



Article

# Conditional-Mean Predictive Precedence and Information Concentration in a Commodity-Dependent Equity Market: Evidence from Petrobras and the Ibovespa, 2005–2026

Alejandro Pérez-y-Soto-Domínguez <sup>1</sup>, Juan Manuel Candelo-Viáfara <sup>2,\*</sup> and Edwin Arango-Espinal <sup>3</sup>

<sup>1</sup> Department of Economics, Faculty of Human and Economic Sciences, Universidad Nacional de Colombia, Sede Medellín, Carrera 65 #59A-110, Medellín 050034, Colombia; aperezy@unal.edu.co

<sup>2</sup> Department of Accounting and Finance, Faculty of Administration Sciences, Universidad del Valle, Sede San Fernando Campus, Calle 4B #36-00, Santiago de Cali 760043, Colombia

<sup>3</sup> Sede Regional Caicedonia, Universidad del Valle, Caicedonia Campus, Caicedonia 762540, Colombia; edwin.arango@correounivalle.edu.co

\* Correspondence: juan.candelo@correounivalle.edu.co

## Abstract

This paper examines whether standard price-discovery measures can reliably identify directional predictive precedence in a highly correlated commodity-equity system. Using 21 years of daily data for Petrobras and the Ibovespa (2005–2026), the study separates a measurement problem in forecast error variance decomposition from the reduced-form question of directional predictability in the conditional mean. The empirical strategy combines Monte Carlo simulation, generalized and Cholesky forecast error variance decompositions, full-sample and rolling-window Granger causality tests, a continuous Granger Leadership Index, Gaussian mixture regime classification, robustness checks, and out-of-sample forecasting validation. The results show that Cholesky-based FEVDs can be systematically misleading in high-correlation settings: at the observed contemporaneous correlation, generalized FEVD symmetry is mechanically induced by row normalization, while Cholesky attribution changes sharply under alternative orderings. By contrast, first-moment predictability reveals a directional asymmetry from Petrobras to the Ibovespa, interpreted as conditional-mean predictive precedence rather than structural informed trading or definitive price discovery. This asymmetry survives alternative lag structures, weekly aggregation, univariate GARCH filtering, within-dataset proxy controls, and a stylized equal-weight ex-Petrobras benchmark. Rolling evidence further identifies five persistent predictive regimes that alternate between firm-led, neutral, and macro-dominant states, indicating that firm-index predictive relations are regime dependent rather than static. Out-of-sample forecasting shows that the identified predictive precedence does not generate exploitable one-step-ahead gains (RMSE ratio = 1.002, OOS- $R^2$  =  $-0.003$ , DM  $p$  = 0.451), thereby delimiting the economic scope of the findings. Overall, the results support a reduced-form interpretation of Petrobras–Ibovespa predictive dynamics and highlight the need to distinguish variance connectedness from conditional-mean predictive content when contemporaneous correlation is high.



Academic Editors: Periklis Gogas and Babacar Sène

Received: 22 April 2026

Revised: 4 June 2026

Accepted: 25 June 2026

Published: 9 July 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

**Keywords:** conditional-mean predictability; predictive precedence; Granger causality; forecast error variance decomposition; commodity-equity markets

**JEL Classification:** C32; C38; G14; G15; Q43

## 1. Introduction

Price discovery in commodity-dependent equity markets raises a central identification problem: when a dominant commodity-linked firm and its benchmark index move together with high contemporaneous correlation, to what extent can standard variance-decomposition methods distinguish variance connectedness from directional predictive content? This question is particularly relevant in the Petrobras–Ibovespa system. Petrobras is highly sensitive to oil-related news and is also a major constituent of the Brazilian benchmark; therefore, observed co-movement may reflect reduced-form predictive precedence, common exposure to external shocks, or mechanical index-composition effects. This paper argues that these possibilities must be separated analytically. The first issue is a measurement problem: whether Cholesky-ordered forecast error variance decompositions (FEVDs) can misattribute directional dominance in highly correlated systems. The second issue is an empirical one: whether Petrobras or the Ibovespa exhibits conditional-mean predictive precedence, under which market conditions, and through which economically plausible channels.

Brazil provides an especially suitable empirical setting for addressing this problem. As a large emerging market with strong exposure to commodity cycles, exchange-rate movements, sovereign-risk repricing, and domestic political shocks, Brazil concentrates several overlapping disturbances that make identification difficult. Petrobras adds a further layer of relevance: it is simultaneously a commodity-sensitive stock and a large-weight index component. This makes Brazil a demanding empirical case rather than merely a convenient one. If reduced-form predictive precedence can be detected in a setting where common shocks, oil dependence, macro-financial spillovers, and benchmark composition interact closely, then the resulting framework may be informative for other commodity-equity systems. In this sense, the Brazilian case is theoretically and methodologically useful precisely because it is empirically difficult.

The theoretical intuition comes from the price discovery and market microstructure literature. Kyle (1985) and Glosten and Milgrom (1985) show that informational advantages may be reflected in trading behavior and price formation, while Hasbrouck (1991, 1995) provides a VAR-based empirical framework through which such dynamics can be studied. Yet that same literature also implies a cautionary point: empirical inference is highly sensitive to identification. In systems with strong contemporaneous dependence, Cholesky orthogonalization imposes an ordering structure that may mechanically assign variance shares to the variable placed first, even when the underlying data-generating process does not justify such asymmetry. As a result, second-moment measures of interconnectedness should not be treated as interchangeable with first-moment measures of predictive content. Granger causality evaluates directional predictability in the conditional mean, whereas FEVD-based measures summarize the allocation of forecast error variance; when contemporaneous correlation is high, the two can lead to sharply different conclusions. Accordingly, throughout this paper, the term leadership is used in a reduced-form predictive sense and should not be interpreted as evidence of structural informed trading, causal dominance, or definitive same-exchange price discovery.

