

Monetary Policy, Exchange Rate Shocks, and Asset Prices in a Small Open Frontier Economy: A Restricted-Sample, Multi-Horizon Distributional Analysis

Isaac Nyame

Kwame Nkrumah University of Science and Technology

Gabriel Osei Forkuo

gabriel.forkuo@unitbv.ro

Transylvania University of Braşov

Research Article

Keywords: Bayesian quantile regression, Ghana Stock Exchange, wavelet coherence, Toda-Yamamoto causality, forecast error variance decomposition, frontier markets

Posted Date: October 6th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10884319/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

We examine how monetary policy shocks transmit to equity returns in Ghana, a frontier economy with a shallow, thinly traded exchange. A central methodological contribution of this paper is a restricted-sample design. We use only Ghana Stock Exchange Composite Index (GSE-CI) and Financial Stock Index (GSE-FSI) observations built from directly reported, independently verified index levels, rather than levels reconstructed by interpolation. This excludes 48 of 132 GSE-FSI months (January 2014 to December 2017) that would otherwise be deterministic by construction. GSE-CI analysis uses the full January 2014 to December 2024 window ($N = 131$ monthly returns). GSE-FSI analysis is restricted to February 2018 through December 2024 ($N = 83$ monthly returns, since the first return in the verified window still depends on an interpolated base level and is also dropped). Within this design we apply Zivot-Andrews structural-break unit root tests, Toda-Yamamoto augmented-VAR causality with Benjamini-Hochberg false discovery rate control, Bayesian quantile regression with an embedded spike-and-slab variable selection prior, structural VAR impulse responses and variance decomposition under a corrected recursive ordering, and continuous Morlet wavelet coherence benchmarked against Monte Carlo red-noise significance. The results tell a modest story. Policy rate shocks show no causal relationship with either index that survives multiple-testing correction, and account for 0.5 to 3.1 percent of twelve-month-ahead GSE-CI forecast error variance and 0.7 to 7.4 percent of GSE-FSI forecast error variance. Money supply growth is the one macro-financial variable with a causal link to GSE-CI that survives false discovery rate control. Oil price shocks carry meaningful posterior support for GSE-CI at upper (bullish) quantiles. Treasury bill rate changes carry posterior support for GSE-FSI at the tenth percentile, though this estimate rests on roughly eight effective observations and should be read as suggestive rather than conclusive. Wavelet coherence between policy shocks and equity returns rises with horizon in raw terms, but once benchmarked against an autoregressive red-noise null, only a small and non-monotonic share of that coherence is statistically distinguishable from noise, which weakens rather than supports a horizon-dependence claim. Taken together, these results indicate that Ghana's equity market is loosely and unevenly coupled to the domestic policy rate, with global commodity prices and short-term funding costs playing a larger, more state-dependent role, and that claims about monetary transmission in thin frontier markets need index-construction and multiple-testing safeguards that are not always applied in this literature.

JEL Classification Codes: E52; G10; G12; G15; C11; C32; O16; O55

1. Introduction

1.1 Background

Monetary policy transmission to asset prices is a well-established channel in advanced-economy macro-finance, but its operation in frontier markets remains far less settled. Frontier equity markets such as Ghana's are characterised by thin trading, a small number of listed firms concentrated in a few sectors, limited institutional participation (Alagidede & Panagiotidis, 2009; Alagidede, 2011), and monetary

regimes that have, at various points, prioritised exchange-rate stabilisation, inflation targeting, and fiscal accommodation in different measure. Whether policy rate changes propagate to equity prices through the standard discount-rate and expected-earnings channels, or whether frontier markets are simply too illiquid and too dominated by a handful of large financial-sector stocks for that channel to operate cleanly, is an empirical question that has not been convincingly settled for Ghana. Prior Ghana-specific evidence on macroeconomic determinants of stock returns and market efficiency is limited and dated (Fama, 1970; Adam & Tweneboah, 2008; Frimpong, 2008), and existing macro-finance research on Ghana has focused primarily on growth and capital-flow linkages rather than monetary-equity transmission (Frimpong & Oteng-Abayie, 2006).

The Ghanaian case is a useful laboratory for this question precisely because it has experienced repeated, large-amplitude monetary cycles within a short span. The Bank of Ghana's policy rate ranged from single digits in the mid-2000s to a peak of 30 percent in late 2023, before easing back toward the high 20s in 2024 (Bank of Ghana, 2023). This occurred against a backdrop of currency depreciation episodes, a 2022–2023 sovereign debt restructuring, and inflation that touched above 50 percent year-on-year before falling sharply (World Bank, 2023). At the same time, the structure of the equity market itself changed materially in this period. The Ghana Stock Exchange replaced its original All-Share Index with two purpose-built indices, the GSE Composite Index (GSE-CI) and the GSE Financial Stock Index (GSE-FSI), with a new base value of 1,000 points as of 31 December 2010. Because these two indices did not exist in their current form before 2011, any study that presents monetary-transmission estimates for a sample starting earlier necessarily blends two structurally different index-construction regimes, which is likely to introduce spurious breaks that have nothing to do with monetary policy. We therefore restrict our sample to January 2014, chosen to allow the new indices several years to establish a stable trading history, through December 2024.

A second, and in this revision more consequential, restriction applies specifically to GSE-FSI. Independently reported GSE-FSI year-end closes are not available for 2014 through 2017; the values used in early drafts of this research were log-linearly interpolated between December year-end anchors, which makes the implied monthly returns over that period deterministic by construction rather than observed (see Section 3.1). We do not carry that interpolated segment into the empirical analysis reported here. GSE-CI analysis uses the full sample, January 2014 to December 2024, giving 132 monthly level observations and 131 usable return observations. GSE-FSI analysis is restricted to independently verified level observations, February 2018 through December 2024, giving 83 usable return observations (the January 2018 return is also excluded because it is computed against a December 2017 base level that is itself interpolated). This restriction is a genuine methodological contribution in its own right: it trades sample length for the assurance that no reported GSE-FSI result rests on constructed rather than observed data, and it means several estimators reported for GSE-CI (medium-to-long-run wavelet bands, tail quantiles) cannot be supported for GSE-FSI on the shorter verified window; we state explicitly, method by method, where that constraint binds.

A third motivation for revisiting this question empirically, rather than relying on priors from advanced-economy transmission studies, is that the effect of policy rate changes on frontier equity markets is theoretically ambiguous in sign and likely to be state-dependent. Higher policy rates raise the discount rate applied to expected future cash flows, a negative effect on prices, but can also signal credible disinflation and currency stabilisation, which may support risk appetite in a market where currency risk is often the dominant concern for both domestic and foreign investors (a potentially offsetting positive effect). Financial-sector equities, which make up a large share of GSE-FSI by construction, carry an additional, competing channel: higher rates can widen net interest margins for banks even as they raise credit risk and funding costs. These competing channels motivate a distributional approach that allows the sign and magnitude of the policy-rate effect to vary across the return distribution, rather than imposing a single average effect.

1.2 Aim, Objectives, and Research Questions

This study aimed to provide a rigorous characterisation of how monetary policy shocks, alongside exchange-rate, money-supply, inflation, and global oil-price shocks, transmit to Ghanaian equity returns across time, frequency, and the return distribution, under a sample design that excludes constructed observations. The specific objectives were to: (i) establish the integration properties of each series using both conventional and structural-break-robust unit root tests, applied under a single, pre-specified classification rule; (ii) test the direction of Granger-causal linkages between monetary and financial variables using a lag-augmentation approach robust to uncertain integration order, with correction for the number of tests conducted; (iii) estimate the distributional response of equity returns to monetary and macroeconomic shocks using a Bayesian framework in which variable selection is embedded in the same posterior as the coefficients; (iv) quantify the dynamic, aggregate importance of monetary shocks relative to other macro-financial shocks using structural VAR impulse responses and forecast error variance decomposition under an economically motivated recursive ordering, with an order-invariant robustness check; and (v) characterise how the strength of co-movement between monetary policy and equity returns varies across time horizons using wavelet coherence benchmarked against a red-noise null.

Correspondingly, the study addressed five research questions. First, are the monetary, exchange-rate, and equity-return series integrated of the same order once structural breaks are accounted for, and does a single, consistently applied decision rule change any classification relative to using multiple tests informally? Second, is there Granger-causal evidence of a directional relationship between monetary policy shocks and equity returns that survives correction for the number of tests conducted, or does the evidence instead point toward other macro-financial channels? Third, do monetary policy shocks affect the upper and lower quantiles of the equity return distribution differently, and how much posterior support do the data provide for including each candidate predictor at each quantile once variable selection is estimated jointly with the coefficients? Fourth, what share of the forecast error variance in equity returns is attributable to monetary shocks specifically, and how sensitive is that share to the identifying ordering assumption? Fifth, does the strength of monetary-equity co-movement vary

systematically across short-, medium-, and long-run horizons once compared against what a red-noise process with no relationship would produce, and is that pattern consistent across the composite and financial-sector indices?

