths above are, by construction, partly a function of the tightness of the shrinkage prior on the coefficients' random-walk step size ($\sigma_\beta = 0.05$). Thus, a robustness check re-estimates the model with a three-fold looser prior ($\sigma_\beta = 0.15$), holding the frozen covariance loadings and all other priors fixed (Table 8). Under the looser prior, the posterior-median coefficient paths show five-to-six-fold more within-sample variation (trade-coefficient range widens from 0.039 to 0.219; FDI-coefficient range widens from 0.041 to 0.258). As for the trade-openness coefficient's point estimate dips most negative (−0.11) in the immediate vicinity of the empirically detected 2002 break. It recovers to a small positive value by 2013, a shape broadly consistent with the regime-shift narrative, though not one the data can pin down precisely. Critically, however, the credible interval excludes zero in no year under either prior specification: the conclusion that neither effect is statistically distinguishable from zero at any point in 1974-2024 is robust to this prior choice, even though the precise magnitude of apparent coefficient drift is not.

Table 9: Robustness check: coefficient innovation-variance prior

Spec.	Trade range	Trade span	Trade sig.	FDI range	FDI span	FDI sig.
Main	[−0.019, 0.020]	0.039000	0	[0.578, 0.618]	0.04100	0
Robust (3x)	[−0.127, 0.093]	0.219000	0	[0.376, 0.634]	0.258000	0

"Sig." counts the number of sample years (out of 51) in which the 90% credible interval for the respective coefficient excludes zero. σ_β: prior standard deviation of the coefficients' random-walk step size (main specification 0.05, robustness check 0.15)

4.7 Stochastic volatility

Unlike the coefficient paths, the estimated stochastic volatilities move a great deal, and the movement lines up with identifiable episodes in Bangladesh's economic history rather than looking like noise (Figure 7). GDP growth volatility peaks in 1974, coinciding with the immediate post-independence and famine-era turmoil visible in Figure 1, and troughs in 2004, a period of comparatively stable growth. Trade-openness change volatility is lowest in 1990 and rises substantially to a peak in 2016. Tracking the RMG export boom's growing exposure to global demand swings. FDI volatility, an order of magnitude smaller in absolute terms, is highest in 1974 and lowest in 2022. It is consistent with FDI flows moving from near-negligible and erratic in the early sample to a larger but comparatively more settled base in the 2020s. These patterns give a genuine empirical case for the stochastic-volatility extension of the model, independent of the inconclusive evidence of time-varying coefficients above.

The constant-parameter VAR benchmark and the data-driven structural break in 2002 together motivate a time-varying specification. The model estimated in this study, Primiceri (2005)-style TVP-VAR-SV, with the covariance loadings frozen to constants following a documented convergence-drive fallback. It converges cleanly on the parameters of substantive interest. The central finding, however, is a qualified null. Neither the trade-openness nor the FDI effect on GDP growth is estimated to vary significantly over 1974-2024, nor is effect statistically distinguishable from zero at any point in the sample, a conclusion that is robust to a three-fold loosening of the coefficient shrinkage prior. This is informative rather than a failure of the exercise. It suggests that, once estimation uncertainty is properly accounted for at annual frequence with T = 52, Bangladesh's growth-trade and growth-FDI elasticities are more plausibly characterized as stable than as episodically shifting with the 1991/2003/2013 policy regime changes. The clear time-variation recovered in the stochastic volatilities, particularly the 1974 peak in GDP growth volatility, indicates that regime change in this data is better captured as a shift in the variance of shocks than as a shift in the transmission coefficients themselves.