2. Literature Review

2.1 Theoretical Framework

2.1.1 Monetary Policy Transmission Mechanisms

The conventional starting point for monetary transmission to asset prices is the present-value channel: a policy rate increase raises the discount rate applied to expected future dividends, mechanically lowering equity valuations, all else equal (Gordon, 1962; Bjørnland & Leitemo, 2009). A complementary channel operates through expected earnings. Tighter policy dampens aggregate demand and, with a lag, corporate profitability (Bernanke & Blinder, 1992; Bernanke & Gertler, 1995; Taylor, 1995; Romer & Romer, 2004). In open economies, a third channel operates through the exchange rate: policy tightening that attracts capital inflows and appreciates the currency can affect firms differently depending on their exposure to imported inputs, foreign-currency debt, and export revenue (Adjasi & Biekpe, 2005). For a commodity-linked, partially dollarised economy such as Ghana's, this exchange-rate channel is plausibly at least as important as the discount-rate channel, and the two can work in opposite directions depending on the source of the shock.

A separate literature emphasises that the strength of monetary transmission to asset prices depends on financial market depth and the credibility of the monetary regime (Adrian & Shin, 2010; Bekaert et al., 2011; Brunnermeier & Pedersen, 2009). In markets with few listed firms, low free float, and a narrow investor base, as in Ghana, price discovery may be slow and thin trading can weaken or delay the standard discount-rate channel, while amplifying sensitivity to idiosyncratic liquidity events unrelated to monetary policy. This motivates treating the strength, and even the sign, of monetary transmission as an empirical rather than a purely theoretical question in this setting.

2.1.2 Distributional and Nonlinear Transmission

A growing body of work argues that monetary policy shocks do not affect the entire distribution of asset returns uniformly. Quantile-based approaches (Koenker & Bassett, 1978) allow the estimated effect of a shock to differ between calm and turbulent, or bullish and bearish, market states, which is important if, for example, policy tightening has a disproportionate effect on returns precisely when the market is already under stress (lower quantiles), or if it primarily curbs speculative upside (upper quantiles) (Longin & Solnik, 2001). Bayesian quantile regression additionally allows formal uncertainty quantification via posterior credible intervals (Yu & Moyeed, 2001) and, when variable selection is embedded directly in the same Gibbs sampler that produces the coefficients (Koop & Korobilis, 2012; Korobilis, 2017), allows the researcher to report how confident the data actually are that a given predictor

belongs in the model at a given quantile, rather than asserting inclusion based on a single specification or a separate, disconnected model-comparison exercise.

2.1.3 Time-Frequency Transmission Dynamics

Because monetary policy transmission is understood to operate with variable and uncertain lags, time-domain methods alone may understate or mischaracterise the relationship between policy shocks and asset prices. Wavelet coherence analysis decomposes the co-movement between two series jointly across time and frequency, equivalently horizon (Rua & Nunes, 2009), making it possible to ask whether monetary-equity co-movement is concentrated at short, medium, or long horizons, and whether that concentration is stable over the sample or specific to particular episodes such as currency crises (Aguiar-Conraria & Soares, 2011). Because the smoothing operator that makes wavelet coherence well defined mechanically averages over fewer effective degrees of freedom at longer scales, raw coherence tends to rise with horizon even under a null of no relationship (Torrence & Compo, 1998; Grinsted et al., 2004); any horizon-dependence claim therefore needs to be benchmarked against a red-noise null before it is interpreted substantively, a point we return to in Section 3.2.5.

2.2 Conceptual Framework

We conceptualise monetary policy transmission to Ghanaian equity returns as operating through three potentially competing channels (Fig. 1): a direct discount-rate/valuation channel, an exchange-rate and imported-inflation channel (Fama, 1981; Geske & Roll, 1983; Kutan & Aksoy, 2003), and a financial-sector net-interest-margin channel specific to bank-heavy indices such as GSE-FSI (Brunnermeier et al., 2020; Acharya et al., 2017), whose relative strength may vary across the return distribution and across time horizons. Rather than imposing a single expected sign or a single dominant channel a priori, our empirical strategy is designed to let the data reveal which channel dominates, under what market conditions, and at what horizon, while treating global oil prices and domestic money-supply growth as additional shocks competing for explanatory power alongside the policy rate.

3. Methodology

3.1 Data Sourcing and Processing

Monthly data span January 2014 to December 2024. GSE-CI analysis uses the full window (N = 132 level observations, 131 return observations after first-differencing). GSE-FSI analysis, for reasons set out below, is restricted to February 2018 through December 2024 (N = 83 return observations). All series are drawn from official primary sources and, where a bulk machine-readable export was unavailable, from primary institutional publications, with every non-directly-exported observation flagged and its source documented in the companion data appendix and in Table 1 below.

The Ghana Stock Exchange Composite Index (GSE-CI) and Financial Stock Index (GSE-FSI) monthly closing levels were compiled from the Ghana Stock Exchange's own monthly Equities Market Report publications, which report current- and prior-year monthly closes side by side; this allowed cross-

validation of overlapping months across multiple independently published reports. Every GSE-CI observation for the full sample, and every GSE-FSI observation for January 2018 through December 2024 (84 of 132 months), matched exactly across at least two independently sourced reports. GSE-FSI monthly values for January 2014 through December 2017 (48 of 132 months) could not be sourced at a monthly frequency from any available report; the only real observations available for that period are the December year-end closes (2013: 1,786.60; 2014: 2,243.63; 2015: 1,930.06; 2016: 1,545.41; 2017: 2,310.58).

An earlier version of this analysis filled the 2014–2017 gap by log-linear interpolation between those year-end anchors. We no longer do this. Interpolating the log index level linearly within each year makes the implied monthly return constant within that year: applied across four consecutive years, this produces a GSE-FSI return series with exactly four distinct values over 2014–2017, each repeated twelve times, and no within-year variation whatsoever. Every downstream statistic computed on that series, descriptive moments, unit root and structural break tests, VAR lag selection and residual diagnostics, quantile regression, and wavelet coherence, would be estimated partly on constructed rather than observed data, and the two cannot be cleanly separated once pooled into a single regression or transform. We therefore exclude January 2014 through December 2017 from all GSE-FSI analysis. Because the first return in the remaining window (January 2018) is computed against the interpolated December 2017 level, we exclude that observation as well. GSE-FSI results in this paper accordingly rest on 83 return observations, February 2018 through December 2024, each computed between two independently verified levels. This is a shorter series than GSE-CI's, and shorter than would be needed to support some of the estimators used elsewhere in this paper; we state explicitly, in Sections 3.2.3 and 3.2.5, where that constraint binds and what we do about it.

Three other series in the panel contain a smaller number of constructed observations, summarised in Table 1. The monetary policy rate (MPR) for January 2014–December 2023 and the 91-day Treasury bill rate (TBR) for the same period were obtained from the Ghana Statistical Service's StatsBank statistical database, itself sourced from the Bank of Ghana. MPR values for January–December 2024 were reconstructed from the Bank of Ghana's own Monetary Policy Committee press releases, which document the exact date, meeting number, and basis-point change or hold decision at each of the six 2024 policy meetings; all twelve 2024 MPR observations are real decisions. TBR for 2024 is real in four months (January, February, June, December, each cross-checked against point-in-time Treasury bill auction rates) and linearly interpolated between adjacent real anchors in the other eight. Broad money supply (M2) for January 2014–July 2024 is sourced from the same StatsBank database; August–November 2024 values (4 of 132 months) are linearly interpolated between the real July 2024 and December 2024 Bank of Ghana Quarterly Statistical Bulletin observations. The GHS/USD exchange rate for January 2014–December 2023 is sourced from StatsBank; 2024 monthly values are interpolated between the real December 2023 observation and a real November 2024 point from the Ghana Stock Exchange's own published economic indicators table, leaving eleven of twelve 2024 exchange-rate months as constructed. The consumer price index level and year-on-year inflation rate are sourced

directly from StatsBank for the full window with no gaps. Brent crude oil spot prices are sourced from the IMF Primary Commodity Price System via the FRED series POILBREUSD, with complete coverage.

Table 1
Constructed (interpolated or reconstructed) observations by series, out of 132 total monthly observations.

Series	Total obs.	Constructed obs.	Share constructed	Treatment in this paper
GSE-FSI (level)	132	48	36.4%	Excluded from all analysis (Section 3.1)
Exchange rate (EXR)	132	11	8.3%	Retained; all constructed obs. fall in 2024
Treasury bill rate (TBR)	132	8	6.1%	Retained; all constructed obs. fall in 2024
Broad money supply (M2)	132	4	3.0%	Retained; constructed obs. are Aug-Nov 2024
GSE-CI (level)	132	0	0.0%	Fully observed
Policy rate (MPR)	132	0	0.0%	Fully observed (2024 values are real MPC decisions)
CPI / inflation	132	0	0.0%	Fully observed
Brent crude oil	132	0	0.0%	Fully observed

Note: 60 of 132 months (45.5%) contain at least one constructed value across the whole panel once GSE-FSI is included; excluding GSE-FSI, 15 of 132 months (11.4%) contain a constructed EXR, TBR, or M2 value, concentrated in 2024. We retain EXR, TBR, and M2 despite their smaller constructed shares because, unlike GSE-FSI, the constructed months are a minority of each series, are concentrated in a single year rather than spread across four, and (for TBR and MPR) are anchored to real, dated policy decisions rather than interpolated between distant endpoints. We nonetheless flag this as a limitation in Section 5.6 and report all results with the caveat that late-2024 estimates involving EXR and TBR draw partly on constructed values.