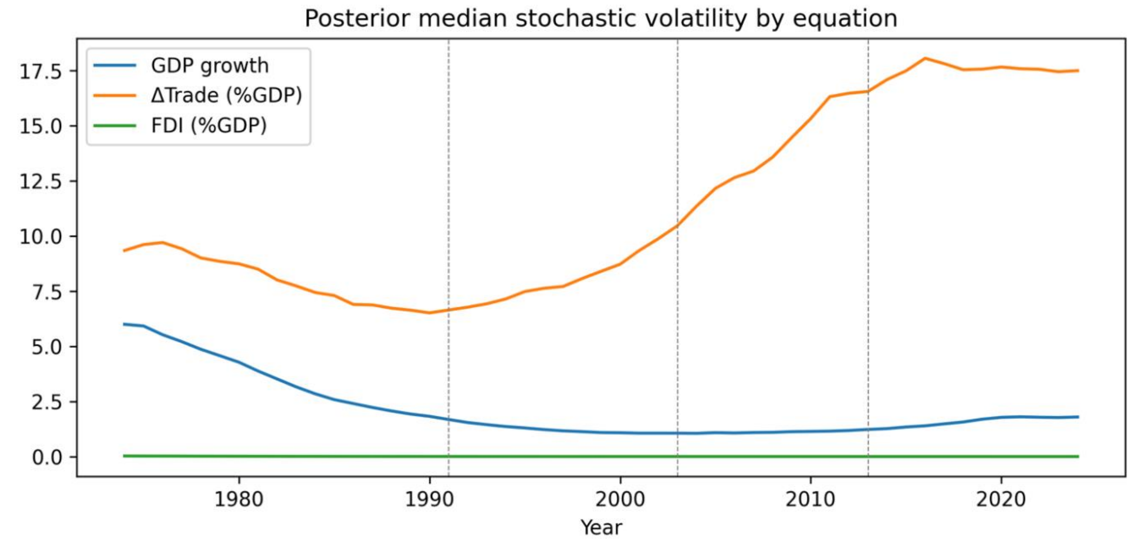

Figure 7: Posterior median stochastic volatility by equation, 1974-2024

5 Discussion

5.1 Reframing the hypotheses

The design was built to detect a strengthening trade elasticity after the 2003 float and a weakening FDI elasticity after the 2013 RMG compliance-cost shock. Neither is supported. From 1974 to 2024 both coefficient paths stay close to flat and zero remains inside the 90% credible interval in every single year, meaning the model cannot rule out a genuinely null effect for either variable at any point in the sample.

The qualified null on trade and FDI elasticities and why volatility absorbs the movement instead

Chan and Eisenstat (2018) found for the US and Australia that most of the gain a full time-varying VAR offers over a constant-parameter one comes from stochastic volatility rather than coefficient drift, with Bayes factors often favouring a constant coefficient VAR with stochastic volatility over the fuller specification. The same pattern appears here: coefficients that barely move, volatility that moves substantially. Terrell et al. (2026) report a similar split across six small open economies, and Malik and Sah (2024) find FDI's growth effect stayed essentially flat across the BRICS economies over two decades. This pattern reflects the Sims–Stock critique: in a constant-variance VAR, unmodelled shifts in shock size are forced into estimated coefficients, creating illusory parameter drift.

Canova et al. (2015) make the complementary theoretical point that this is not a disappointing outcome. A flat coefficient path is a genuine finding about the transmission mechanism, not a wasted modelling exercise, because the structural dynamics implied by a time-varying model, impulse included, are the same as those implied by a fixed-coefficient model whenever the underlying parameter variation is exogenous to the system.

Mendoza et al. (2026) adds a robustness check this study relies on and that is if you loosen the prior on coefficient innovation variance several times over and the paths still stay flat, that actually strengthens the stability finding because it rules out the tight prior as the reason. That's exactly what happens here, the credible interval excludes zero in no year, whether under the main prior (standard deviation 0.05) or the loosened one (0.15). Umar and Aldio (2026) looking at Indonesia, does find genuine coefficient movement around the global financial crisis and COVID-19 but on a quarterly series with well over a hundred observations. At T = 52 and annual frequency, Bangladesh's model has a fraction of that information to work with, which could just as easily explain the difference as anything about the two economies.

5.2 Volatility as the locus of regime change

The 1974 peak in GDP growth volatility lines up with the immediate post-independence period and the 1974 famine, not with any of the three regime markers the study was originally built around. Plakandaras et al. (2018) find something similar in their long-run UK data like volatility spikes around political turmoil and the collapse of a fixed exchange rate regime not around gradual policy shifts and that description fits 1974 in Bangladesh better than it fits 1991.

From its 1990 low to its 2016 peak the rise in trade openness volatility is a different kind of movement as it was gradual rather than sudden. It fits with Bangladesh's exports becoming more exposed to global demand swings as trade's share of GDP grew through liberalisation and the RMG boom. FDI volatility follows a related but smaller scale pattern. High in 1974 and lowest in 2022, reflecting decades of small, unremarkable FDI flows followed by a period of larger, choppier inflows as EPZ driven investment picked up.

5.3 Domain-specific mechanisms behind the null

Two features specific to Bangladesh make a flat elasticity easier to believe than a shifting one. First, the FDI coefficient's point estimate sits around 0.6 throughout the sample. It is positive and not small, but it is never statistically distinguishable from zero. FDI net inflows averaged just 0.39% of GDP over the period and staying near zero until the mid-1990s. With that little variation in the regressor itself, the model may simply not have enough to work with to detect a moving coefficient, whatever the true elasticity is doing. On top of that, most of this FDI has gone into energy and export processing zones. Sectors with limited links back to the rest of the economy which works against the kind of economy wide spillover that would make an elasticity shift alongside policy.

Second, the trade coefficient's flat, near-zero path has a more structural explanation. Bangladesh's exports are dominated by ready-made garments, and the sector relies heavily on back-to-back letters of credit, where export earnings are largely cancelled out within the same transaction by imported fabric, yarn, and accessories. Huchet-Bourdon et al. (2018) make the broader point that gross trade volume is a weak proxy for growth impact when a country's exports carry a lot of imported content. That is basically a description of the RMG model.

First, on export diversification Gnangnon (2019) finds across a 20-year cross-country panel that multilateral trade liberalization raises export product diversification more strongly in less-developed economies than in advanced ones and links greater diversification to lower volatility of the growth rate itself which speaks directly to both halves of this paper's finding, a flat elasticity and a volatile shock variance. Second on investment quality, the discussion above suggests FDI's weak, if positive, growth linkage in Bangladesh owes something to its concentration in enclave-style energy and EPZ projects investment promotion policy that screens for backward linkages and local sourcing requirements rather than headline inflow volume would target the actual mechanism this paper identifies as weak. Third, on industrial upgrading, within the RMG sector specifically, moving beyond cut-make-trim contracting toward higher-value design, fabric sourcing and branding activity would reduce the back-to-back import content that currently caps trade's domestic value added, the structural reason the trade coefficient sits near zero. Fourth, on innovation and absorptive capacity, the volatility findings, not the elasticity findings are where shock-absorption policy has purchased, exchange-rate and trade finance buffers sized to the trade and FDI volatility regimes identified here, rather than to the assumption of a shifting transmission mechanism the calendar narrative implied. None of these four levers requires believing the growth elasticity of trade or FDI has moved or will move. Each targets the composition and resilience of trade and investment rather than their volume.

5.4 Methodological and data caveats

A null finding is only as credible as the estimation behind it, and this one didn't converge cleanly on the first pass. The full Primiceri specification with time-varying contemporaneous covariance loadings failed the pre-registered convergence gate which is a maximum R-hat of 1.029 and an effective sample size as low as 111, even though the coefficients and volatilities converged fine on their own.

Data treatment choices also carry trade-offs. Keeping GDP growth and FDI in levels aligns with Perron (1988) and bounded-ratio logic, though residual non-stationarity in FDI may contribute to the convergence soft spot in its volatility path.

Taken together these findings all point the same way. The proposal's calendar-based hypotheses don't survive contact with the data as coefficient shifts but they are not simply wrong either. 2002 shows up as a break in the joint volatility structure which is close enough to the 2003 float to be recognisable and 1974 and the mid 2010s show up as distinct volatility regimes tied to identifiable moments in Bangladesh's history. What this study adds to the existing literature on Bangladesh is not a rejection of the idea that things changed over the past five decades. It is a more precise account of where that change lives, in the size of the shocks hitting trade and FDI not in the mechanism that turns them into growth.