Return and change series are constructed as simple monthly percentage changes for level variables expressed as index or price levels (GSE-CI, GSE-FSI, M2, the exchange rate, and Brent oil), and as simple percentage-point changes for variables already expressed as rates (MPR, TBR, and year-on-year inflation). We adopt percentage-point differencing for rate variables specifically to avoid the scale distortion that arises from taking percentage changes of a rate that is itself already a percentage.

3.2 Econometric Framework

3.2.1 Unit Root and Structural Break Tests

We test each series for a unit root using the augmented Dickey-Fuller (ADF) test with both a constant-only and a constant-plus-trend specification, and using the Zivot-Andrews (ZA) test (Zivot & Andrews,

1992), which allows for a single endogenously determined structural break in the level (building on Perron, 1989) and is therefore informative for a sample containing the 2014–2015 currency crisis, the COVID-19 shock, and the 2022–2023 debt and inflation crisis. Lag length in each test is selected by the Bayesian Information Criterion.

A series is classified as integrated of order one, $I(1)$, using a single, pre-specified rule: the constant-plus-trend ADF test cannot reject a unit root in levels at the 5 percent level, and the same test rejects a unit root in first differences at the 5 percent level. We use ADF, rather than ZA, as the primary classification test because ZA tests a different null (a unit root allowing for one break) against a different alternative, and a ZA rejection in levels does not by itself imply the ADF null would also be rejected; treating either test's rejection as sufficient, as in earlier drafts of this analysis, produced an internally inconsistent classification for the policy rate, which rejects the ZA level null ($p = 0.020$) but not the ADF level null ($p = 0.833$). Under the single-rule classification, the policy rate is $I(1)$; we report the ZA result alongside it for corroboration and structural-break dating, not as an alternative sufficient condition for classification.

3.2.2 Toda-Yamamoto Causality Testing

To test Granger-causal ordering without pre-testing for cointegration, we implement the Toda and Yamamoto (1995) augmented-lag procedure. For each pair of variables we estimate a level VAR with lag order p selected by Akaike's Information Criterion separately for that bivariate system, augmented with $d_{\max} = 1$ additional lag reflecting the maximum order of integration found in Section 3.2.1, and test the joint significance of the coefficients on the first p (non-augmented) lags of the candidate causal variable using a Wald test. We report the selected lag order p explicitly for every bivariate system in section 4.3, since it varies across pairs and need not match the lag order selected for the full six-variable VAR reported in Section 3.2.4; the two are different models estimated for different purposes; conflating them was a source of an apparent inconsistency in earlier drafts of this paper. We report this alongside a conventional Granger causality test (Granger, 1969; Geweke, 1982) on the differenced (stationary) series as a cross-check. Because we conduct 24 Toda-Yamamoto tests (twelve macro-financial variables, two directions, two equity indices), we apply the Benjamini-Hochberg (1995) procedure to control the false discovery rate across this family of tests, and report both the raw p -value and the resulting q -value for every test in section 4.3.

3.2.3 Bayesian Quantile Regression with Embedded Variable Selection

To characterise how the effect of monetary and macro-financial shocks on equity returns varies across the return distribution, we estimate Bayesian quantile regression (BQR) using the asymmetric Laplace likelihood and its normal-exponential mixture representation (Kozumi & Kobayashi, 2011), via Gibbs sampling with 4,500 post-burn-in draws following a 1,500-draw burn-in, run from two independent starting points to compute Gelman-Rubin convergence diagnostics (R -hat) for every coefficient at every quantile; all reported R -hat values are below 1.01. We estimate at five quantiles, $\tau = 0.10, 0.25, 0.50, 0.75, 0.90$, a symmetric grid chosen to replace an earlier, unexplained asymmetric grid (0.05, 0.10, 0.25,

0.30, 0.50, 0.75, 0.80, 0.90) that mixed adjacent quantiles at some points while omitting the upper-tail counterpart to $\tau = 0.05$ entirely. We drop $\tau = 0.05$ throughout: with an effective sample size of roughly $\min(\tau, 1-\tau)$ times N informing extreme-quantile estimation, $\tau = 0.05$ leaves approximately 4 effective observations for GSE-FSI ($N = 83$) and 7 for GSE-CI ($N = 131$) against six candidate predictors, too few to support a stable posterior. We report the effective sample size at every retained quantile alongside the results, so readers can weigh the $\tau = 0.10$ estimates (roughly 8 effective observations for GSE-FSI, 13 for GSE-CI) accordingly.

Table 2 shows Treasury bill rate changes have a skewness of -5.57 and a kurtosis of 54.37, indicating one or two extreme observations, plausibly associated with the 2022–2023 debt restructuring, that could otherwise dominate a tail-quantile estimate on their own. We therefore winsorize TBR change at the first and 99th percentiles before it enters the quantile regressions (all other predictors are used unwinsorized). Both the dependent variable and all six candidate predictors (MPR shock, winsorized TBR change, money-supply growth, exchange-rate return, inflation change, and oil-price return) are standardised to zero mean and unit variance prior to estimation, so that reported coefficients represent the effect, in standard deviations of the equity return, of a one-standard-deviation change in the predictor.

To address model uncertainty over which of the six candidate predictors belong in the model at each quantile, we embed a spike-and-slab stochastic search variable selection (SSVS) prior directly in the Gibbs sampler, rather than computing posterior inclusion probabilities from a separate procedure. Each slope coefficient carries a binary inclusion indicator; conditional on inclusion the coefficient is drawn from a diffuse “slab” normal prior, and conditional on exclusion from a narrow “spike” normal prior concentrated near zero, following George and McCulloch (1993) and, for the quantile regression setting specifically, the general logic of Koop and Korobilis (2012). The inclusion indicators are updated within the same Markov chain that draws the coefficients, so the reported posterior inclusion probability (PIP) for each predictor is a draw from the actual estimated posterior, not an approximation computed from a disjoint frequentist search over predictor subsets, which was the approach used in earlier drafts of this paper and which does not correspond to the posterior of the model whose coefficients are reported elsewhere in the same table. Following standard conventions, we treat PIP values above 0.95 as very strong evidence of inclusion, 0.75–0.95 as substantial, 0.50–0.75 as weak-to-moderate, and below 0.50 as weak evidence against inclusion; we report PIP for every predictor at every quantile rather than asserting inclusion selectively. We note that embedding selection in the sampler also provides a principled, model-based treatment of the multiple-comparisons problem raised in Section 3.2.2: the spike prior shrinks weakly supported predictors toward zero across all quantile-predictor combinations jointly, rather than testing each combination as an independent frequentist hypothesis.

3.2.4 Vector Autoregression, Impulse Responses, and Forecast Error Variance Decomposition

To quantify the dynamic, aggregate importance of monetary shocks relative to other macro-financial shocks, we estimate a six-variable VAR comprising the oil-price return, the inflation change, the policy rate shock, the Treasury bill rate change, the exchange-rate return, and the equity return (GSE-CI or GSE-FSI, estimated separately on their respective samples), with lag order selected by AIC. We identify structural shocks via a Cholesky decomposition with variables ordered from slowest-moving and least contemporaneously responsive to fastest-moving and most contemporaneously responsive: the oil price first, as an internationally determined price that domestic conditions cannot affect within the month; domestic inflation second, as a slow-moving series that is sticky within the month; the policy rate third, reflecting the standard recursive assumption that the central bank observes and responds to current oil-price and inflation conditions within the month while its own policy decisions affect the macro-financial variables ordered after it only with a lag; the Treasury bill rate fourth, as a money-market rate that adjusts to policy decisions within the month; the exchange rate fifth; and the equity return last, as the fastest-moving variable, which absorbs the contemporaneous effect of every other shock in the system. An earlier draft of this paper ordered the policy rate first, on the stated grounds that policy responds to macro-financial conditions with a lag while affecting them contemporaneously; that is the reverse of the standard assumption used here, under which the central bank is normally taken to respond within the month to the conditions it observes, while those conditions respond to policy with a lag. We correct that ordering in this revision. Because results from a recursive VAR can be sensitive to the ordering assumption, we additionally report the generalised impulse response function (GIRF) of Pesaran and Shin (1998), which is invariant to variable ordering, and we report impulse responses and variance decompositions under the original (policy-rate-first) ordering alongside the corrected ordering, so the sensitivity of the result to this identifying assumption is visible rather than hidden. We verify that the estimated VAR is stable, all companion-matrix eigenvalue moduli below unity, in every specification before proceeding, and report orthogonalised impulse responses and forecast error variance decomposition at one-, three-, six-, and twelve-month horizons.

3.2.5 Continuous Morlet Wavelet Coherence Analysis, Benchmarked Against a Red-Noise Null

To examine how the strength of co-movement between the policy rate shock and equity returns varies jointly across time and frequency, we compute continuous wavelet coherence using the Morlet wavelet, following Torrence and Compo (1998), with time and scale smoothing following Grinsted, Moore and Jevrejeva (2004). Because the smoothing operator required for coherence to be well defined averages over fewer effective degrees of freedom at longer scales, raw band-averaged coherence is expected to rise with horizon under a null of no relationship between the two series; a monotonic rise in raw coherence therefore cannot by itself support a substantive horizon-dependence claim. We benchmark observed coherence against a Monte Carlo null generated from 150 pairs of AR(1) surrogate series matched to the autoregressive coefficient and innovation variance of each observed series, following Grinsted, Moore and Jevrejeva (2004), and report both the raw band-averaged coherence and the share of coherence, restricted to points outside the cone of influence, that exceeds the 95th percentile of the Monte Carlo null at each point. We summarise coherence within three conventional horizon bands: short-

run (2–6 months), medium-run (6–18 months), and long-run (18–36 months). For GSE-FSI, whose verified sample (83 months) leaves very little area outside the cone of influence at the 18–36 month scale, we do not report the long-run band; this is a direct consequence of the sample restriction described in Section 3.1, not an oversight.

4. Results

4.1 Descriptive Statistics

Table 2 reports descriptive statistics for level and return/change series. GSE-CI averaged 2,525.1 points (SD = 642.6) over the full 2014–2024 sample (Fig. 2). GSE-FSI, restricted to its verified February 2018–December 2024 window, averaged 2,071.1 points (SD = 312.7). Both level series show positive skewness and excess kurtosis. At the return level, mean monthly GSE-CI returns are 0.71 percent (SD = 5.01 percent, N = 131); mean monthly GSE-FSI returns over the verified window are – 0.11 percent (SD = 4.37 percent, N = 83), reflecting the different sub-period covered once the constructed 2014–2017 segment is excluded rather than a change in measurement. Jarque-Bera tests reject normality for every return series (all $p < 0.01$), motivating the quantile-based approach used in Section 4.4. The policy rate shock (MPR) has a mean monthly change of essentially zero (0.084 percentage points, SD = 0.755), consistent with a rate held constant most months and adjusted in discrete steps at scheduled Monetary Policy Committee meetings. Treasury bill rate changes show pronounced excess kurtosis (54.37) and negative skewness (-5.57), consistent with one or two extreme observations around the 2022–2023 debt restructuring; we address this directly in Section 3.2.3 by winsorizing this series before it enters the quantile regressions.

Table 2
Descriptive statistics for level and return/change series.

Variable	N	Mean	SD	Skewness	Kurtosis	Jarque-Bera	p-value
GSE-CI (level)	132	2525.1	642.6	1.47	5.48	81.4	< 0.001
GSE-CI return (%)	131	0.71	5.01	1.07	5.24	52.3	< 0.001
GSE-FSI (level, 2018–2024)	83	2071.1	312.7	1.62	6.20	71.8	< 0.001
GSE-FSI return (% , 2018–2024)	83	-0.11	4.37	-0.31	5.14	17.2	< 0.001
MPR change (pp)	131	0.084	0.755	1.07	7.50	135.8	< 0.001
TBR change (pp)	131	0.063	1.689	-5.57	54.37	15079.7	< 0.001
M2 growth (%)	131	1.9	2.6	0.55	3.80	10.2	0.006
EXR return (%)	131	1.6	4.7	0.35	18.78	1362.5	< 0.001
Inflation change (pp)	131	0.076	1.846	-0.02	13.68	623.1	< 0.001
Oil return (%)	131	0.1	9.2	-0.54	5.14	31.3	< 0.001
<i>Note: GSE-FSI figures use the verified February 2018-December 2024 sample only (Section 3.1); all other series use the full January 2014-December 2024 sample. Source: authors' computation; see Table 1 for the source and construction status of each series.</i>							

4.2 Unit Root and Structural Break Results

Table 3 reports ADF and Zivot-Andrews results for each series in levels and first differences, together with the classification under the single decision rule set out in Section 3.2.1 (constant-plus-trend ADF, applied consistently rather than combined with Zivot-Andrews on an either-test basis). GSE-CI, GSE-FSI, MPR, TBR, M2, EXR, and oil all fail to reject the unit-root null in levels under this rule and reject it in first differences at the 1 percent level, and are classified $I(1)$. Inflation rejects the level null under constant-plus-trend ADF ($p = 0.048$) and is classified $I(0)$. The policy rate illustrates why a single rule matters: it rejects the Zivot-Andrews level null ($p = 0.020$), which on its own might suggest trend-stationarity around a break, but does not reject the ADF level null ($p = 0.833$); under the pre-specified rule it is classified $I(1)$, and we proceed with $d_{\max} = 1$ for the Toda-Yamamoto procedure in Section 3.2.2, which remains valid under this classification.

Table 3

Unit root tests: ADF (constant plus trend) and Zivot-Andrews (single endogenous break), levels and first differences.

Variable	ADF level (ct) stat	p	ZA level stat	p	ADF Δ stat	p	Classification
GSE-CI	0.35	0.996	-1.44	0.999	-9.58	< 0.001	I(1)
GSE-FSI (2018–2024)	-2.53	0.314	-2.90	0.936	-7.66	< 0.001	I(1)
MPR	-1.49	0.833	-5.11	0.020	-3.76	0.003	I(1)
TBR	-1.90	0.653	-5.13	0.019	-8.49	< 0.001	I(1)
M2	4.55	1.000	2.70	0.999	-4.18	< 0.001	I(1)
EXR	0.47	0.997	-3.22	0.832	-10.96	< 0.001	I(1)
Inflation	-3.42	0.048	-4.51	0.117	-3.30	0.015	I(0)
Oil	-3.12	0.102	-3.83	0.463	-8.36	< 0.001	I(1)

Source: authors' estimation. GSE-FSI uses the verified 2018–2024 sample.

4.3 Causality Test Results

Table 4 reports standard Granger causality and Toda-Yamamoto results for both directions between each macro-financial variable and each equity index, with the lag order used in each bivariate system reported explicitly (it varies by pair and need not, and generally does not, match the lag order of the six-variable VAR reported in Section 4.5) and with Benjamini-Hochberg q-values reported alongside raw p-values for the 24 Toda-Yamamoto tests.

The one relationship that is both individually significant at conventional levels and survives false discovery rate correction is money-supply growth to GSE-CI (Toda-Yamamoto chi-squared = 50.72, df = 8, $p < 0.001$, $q < 0.001$, an eight-lag bivariate system). The standard Granger test on the differenced series does not reproduce this ($F = 1.34$, $p = 0.249$). We do not read the divergence as unambiguous evidence for the level-VAR approach. Table 3 shows M2 has an ADF statistic of 8.50 in levels, indicating a strongly trending series; an eight-lag bivariate level VAR relating two trending series on a sample of this length is also well suited to fitting shared trend dynamics that a differenced-series test would not pick up. We report the result, note that it survives multiple-testing correction, and flag this alternative reading explicitly as a limitation on how causally the result should be read (Section 5.6).

A second result, exchange-rate-to-GSE-FSI causality, is significant in the restricted GSE-FSI sample (Toda-Yamamoto chi-squared = 18.81, df = 5, $p = 0.002$, $q = 0.025$, a five-lag system) and survives false

discovery rate correction; this result was not detectable in earlier drafts that used the interpolated GSE-FSI series, and we read it as a genuine feature of the corrected data rather than as revisiting a previously reported finding. No other macro-financial variable shows a Toda-Yamamoto relationship with either equity index that survives correction for multiple testing; in particular, the policy-rate-to-equity relationship is not significant for either index under either test (GSE-CI: chi-squared = 3.74, df = 4, p = 0.442; GSE-FSI: chi-squared = 9.91, df = 6, p = 0.128).

Table 4
Toda-Yamamoto and Granger causality tests, both directions, with per-system lag order and FDR-adjusted q-values.