6. Conclusion

This study asked whether the single, constant elasticity that every existing Bangladesh-specific study of the FDI-trade-growth nexus has assumed is defensible. By letting the coefficients linking trade openness and FDI to GDP growth evolve as latent processes, jointly with shock volatility, within a Primiceri-style TVP-VAR-SV estimated over 1972-2024. The answer is not the one the design anticipated. But it is no less informative. The growth elasticities of trade openness and FDI show no statistically detectable time variation anywhere in the sample, while the volatility of the shocks themselves moves substantially and in economically legible ways.

The trade-openness coefficient has a narrow band around zero throughout. And the FDI coefficient, though consistently positive and economically non-trivial at roughly 0.6, is never distinguishable from zero at conventional credibility levels in any year. This flat-coefficient result holds under a three-fold loosening of the coefficient shrinkage prior. This strengthens rather than weakens confidence that the finding reflects the information content of a $T = 52$ annual sample rather than an artefact of prior tightness. Volatility tells a different story: GDP growth volatility peaks in 1974, near the post-independence contraction, and troughs in 2004; trade-openness volatility rises from a 1990 trough to a 2016 peak; FDI volatility is highest in 1974 and lowest in 2022. A data-driven structural break search corroborates this reading. It locates a single stable break in 2002, one year ahead of the exchange-rate float while finding nothing at 1991 or 2013.

Bangladesh’s growth-transmission mechanism therefore looks more stable than its policy history might suggest. The regime shifts long treated as watersheds for the trade-FDI-growth relationship are better understood as episodes of shifting shock variance than as changes in the elasticities themselves. This reframes rather than contradicts the existing literature, which remains correct about the level of trade and FDI ratios rising after 1991. Whether the coefficient linking them to growth moved is a separate question. The cross-country evidence that most of the gain from a fully time-varying VAR comes from volatility rather than coefficient drift suggests Bangladesh is not an outlier in that respect.

For policy, a flat and statistically imprecise trade elasticity implies that expanding trade volume through new market-access agreements is unlikely, on its own, to shift the growth relationship. In particular as LDC graduation removes preferential access. The concrete levers set out in the Discussion, export diversification, investment screened for backward linkages, RMG upgrading beyond cut-make-trim, and shock-absorption buffers sized to the identified volatility regimes, target that composition and resilience question directly. This is independent of any claim about whether the underlying elasticity itself has moved.

These conclusions should be read within the bounds of a single-country, annual frequency sample of T = 52 observations, in which the full time-varying covariance specification did not meet the pre-registered convergence gate. It was replaced by a fallback holding the contemporaneous covariance loadings constant. A residual convergence soft spot in the FDI volatility path, and judgment calls in classifying GDF growth and FDI as level series against mixed unit-root evidence are additional sources of uncertainty. These findings should not be extrapolated beyond the sample period or read as evidence that the transmission coefficients are structurally incapable of shifting. Quarterly or higher-frequency Bangladesh data, or a fully converged time-varying covariance structure, is the natural next step for detecting coefficient movement should it exist and for tracing how shock variance is transmitted across the trade, FDI and growth equations.

This study relocates, rather than settles, the question of how trade and FDI shape Bangladesh' growth path: the evidence says structural change lives in the variance of the shocks, not in the mechanism converting them into growth.