Equity	Macro	TY lag	TY chi2 (Macro to Eq)	p	q	TY lag	TY chi2 (Eq to Macro)	p	q
GSE-CI	MPR	4	3.74	0.442	0.712	4	6.42	0.170	0.560
GSE-CI	TBR	2	1.15	0.562	0.712	2	4.25	0.119	0.560
GSE-CI	MSG	8	50.72	< 0.001	< 0.001	8	9.91	0.271	0.651
GSE-CI	EXR	3	4.81	0.187	0.560	3	0.51	0.917	0.917
GSE-CI	INF	8	6.24	0.620	0.712	8	6.92	0.546	0.712
GSE-CI	OIL	2	0.73	0.696	0.759	2	0.95	0.623	0.712
GSE-FSI	MPR	6	9.91	0.128	0.560	6	4.78	0.573	0.712
GSE-FSI	TBR	2	3.46	0.178	0.560	2	1.69	0.430	0.712
GSE-FSI	MSG	1	0.24	0.621	0.712	1	0.46	0.499	0.712
GSE-FSI	EXR	5	18.81	0.002	0.025	5	1.81	0.874	0.912
GSE-FSI	INF	8	13.16	0.107	0.560	8	10.71	0.219	0.584
GSE-FSI	OIL	2	1.51	0.470	0.712	2	2.10	0.351	0.712

Source: authors' estimation, following Toda and Yamamoto (1995); q-values from Benjamini-Hochberg correction across all 24 tests. GSE-FSI uses the verified 2018–2024 sample throughout.

Table 5 reports Breusch-Godfrey serial-correlation diagnostics from the six-variable VAR underlying Section 4.5, for both equity systems. All Breusch-Godfrey statistics are statistically insignificant (all p > 0.19), indicating no evidence of residual serial correlation in either system. The GSE-CI system selects a

lag order of 5 by AIC; the GSE-FSI system, estimated on a shorter sample, selects a lag order of 8. Both systems satisfy the VAR stability condition (all companion-matrix eigenvalue moduli below unity).

Table 5
Breusch-Godfrey serial-correlation diagnostics, six-variable VAR (corrected ordering).

System	N	VAR lag (AIC)	Stable	Variable	BG stat	BG p
GSE-CI	131	5	Yes	OIL	0.25	0.993
GSE-CI	131	5	Yes	INF	3.02	0.554
GSE-CI	131	5	Yes	MPR	0.48	0.975
GSE-CI	131	5	Yes	TBR	0.48	0.975
GSE-CI	131	5	Yes	EXR	0.14	0.998
GSE-CI	131	5	Yes	EQ	0.27	0.992
GSE-FSI (2018–2024)	83	8	Yes	OIL	1.62	0.805
GSE-FSI (2018–2024)	83	8	Yes	INF	4.05	0.400
GSE-FSI (2018–2024)	83	8	Yes	MPR	1.44	0.837
GSE-FSI (2018–2024)	83	8	Yes	TBR	3.65	0.456
GSE-FSI (2018–2024)	83	8	Yes	EXR	6.08	0.193
GSE-FSI (2018–2024)	83	8	Yes	EQ	0.80	0.939

Source: authors' estimation from the six-variable VAR systems underlying Section 4.5.

4.4 Bayesian Quantile Regression Results

Figure 3 presents the Bayesian quantile regression standardised coefficients, while Table 6 reports posterior means, 95 percent credible intervals, and posterior inclusion probabilities (PIP) from the spike-and-slab embedded variable-selection procedure described in Section 3.2.3, at five quantiles ($\tau = 0.10, 0.25, 0.50, 0.75, 0.90$), together with the effective sample size at each quantile. All coefficients are standardised. All reported Gelman-Rubin R -hat statistics are below 1.01.

For GSE-CI, the posterior mean effect of the standardised policy rate shock is small and inconsistent in sign across the distribution, and its PIP never exceeds 0.09 at any quantile, indicating weak posterior support for including the policy rate shock as a predictor anywhere in the distribution. The standardised oil-price return shows a materially different pattern: its posterior mean effect grows from small and imprecise at the median (0.05) to 0.30 at $\tau = 0.75$ (95% CI: 0.15 to 0.45) and 0.36 at $\tau = 0.90$ (95% CI: 0.05 to 0.54), with PIP of 0.98 at $\tau = 0.75$ ("very strong" evidence) and 0.94 at $\tau = 0.90$ ("substantial" evidence). This indicates global oil-price shocks, rather than the domestic policy rate, carry the strongest

and most consistent posterior support for explaining GSE-CI returns at upper (bullish) quantiles, plausibly reflecting Ghana's status as an oil producer since 2011.

For GSE-FSI, the policy rate shock again shows weak PIP throughout (maximum 0.17, at $\tau = 0.90$). The Treasury bill rate shows its largest PIP at $\tau = 0.10$ (0.64, "weak-to-moderate" evidence under the conventions of Section 3.2.3, not the "very strong" support reported for the uncorrected series in earlier drafts), with a posterior mean of 0.24 (95% CI: -0.05 to 0.51) that spans zero. This estimate is informed by an effective sample of roughly 8 observations at $\tau = 0.10$, and we read it as suggestive of a funding-cost channel concentrated in the lower tail of GSE-FSI returns, consistent with the same channel discussed for the uncorrected data in earlier work, but as noticeably weaker and less certain evidence than previously reported, once it is estimated on a clean sample with variable selection embedded in the sampler rather than computed separately. Oil-price shocks also carry moderate support for GSE-FSI at $\tau = 0.90$ (PIP = 0.63), a pattern not previously reported for this index.

Table 6

Bayesian quantile regression: posterior mean, 95% credible interval, and posterior inclusion probability, by quantile and predictor.

Equation	tau	Eff. N	Predictor	Mean	95% CI	PIP
GSE-CI	0.10	13	OIL	0.01	(-0.07, 0.12)	0.05
GSE-CI	0.10	13	MPR	0.01	(-0.08, 0.10)	0.05
GSE-CI	0.10	13	TBR	-0.03	(-0.22, 0.07)	0.09
GSE-CI	0.50	66	OIL	0.05	(-0.03, 0.18)	0.11
GSE-CI	0.50	66	MPR	0.00	(-0.07, 0.07)	0.04
GSE-CI	0.75	33	OIL	0.30	(0.15, 0.45)	0.98
GSE-CI	0.75	33	MPR	-0.01	(-0.10, 0.07)	0.05
GSE-CI	0.90	13	OIL	0.36	(0.05, 0.54)	0.94
GSE-CI	0.90	13	MPR	-0.02	(-0.18, 0.08)	0.09
GSE-FSI	0.10	8	TBR	0.24	(-0.05, 0.51)	0.64
GSE-FSI	0.10	8	MPR	-0.05	(-0.35, 0.07)	0.17
GSE-FSI	0.10	8	INF	0.16	(-0.06, 0.57)	0.41
GSE-FSI	0.50	42	TBR	0.02	(-0.07, 0.13)	0.06
GSE-FSI	0.90	8	OIL	0.24	(-0.04, 0.53)	0.63
GSE-FSI	0.90	8	MPR	-0.05	(-0.30, 0.07)	0.16

Note: table shows the predictors of substantive interest at each quantile; full results for all six predictors at all five quantiles are in the companion figshare deposit. PIP values above 0.95 are treated as very strong evidence of inclusion, 0.75–0.95 as substantial, 0.50–0.75 as weak-to-moderate, below 0.50 as weak evidence against inclusion (Section 3.2.3). Source: authors' estimation.

4.5 Impulse Response Functions and Forecast Error Variance Decomposition

Under the corrected recursive ordering (Section 3.2.4: oil, inflation, policy rate, Treasury bill rate, exchange rate, equity return), the orthogonalised impulse response of GSE-CI returns to a one-standard-deviation policy rate shock is small in magnitude and alternates in sign across the twelve-month horizon rather than showing a sustained directional effect (Fig. 4). We also report the generalised impulse response function (GIRF), which is invariant to ordering, and the impulse response under the original, policy-rate-first ordering used in earlier drafts, alongside the corrected result; for GSE-CI the three are broadly similar in shape, but for GSE-FSI the sign of the impact response differs between the corrected ordering and the GIRF, illustrating that the ordering assumption is not innocuous for this system and should be reported alongside, not instead of, the order-invariant robustness check.

Forecast error variance decomposition (Table 7, Fig. 5) shows the policy rate shock accounts for 0.5 percent of GSE-CI return variance at a one-month horizon, rising to 3.1 percent at twelve months; for GSE-FSI the corresponding range is 0.7 percent at one month to 7.4 percent at twelve months, a larger share than for GSE-CI, and higher than reported for either index in earlier drafts that used the uncorrected ordering and, for GSE-FSI, the interpolated series. The equity index's own shock continues to account for the majority of its variance even at twelve months (76.1 percent for GSE-CI, 54.0 percent for GSE-FSI). Among non-own shocks, oil price shocks account for the largest share of twelve-month GSE-CI variance (7.0 percent), while for GSE-FSI, inflation (8.9 percent), the Treasury bill rate (9.8 percent), and the exchange rate (9.7 percent) all contribute more than the policy rate at the twelve-month horizon. We report both the corrected-ordering and original-ordering FEVD tables in full for comparison in the companion figshare deposit.

Table 7
Forecast error variance decomposition (%), corrected ordering, by horizon.

Equation	Horizon (months)	Oil	Inflation	MPR	TBR	EXR	Own (EQ)
GSE-CI	1	2.3	0.1	0.5	1.2	1.2	94.7
GSE-CI	3	5.3	0.5	2.2	2.6	1.9	87.5
GSE-CI	6	6.3	0.8	2.8	5.0	4.8	80.2
GSE-CI	12	7.0	2.9	3.1	4.8	6.0	76.1
GSE-FSI (2018–2024)	1	0.0	9.6	0.7	1.8	3.5	84.3
GSE-FSI (2018–2024)	3	2.7	10.0	2.4	3.2	4.0	77.8
GSE-FSI (2018–2024)	6	6.0	9.8	4.1	4.5	9.9	65.8
GSE-FSI (2018–2024)	12	10.3	8.9	7.4	9.8	9.7	54.0

Source: authors' estimation, corrected Cholesky ordering (Section 3.2.4). Original-ordering results are reported in the companion figshare deposit for comparison.

4.6 Wavelet Coherence Results, Benchmarked Against a Red-Noise Null

Raw band-averaged wavelet coherence between the policy rate shock and GSE-CI returns rises with horizon: 0.48 in the short-run band (2–6 months), 0.56 in the medium-run band (6–18 months), and 0.56 in the long-run band (18–36 months) (Fig. 6). Taken at face value, this reproduces the pattern reported in earlier drafts. However, once benchmarked against a Monte Carlo null of 150 AR(1) surrogate pairs matched to each series' own persistence (Section 3.2.5), the share of coherence that clears the 95th-percentile significance surface, restricted to points outside the cone of influence, is low and non-monotonic across bands: 3.4 percent short-run, 7.4 percent medium-run, and 1.2 percent long-run. This is close to the roughly 5 percent false-positive rate one would expect by construction under a true null of

no relationship, and it does not rise monotonically with horizon the way the raw coherence does. We read this as evidence that the raw monotonic rise reported previously is substantially, though perhaps not entirely, an artefact of the scale-dependent smoothing the coherence estimator requires, rather than evidence of strengthening policy-equity co-movement at longer horizons. For GSE-FSI, restricted to its verified 83-month sample, raw short- and medium-run coherence are 0.50 and 0.51 respectively, with 3.3 percent and 1.9 percent of points clearing the significance surface; the long-run band cannot be estimated on this shorter sample because the cone of influence leaves too little usable area at that scale (Section 3.2.5).

5. Discussion

5.1 A More Modest Role for the Policy Rate Than Commonly Assumed

Across every method employed, Toda-Yamamoto causality with false discovery rate correction, Bayesian quantile regression with embedded variable selection, and forecast error variance decomposition, the domestic policy rate shows a smaller and less consistent relationship with Ghanaian equity returns than a discount-rate-channel prior would suggest, and this conclusion is, if anything, reinforced once the corrections in this revision are applied: no policy-rate causality result approaches significance after multiple-testing correction, and PIP for the policy rate shock does not exceed 0.23 at any quantile for either index. We do not read this as evidence that monetary policy is irrelevant to Ghana's financial system; rather, it is consistent with a market too thin, too concentrated, and too dominated by idiosyncratic firm- and liquidity-level dynamics for the policy rate specifically to leave a statistically detectable, consistent fingerprint on aggregate index returns at monthly frequency. The two stronger and more consistent signals we do find, for money-supply growth (subject to the trend-fitting caveat in Section 4.3) and, in the corrected data, for the exchange rate specifically for GSE-FSI, suggest that liquidity conditions and currency conditions, rather than the headline policy rate itself, may be the more proximate channels through which monetary conditions reach Ghanaian equities.

5.2 Oil Prices as a Driver of Upper-Quantile GSE-CI Returns

The finding that oil-price shocks carry very strong posterior support for GSE-CI returns specifically at upper (bullish) quantiles has received limited attention in the Ghana-specific equity literature, which has tended to focus on the policy rate and the exchange rate; we frame this as an under-explored relationship rather than a claim of priority, since the broader literature on commodity-linked frontier equity markets is substantial. Ghana has been a net oil producer since the Jubilee field began production in December 2010, and the finding that global oil-price strength is most closely associated with the market's best-performing states, rather than exerting a uniform effect throughout the distribution, is consistent with an interpretation in which oil-revenue-driven optimism about government finances and the broader economy amplifies bullish sentiment specifically, without necessarily providing equivalent downside protection during bearish episodes. This quantile-specific pattern also helps reconcile an apparent

tension in our own results: oil shows no significant Toda-Yamamoto causal relationship with GSE-CI in either direction ($p = 0.696$) and a moderate forecast error variance share (7.0 percent at twelve months), figures that would look unremarkable in a mean-based framework, precisely because its explanatory power is concentrated at upper quantiles rather than distributed across the whole return distribution; a method that only examines the conditional mean or an average causal relationship will understate a channel that operates specifically, and strongly, in the upper tail.

5.3 A Funding-Cost Channel for GSE-FSI, More Tentative Than Previously Reported

The Treasury bill rate carries its largest posterior support for GSE-FSI at the tenth percentile (PIP = 0.64), consistent with a funding-cost channel concentrated in bearish, financial-sector-stress states. We report this more cautiously than earlier drafts of this paper did. The estimate now rests on a clean, verified 83-month GSE-FSI sample rather than a series contaminated by four years of interpolated, artificially smooth returns, on a symmetric quantile grid that does not push estimation into a region the data cannot support, on a winsorized Treasury bill rate series that limits the influence of one or two extreme 2022–2023 observations, and on posterior inclusion probabilities drawn from the same sampler that produces the coefficients rather than a separate calculation. Under all of these corrections, the finding survives, but at “weak-to-moderate” rather than “very strong” evidential strength, and on an effective sample of roughly eight observations. We read this as a plausible funding-cost channel worth further investigation with a longer verified sample, not as a settled result.

5.4 Wavelet Coherence: A Weak and Non-Monotonic Signal, Not a Clear Horizon-Dependence Finding

Earlier drafts of this paper reported that co-movement between the policy rate shock and GSE-CI returns strengthens monotonically with horizon, and treated this as a substantive finding distinguishing Ghana from advanced-economy results showing medium-run dominance. That reading does not survive contact with a significance test. Raw band-averaged coherence does rise with horizon, but the wavelet smoothing operator that makes coherence well defined is expected to produce exactly this pattern under a null of no relationship, because it averages over fewer effective degrees of freedom at longer scales. Once we benchmark observed coherence against an AR(1) red-noise Monte Carlo null, the share of coherence that clears the 95 percent significance surface is low, close to the false-positive rate expected by construction, and does not rise monotonically with horizon (3.4, 7.4, and 1.2 percent for GSE-CI’s short-, medium-, and long-run bands respectively). We therefore no longer characterise horizon-increasing coherence as a finding of this paper. The most defensible reading of the wavelet evidence is that policy-equity co-movement in Ghana is weak and does not show a clear, statistically distinguishable concentration at any particular horizon band, for either index, over the sample available; this is itself informative, since it means claims of horizon-specific transmission in this market should be treated with caution until tested against an appropriate null, a methodological point that may generalise to similar studies of other frontier markets.

5.5 Practical and Policy Implications

These findings carry distinct implications for three audiences with a stake in Ghana's equity market: the Bank of Ghana, macroprudential regulators, and investors.

For monetary policymakers, the central implication is a cautionary one. A substantial body of work in advanced-economy settings treats equity prices as a fast, informative channel through which policy stances are transmitted and can be read back by the central bank (Bernanke & Kuttner, 2005; Rigobon & Sack, 2004). Our results suggest this channel operates only weakly, if at all, in Ghana's equity market: policy rate shocks show no causal relationship with either index that survives correction for multiple testing, weak posterior support throughout the quantile regressions, and a forecast error variance contribution that remains below 4 percent for GSE-CI and below 8 percent for GSE-FSI at a twelve-month horizon. This is broadly consistent with the view that monetary transmission through asset prices depends on financial market depth and credibility (Mishkin, 1996), both of which remain constrained in Ghana's frontier setting. Practically, this means the Bank of Ghana should not expect the policy rate to function as a reliable lever for managing equity market conditions, nor should it treat GSE movements as a dependable real-time signal of how a given rate decision is being received by market participants.

For macroprudential regulators, the more policy-relevant signal in our results is the concentration of Treasury-bill-rate sensitivity in the lower quantiles of GSE-FSI returns, though we now report this as tentative evidence rather than a strong finding (Section 5.3). This is consistent with the broader macroprudential literature's emphasis on funding-cost and liquidity channels as sources of financial-sector fragility, distinct from solvency risk per se (Claessens et al., 2013; Cecchetti & Kohler, 2014). Because this sensitivity, if real, is state-dependent rather than uniform across market conditions, standard stress-testing exercises that examine only average-case rate sensitivity may understate the risk that short-term funding-cost increases pose specifically during periods when the financial sector is already under pressure, which argues for macroprudential monitoring frameworks that track banking-sector equity valuations conditional on prevailing market stress, alongside the tentative nature of the current estimate.