References

Adeleye, B. N., Nathaniel, S., Ogunrinola, I., & Ikuemonisan, E. (2021). Financial Deepening and Income Inequality in Nigeria: Evidence from Zivot-Andrews and Gregory-Hansen Structural Break Analyses. *Studies of Applied Economics*, *39*(8). https://doi.org/10.25115/eea.v39i8.4348

Adhikary, B. K. (2010). FDI, Trade Openness, Capital Formation, and Economic Growth in Bangladesh: A Linkage Analysis. *International Journal of Business and Management*, *6*(1), p16. https://doi.org/10.5539/ijbm.v6n1p16

Ahmad, Y., Check, A., & Lo, M. C. (2024). Unit Roots in Macroeconomic Time Series: A Comparison of Classical, Bayesian and Machine Learning Approaches. *Computational Economics*, *63*(6), 2139–2173. https://doi.org/10.1007/s10614-023-10397-0

Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, *18*(1), 1–22. https://doi.org/10.1002/jae.659

Banbura, M., Giannone, D., & Reichlin, L. (2010). Large Bayesian VARs. *SSRN Electronic Journal*, *25*(1), 71–92. https://doi.org/10.2139/ssrn.1292332

Canova, F., Ferroni, F., & Matthes, C. (2015). Approximating Time Varying Structural Models with Time Invariant Structures. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2697582

Chan, J. C. C. (2022). *Large Hybrid Time-Varying Parameter VARs* (arXiv:2201.07303). arXiv. https://doi.org/10.48550/arXiv.2201.07303

Chan, J. C. C., & Eisenstat, E. (2015). Bayesian Model Comparison for Time-Varying Parameter VARs with Stochastic Volatility. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2642091

Chan, J. C. C., & Eisenstat, E. (2018). Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *Journal of Applied Econometrics*, *33*(4), 509–532. https://doi.org/10.1002/jae.2617

Cieślik, A., & Tarsalewska, M. (2011). External Openness and Economic Growth in Developing Countries. *Review of Development Economics*, *15*(4), 729–744. https://doi.org/10.1111/j.1467-9361.2011.00639.x

Clark, T. E., & Ravazzolo, F. (2015). The Macroeconomic Forecasting Performance of Autoregressive Models with Alternative Specifications of Time-Varying Volatility. *Journal of Applied Econometrics*, *30*(4), 551–575.

Clements, M. P., & Mizon, G. E. (1991). Empirical analysis of macroeconomic time series. *European Economic Review*, *35*(4), 887–917. https://doi.org/10.1016/0014-2921(91)90042-H

Cogley, T., & Sargent, T. J. (2005). *Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. 8*(2), 262–302. https://doi.org/https://doi.org/10.1016/j.red.2004.10.009

Gnangnon, S. K. (2019). Multilateral trade liberalisation helps promote export product diversification: Trade tensions damage the prospects of the poorest economies. *Economic Affairs*, *39*(3), 363–380. https://doi.org/10.1111/ecaf.12366

Haque, S., & Amin, S. B. (2018). Foreign Direct Investment, Trade Openness and Economic Growth: The Case of Bangladesh. *World Review of Business Research*, *8*(3), 104–125.

He, C., Teräsvirta, T., & González, A. (2008). Testing Parameter Constancy in Stationary Vector Autoregressive Models Against Continuous Change. *Econometric Reviews*, *28*(1–3), 225–245. https://doi.org/10.1080/07474930802388041

Huchet-Bourdon, M., Le Mouël, C., & Vijil, M. (2018). The relationship between trade openness and economic growth: Some new insights on the openness measurement issue. *The World Economy*, *41*(1), 59–76. https://doi.org/10.1111/twec.12586

Islam, S. (2016). EXCLUSION OF GENERALIZED SYSTEM OF PREFERENCES (GSP): NEGATIVE IMPACT ON FOREIGN DIRECT INVESTMENT (FDI) FOR BANGLADESH. *Journal of Global Economics, Management and Business Research*, *7*(1), 52–60.