For investors and portfolio managers, our results point away from the policy rate and toward two more specific risk factors. The strong upper-quantile association between oil price shocks and GSE-CI returns suggests that global commodity price risk, rather than domestic monetary policy risk, is the more material factor for anyone holding broad Ghanaian equity exposure during favourable market conditions. For financial-sector-specific exposure, the tentative concentration of Treasury-bill-rate sensitivity in bearish states implies that short-term funding-cost indicators may offer more useful downside risk signals for GSE-FSI than the policy rate itself, a distinction that would not be visible from a conventional mean-based risk model (Jansen & Tsai, 2010; Kontonikas et al., 2013), though we emphasise this specific estimate should be treated with the same caution as in Section 5.3. More generally, the weak and inconsistent policy-rate effect we document across quantiles reinforces the case, made elsewhere

in the quantile and model-averaging literature, for distributional rather than average-effect risk assessment in frontier markets characterised by regime-dependent dynamics (Koop & Korobilis, 2012).

5.6 Limitations and Directions for Future Research

Several limitations warrant explicit acknowledgment, in addition to the corrections made in this revision. First, although we have excluded the deterministic 2014–2017 GSE-FSI segment entirely, three other series retain a smaller number of constructed observations: 11 of 132 exchange-rate months, 8 of 132 Treasury bill rate months, and 4 of 132 money-supply months, concentrated in 2024 (Table 1). Results involving these series for late 2024 specifically should be read with this in mind. Second, the GSE-FSI sample used throughout this paper (83 return observations) is markedly shorter than the GSE-CI sample (131 observations), which is why the long-run wavelet band and the $\tau = 0.05$ quantile are not reported for GSE-FSI; extending the verified GSE-FSI series backward, should independently sourced monthly values become available for 2014–2017, would allow these to be revisited. Third, our sample, while covering a rich set of macro-financial episodes, spans at most eleven years and, for GSE-FSI, seven, which is modest by the standards of time-series macro-finance and limits the precision of several estimators beyond the wavelet long-run band specifically. Fourth, the variable-selection exercise is conducted over a fixed candidate set of six macro-financial predictors; extending this to a larger candidate set, including firm-level and sectoral variables, is a natural direction for future work. Fifth, while we have applied Benjamini-Hochberg correction to the Toda-Yamamoto family of tests and embedded shrinkage-based selection in the quantile regressions, the wavelet band comparisons and the impulse-response and variance-decomposition results are not subject to a formal multiple-testing correction beyond the Monte Carlo significance benchmark described in Section 3.2.5 and the ordering-sensitivity check reported in Section 4.5; readers should weigh the full set of reported results, rather than any single estimate in isolation, accordingly. Finally, our Cholesky identification in the VAR, even under the corrected ordering used in this revision, imposes a specific contemporaneous structure; we report the generalised, order-invariant impulse response and the original-ordering result alongside it precisely so this sensitivity is visible, and future work testing additional identification schemes, once a longer post-2014 sample is available, would be a natural extension.

6. Conclusion

This study re-examines monetary policy transmission to Ghanaian equity returns using monthly data spanning January 2014 to December 2024 for GSE-CI and February 2018 to December 2024 for GSE-FSI, the latter window restricted specifically to exclude index values that would otherwise be constructed by interpolation rather than observed. Contrary to a discount-rate-channel prior, we find limited, statistically inconsistent evidence that the domestic policy rate itself drives equity returns in this market, whether assessed through Granger-causal ordering with multiple-testing correction, distributional quantile regression with embedded variable selection, or forecast error variance decomposition under a corrected identification ordering. Instead, our results point to two more specific channels, reported with different degrees of confidence: global oil-price shocks are strongly associated with bullish upper-

quantile GSE-CI returns, a result that survives every robustness check applied in this revision; and short-term Treasury bill rates carry a state-dependent, but more tentative than previously reported, relationship with GSE-FSI returns concentrated in the lower (bearish) quantiles. Wavelet coherence, once benchmarked against an appropriate statistical null, does not support a clear claim about how policy-equity co-movement varies with horizon for either index; this null result is itself a corrected contribution of this revision, since the raw, unbenchmarked pattern reported in earlier work does not survive testing. Collectively, these findings suggest that policymakers, investors, and regulators focused on Ghana's equity market may derive more diagnostic value from monitoring global commodity prices and short-term funding costs than from the policy rate in isolation, and that studies of monetary transmission in thin, structurally discontinuous frontier markets benefit from index-construction and multiple-testing safeguards that are not yet standard practice in this literature.

Declarations

CRedit authorship contribution statement

Conceptualization, I.N. and G.O.F.; methodology, I.N. and G.O.F.; software, G.O.F.; validation, I.N. and G.O.F.; formal analysis, I.N. and G.O.F.; investigation, I.N. and G.O.F.; resources, I.N. and G.O.F.; writing, original draft preparation, I.N. and G.O.F.; writing, review and editing, I.N. and G.O.F.; supervision, G.O.F. All authors have read and agreed to the publication of the manuscript.

Conflicts of Interest:

The authors declare no conflicts of interest.

Ethics approval statement: Not applicable.

Funding:

This research received no external funding.

Author Contribution

Conceptualization, I.N. and G.O.F.; methodology, I.N. and G.O.F.; software, G.O.F.; validation, I.N. and G.O.F.; formal analysis, I.N. and G.O.F.; investigation, I.N. and G.O.F.; resources, I.N. and G.O.F.; writing, original draft preparation, I.N. and G.O.F.; writing, review and editing, I.N. and G.O.F.; supervision, G.O.F. All authors have read and agreed to the publication of the manuscript.

Acknowledgement

The authors would like to thank the Kwame Nkrumah University of Science and Technology, Kumasi (KNUST), and the Department of Forest Engineering, Forest Management Planning, and Terrestrial Measurements, Faculty of Silviculture and Forest Engineering, Transilvania University of Brasov, for providing some of the resources needed for this study.

Data Availability

The dataset and Python scripts used for data analysis and visualization are available as supplementary files on the figshare repository at <https://doi.org/10.6084/m9.figshare.32568465>.

References

1. Acharya VV, Pedersen LH, Philippon T, Richardson M (2017) Measuring systemic risk. *Rev Financial Stud* 30(1):2–47. <https://doi.org/10.1093/rfs/hhw088>
2. Adam AM, Tweneboah G (2008) *Macroeconomic factors and stock market movement: Evidence from Ghana* (MPRA Paper No. 11256). Munich Personal RePEc Archive. <https://mpra.ub.uni-muenchen.de/11256/>
3. Adjasi CKD, Biekpe NB (2005) Stock market returns and exchange rate dynamics in selected African countries: A bivariate analysis. *Afr Financ J* 7(2):17–28
4. Adrian T, Shin HS (2010) Liquidity and leverage. *J Financ Intermediation* 19(3):418–437. <https://doi.org/10.1016/j.jfi.2008.12.002>
5. Aguiar-Conraria L, Soares MJ (2011) Oil and the macroeconomy: Using wavelets to analyze old issues. *Empirical Economics* 40(3):645–655. <https://doi.org/10.1007/s00181-010-0371-x>
6. Alagidede P (2011) Return behaviour in Africa's emerging equity markets. *Q Rev Econ Finance* 51(2):133–140. <https://doi.org/10.1016/j.qref.2011.01.004>
7. Alagidede P, Panagiotidis T (2009) Modelling stock returns in Africa's emerging equity markets. *Int Rev Financial Anal* 18(1–2):1–11. <https://doi.org/10.1016/j.irfa.2009.02.001>
8. Bai J, Perron P (2003) Computation and analysis of multiple structural change models. *J Appl Econom* 18(1):1–22. <https://doi.org/10.1002/jae.659>
9. Bank of Ghana (2023) *Annual report and financial statements*. Bank of Ghana
10. Bekaert G, Harvey CR, Lundblad C (2011) Financial openness and productivity. *World Dev* 39(1):1–19. <https://doi.org/10.1016/j.worlddev.2010.06.016>
11. Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J Roy Stat Soc: Ser B (Methodol)* 57(1):289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
12. Bernanke BS, Blinder AS (1992) The federal funds rate and the channels of monetary transmission[1]. *Am Econ Rev* 82(4):901–921