Killick, R., & Eckley, I. A. (2014). changepoint: An *R* Package for Changepoint Analysis. *Journal of Statistical Software*, *58*(3). https://doi.org/10.18637/jss.v058.i03

Kinfack, E., & Bonga-Bonga, L. (2023). The growth effect of trade openness on African countries: Evidence from using an instrumental variable panel smooth transition model. *Bulletin of Economic Research*, *75*(2), 366–379. https://doi.org/10.1111/boer.12360

Majumder, S. C., & Rana, Md. M. (2016). Trade Liberalization and Its Effects on the Economic Growth of Bangladesh: An Empirical Analysis. *American Journal of Trade and Policy*, *3*(2), 61–70. https://doi.org/10.18034/ajtp.v3i2.402

Malik, A., & Sah, A. N. (2024). Can stronger governance and institutional quality drive growth through inward foreign direct investments in BRICS nations? *South African Journal of Economic and Management Sciences*, *27*(1), a5785. https://doi.org/10.4102/SAJEMS.v27i1.5785

Mendoza, P., Quineche, R., & Zapata, J. (2026). THE TWO SIDES OF INFLATION IN PERU: DYNAMICS AND MONETARY POLICY TRANSMISSION. *International Economics*, 100733. https://doi.org/10.2139/ssrn.6290358

Phillips, P. C. B., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, *75*(2), 335–346. https://doi.org/10.1093/biomet/75.2.335

Plakandaras, V., Gupta, R., & Wohar, M. E. (2018). UK macroeconomic volatility: Historical evidence over seven centuries. *Journal of Policy Modeling*, *40*(4), 767–789. https://doi.org/10.1016/j.jpolmod.2018.04.002

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *The Review of Economic Studies*, *72*(3), 821–852. https://doi.org/https://doi.org/10.1111/j.1467-937X.2005.00353.x

Razzaque, M. A., Bidisha, S. H., & Khondker, B. H. (2017). Exchange Rate and Economic Growth: An Empirical Assessment for Bangladesh. *Journal of South Asian Development*, *12*(1), 42–64. https://doi.org/10.1177/0973174117702712

Ridwan, M., Tahsin, M. S., Al-Absy, M. S. M., Eleais, Md., Ridzuan, A. R., & Mukthar, K. P. J. (2025). Economic Alchemy: Unraveling the Nexus between Trade Openness, Inflation, Exchange Rates, and Economic Growth in Bangladesh. *International Journal of Economics and Financial Issues*, *15*(3), 244–253. https://doi.org/10.32479/ijefi.17779

Roy, V. (2019). *Convergence diagnostics for Markov chain Monte Carlo* (arXiv:1909.11827). arXiv. https://doi.org/10.48550/arXiv.1909.11827

Saleem, H., Shabbir, M. S., & Bilal Khan, M. (2020). The short-run and long-run dynamics among FDI, trade openness and economic growth: Using a bootstrap ARDL test for co-integration in selected South Asian countries. *South Asian Journal of Business Studies*, 9(2), 279–295. https://doi.org/10.1108/SAJBS-07-2019-0124

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, *48*(1), 1. https://doi.org/10.2307/1912017

Tabassum, N., & Ahmed, S. P. (2014). Foreign Direct Investment and Economic Growth: Evidence from Bangladesh. *International Journal of Economics and Finance*, *6*(9), p117. https://doi.org/10.5539/ijef.v6n9p117

Terrell, M., Haque, Q., Cross, J. L., & Doko Tchatoka, F. (2026). Monetary Policy Shocks and Exchange Rate Dynamics in Small Open Economies. *Journal of Applied Econometrics*, jae.70051. https://doi.org/10.1002/jae.70051

Umar, A. H., & Aldio, K. M. (2026). *Macroeconomic Policy Coordination and Economic Resilience in Indonesia: A Time-Varying Analysis under Global Uncertainty*. In Review. https://doi.org/10.21203/rs.3.rs-9663421/v1

Vats, D., & Knudson, C. (2021). Revisiting the Gelman–Rubin Diagnostic. *Statistical Science*, *36*(4). https://doi.org/10.1214/20-STS812

Wang, X., Sui, X., Wang, Y., & Zhang, X. (2025). A Randomised Test to Distinguish Time-Varying Coefficient Models. *Oxford Bulletin of Economics and Statistics*, *87*(6), 1227–1241. https://doi.org/10.1111/obes.12680