13. Bernanke BS, Gertler M (1995) Inside the black box: The credit channel of monetary policy transmission[2]. *J Economic Perspect* 9(4):27–48. <https://doi.org/10.1257/jep.9.4.27>
14. Bernanke BS, Kuttner KN (2005) What explains the stock market's reaction to Federal Reserve policy? *J Finance* 60(3):1221–1257. <https://doi.org/10.1111/j.1540-6261.2005.00760.x>
15. Bjørnland HC, Leitemo K (2009) Identifying the interdependence between US monetary policy and the stock market. *J Monet Econ* 56(2):275–282. <https://doi.org/10.1016/j.jmoneco.2008.12.001>
16. Brunnermeier MK, Dong GN, Palia D (2020) Banks' noninterest income and systemic risk. *Rev Corp Finance Stud* 9(2):229–255. <https://doi.org/10.1093/rcfs/cfaa006>
17. Brunnermeier MK, Pedersen LH (2009) Market liquidity and funding liquidity. *Rev Financial Stud* 22(6):2201–2238. <https://doi.org/10.1093/rfs/hhn098>
18. Cecchetti SG, Kohler M (2014) When capital adequacy and interest rate policy are substitutes (and when they are not)[3]. *Int J Cent Bank* 10(3):205–231
19. Claessens S, Ghosh SR, Mihet R (2013) Macro-prudential policies to mitigate financial system vulnerabilities. *J Int Money Finance* 39:153–185. <https://doi.org/10.1016/j.jimonfin.2013.06.023>
20. Fama EF (1970) Efficient capital markets: A review of theory and empirical work[4]. *J Finance* 25(2):383–417. <https://doi.org/10.1111/j.1540-6261.1970.tb00518.x>
21. Fama EF (1981) Stock returns, real activity, inflation, and money[5][6]. *Am Econ Rev* 71(4):545–565
22. Frimpong JM (2008) Capital market efficiency: An analysis of weak-form efficiency on the Ghana Stock Exchange[7]. *J Money Invest Bank* 5:5–13
23. Frimpong JM, Oteng-Abayie EF (2006) *Bivariate causality analysis between FDI inflows and economic growth in Ghana* (MPRA Paper No. 351). Munich Personal RePEc Archive. <https://mpra.ub.uni-muenchen.de/351/>
24. George EI, McCulloch RE (1993) Variable selection via Gibbs sampling. *J Am Stat Assoc* 88(423):881–889. <https://doi.org/10.1080/01621459.1993.10476353>
25. Geske R, Roll R (1983) The fiscal and monetary linkage between stock returns and inflation. *J Finance* 38(1):1–33. <https://doi.org/10.1111/j.1540-6261.1983.tb03623.x>
26. Geweke J (1982) Measurement of linear dependence and feedback between multiple time series. *J Am Stat Assoc* 77(378):304–313. <https://doi.org/10.1080/01621459.1982.10477803>
27. Gordon MJ (1962) The investment, financing, and valuation of the corporation. Richard D. Irwin
28. Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3):424–438. <https://doi.org/10.2307/1912791>
29. Grinsted A, Moore JC, Jevrejeva S (2004) Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Process Geophys* 11(5/6):561–566. <https://doi.org/10.5194/npg-11-561-2004>
30. Gürkaynak RS, Sack B, Swanson E (2005) The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *Am Econ Rev* 95(1):425–436. <https://doi.org/10.1257/0002828053828446>

31. Jansen DW, Tsai C-L (2010) Monetary policy and stock returns: Financing constraints and asymmetries in bull and bear markets. *J Empir Finance* 17(5):981–990. <https://doi.org/10.1016/j.jempfin.2010.08.002>
32. Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–292. <https://doi.org/10.2307/1914185>
33. Koenker R, Bassett G (1978) Regression quantiles. *Econometrica* 46(1):33–50. <https://doi.org/10.2307/1913643>
34. Kontonikas A, MacDonald R, Saggiu A (2013) Stock market reaction to Fed funds rate surprises: State dependence and the financial crisis. *J Banking Finance* 37(10):4025–4037. <https://doi.org/10.1016/j.jbankfin.2013.06.010>
35. Koop G, Korobilis D (2012) Forecasting inflation using dynamic model averaging. *Int Econ Rev* 53(3):867–886. <https://doi.org/10.1111/j.1468-2354.2012.00704.x>
36. Korobilis D (2017) Quantile regression forecasts of inflation under model uncertainty. *Int J Forecast* 33(1):11–20. <https://doi.org/10.1016/j.ijforecast.2016.07.005>
37. Kozumi H, Kobayashi G (2011) Gibbs sampling methods for Bayesian quantile regression. *J Stat Comput Simul* 81(11):1565–1578. <https://doi.org/10.1080/00949655.2010.496117>
38. Kutun AM, Aksoy T (2003) Public information arrival and the Fisher effect in emerging markets: Evidence from stock and bond markets in Turkey. *J Financial Serv Res* 23(3):225–239. <https://doi.org/10.1023/A:1024671612784>
39. Longin F, Solnik B (2001) Extreme correlation of international equity markets. *J Finance* 56(2):649–676. <https://doi.org/10.1111/0022-1082.00340>
40. Mishkin FS (1996) *The channels of monetary transmission: Lessons for monetary policy* (NBER Working Paper No. 5464). National Bureau of Economic Research. <https://doi.org/10.3386/w5464>
41. Perron P (1989) The great crash, the oil price shock, and the unit root hypothesis. *Econometrica* 57(6):1361–1401. <https://doi.org/10.2307/1913712>
42. Pesaran HH, Shin Y (1998) Generalized impulse response analysis in linear multivariate models. *Econ Lett* 58(1):17–29. [https://doi.org/10.1016/S0165-1765\(97\)00214-0](https://doi.org/10.1016/S0165-1765(97)00214-0)
43. Rigobon R, Sack B (2004) The impact of monetary policy on asset prices. *J Monet Econ* 51(8):1553–1575. <https://doi.org/10.1016/j.jmoneco.2004.02.004>
44. Rodrigues J, Andrade A (2013) Instantaneous Granger causality with the Hilbert-Huang transform. *International Scholarly Research Notices*, 2013, Article 374064. <https://doi.org/10.1155/2013/374064>
45. Romer CD, Romer DH (2004) A new measure of monetary shocks: Derivation and implications. *Am Econ Rev* 94(4):1055–1084. <https://doi.org/10.1257/0002828042002651>
46. Rua A, Nunes LC (2009) International comovement of stock market returns: A wavelet analysis. *J Empir Finance* 16(4):632–639. <https://doi.org/10.1016/j.jempfin.2009.02.002>

47. Said SE, Dickey DA (1984) Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika* 71(3):599–607. <https://doi.org/10.1093/biomet/71.3.599>
48. Taylor JB (1995) The monetary transmission mechanism: An empirical framework. *J Economic Perspect* 9(4):11–26. <https://doi.org/10.1257/jep.9.4.11>
49. Toda HY, Yamamoto T (1995) Statistical inference in vector autoregressions with possibly integrated processes. *J Econ* 66(1–2):225–250. [https://doi.org/10.1016/0304-4076\(94\)01616-8](https://doi.org/10.1016/0304-4076(94)01616-8)
50. Torrence C, Compo GP (1998) A practical guide to wavelet analysis. *Bull Am Meteorol Soc* 79(1):61–78. [https://doi.org/10.1175/1520-0477\(1998\)079%3C0061:APGTWA%3E2.0.CO;2](https://doi.org/10.1175/1520-0477(1998)079%3C0061:APGTWA%3E2.0.CO;2)
51. World Bank (2023) Ghana economic update: Restoring stability and resilience. World Bank
52. Yu K, Moyeed RA (2001) Bayesian quantile regression. *Stat Probab Lett* 54(4):437–447. [https://doi.org/10.1016/S0167-7152\(01\)00124-9](https://doi.org/10.1016/S0167-7152(01)00124-9)
53. Zivot E, Andrews DWK (1992) Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *J Bus Economic Stat* 10(3):251–270. <https://doi.org/10.1080/07350015.1992.10509904>

Figures

Conceptual Framework: Monetary Policy Transmission to Ghanaian Equity Returns




Figure 1

Conceptual framework linking monetary policy rate shocks to Ghanaian equity returns through three competing transmission channels: a discount-rate/valuation channel, an exchange-rate and imported-inflation channel, and a financial-sector net-interest-margin channel relevant to bank-heavy indices such as GSE-FSI, alongside global oil price shocks and domestic money supply growth as competing sources of explanatory power. *Note: The relative strength of each channel is treated as an empirical question rather than assumed a priori, and is allowed to vary across the position in the return distribution (quantile) and across time horizons (frequency band) before manifesting in returns on the GSE Composite Index (GSE-CI) and GSE Financial Stock Index (GSE-FSI).*




Figure 2

GSE-CI (2014-2024) and GSE-FSI (2018-2024 verified window) monthly closing levels, with major crisis episodes shaded.




Figure 3

Bayesian quantile regression standardised coefficients (posterior mean $\pm$ 95% credible interval) across quantiles: OIL and MPR for GSE-CI; TBR and MPR for GSE-FSI.




Figure 4

Orthogonalised impulse response of GSE-CI and GSE-FSI returns to a one-standard-deviation MPR shock, 0-12 month horizon, corrected ordering, with the generalised (order-invariant) impulse response shown for comparison.




Figure 5

Forecast error variance decomposition of GSE-CI and GSE-FSI returns (%), 1/3/6/12-month horizons, corrected ordering.

Wavelet coherence, MPR shock vs equity return (white contour = 95% AR(1) Monte Carlo significance; hatched region = cone of influence)




Figure 6

Wavelet coherence between the MPR shock and GSE-CI / GSE-FSI returns, with the 95% Monte Carlo AR(1) red-noise significance contour overlaid and the cone of influence shaded. Dashed lines mark the 6- and 18-month band boundaries; the 18-36 month band is not shown for GSE-FSI.