This distinction is especially consequential in firm-index systems exposed to commodity shocks, because co-movement may arise through several economically plausible transmission channels that do not imply the same form of predictive relation. This paper emphasizes four channels. First, a commodity-fundamental channel: oil-price shocks, reserve news, export expectations, and Petrobras-specific earnings information may be incorporated into Petrobras and then be reflected in the broader index. Second, a macro-financial channel: country-risk repricing, exchange-rate shocks, domestic political news, and monetary conditions may affect the Ibovespa first and only later propagate to Petrobras as part

of a market-wide repricing of Brazilian equities. Third, a portfolio and index-composition channel: because Petrobras carries substantial weight in the Ibovespa, part of the measured predictability may be mechanical, driven by benchmark arithmetic, passive flows, or rebalancing effects rather than firm-level information. Fourth, an information-processing and liquidity channel: depending on market stress, trading intensity, and relative depth, either the single stock or the broad index may absorb and disseminate information more rapidly. Since the empirical design is reduced-form, these channels are interpreted as economically grounded mechanisms consistent with the data rather than as fully structurally identified causal effects.

The paper builds on the idea that predictive relations are not necessarily fixed, but state-dependent. Back et al. (2000) show that competition among informed traders can generate outcomes ranging from relatively symmetric interaction to concentrated leadership, which provides a useful conceptual analogy for the paper's empirical regime approach. Rather than assuming a stable hierarchy between Petrobras and the Ibovespa, the paper treats predictive precedence as something that may alternate between more balanced configurations and more concentrated ones. This view is consistent with the broader literature showing that oil-equity linkages and predictive relations are asymmetric and time varying, particularly in emerging and commodity-exposed markets (Basher et al., 2012; Antonakakis et al., 2017; Degiannakis et al., 2018; Shi et al., 2018; Hong et al., 2024). It is also consistent with Lo's (2004) adaptive markets hypothesis, under which market efficiency may vary over time rather than remain constant.

Against this background, the paper addresses three gaps in the literature. First, although Cholesky-based FEVDs are widely used in price discovery and connectedness research, the literature rarely quantifies how the ordering artifact scales with contemporaneous correlation. Second, the rolling Granger causality literature lacks a continuous directional measure capable of tracking the intensity of conditional-mean predictive precedence over time in a parsimonious way. Third, the commodity-equity literature has not sufficiently characterized the recurrent alternation between concentrated and dispersed predictive regimes in firm-index systems, nor documented whether such regimes exhibit persistence beyond their unconditional frequencies.

The paper makes five contributions. First, it provides Monte Carlo evidence, based on 1000 replications and a fully specified bivariate VAR data-generating process, showing that the Cholesky FEVD ordering artifact—measured as the monotone gap between Cholesky and generalized FEVD attribution—increases with contemporaneous correlation. This supplies a diagnostic for empirical studies that rely on variance-decomposition measures in highly correlated systems. Second, it introduces a continuous Granger Leadership Index, defined as the difference between the rolling F-statistics for Petrobras-to-Ibovespa and Ibovespa-to-Petrobras predictive effects; under a fixed rolling window and constant degrees of freedom, this index provides a first-moment measure of directional predictive asymmetry over time. Third, it proposes a data-driven five-regime taxonomy based on Gaussian mixture modeling in the standardized space of bilateral rolling F-statistics, allowing the paper to classify distinct predictive configurations. Fourth, it shows that these regimes are persistent, with self-transition probabilities above their stationary frequencies at both daily and weekly horizons. Fifth, it provides an out-of-sample validation showing that the identified predictive precedence does not generate exploitable one-step-ahead Ibovespa return forecasts (RMSE ratio = 1.002, OOS- $R^2$  = −0.003, DM-HLN  $p$  = 0.451), thereby delimiting the economic scope of the Granger-based findings and aligning them with semi-strong-form market efficiency. Together, these contributions clarify when variance-decomposition evidence is informative and when conditional-mean predictability should be analyzed separately.

These arguments lead to four testable hypotheses.

**H1.** *The measurement hypothesis, states that in bivariate commodity-equity systems, the distortion induced by Cholesky-ordered FEVDs increases with contemporaneous correlation; therefore, directional dominance inferred from Cholesky variance decompositions becomes progressively less reliable as correlation rises.*

**H2.** *The directional predictive-precedence hypothesis, states that Petrobras and the Ibovespa do not exhibit constant symmetric interaction; instead, directional predictive asymmetry emerges intermittently and depends on the dominant source of information. Petrobras is expected to exhibit stronger predictive content when firm-specific and commodity-fundamental news dominates, whereas the Ibovespa is expected to exhibit stronger predictive content when aggregate macro-financial and market-wide shocks dominate.*

**H3.** *The regime-dependence hypothesis, states that predictive precedence is state dependent rather than constant; the Petrobras–Ibovespa system alternates between concentrated predictive regimes and more balanced regimes, consistent with shifts between Stackelberg-like and Nash-like informational configurations.*

**H4.** *The persistence hypothesis, states that once predictive regimes emerge, they persist beyond what would be expected from unconditional frequencies alone, implying that directional predictability is not merely transitory noise but a recurrent market state.*

To test these hypotheses, the paper uses 21 years of daily data for Petrobras and the Ibovespa, from 2005 to 2026, in a multi-stage empirical design. The first stage uses Monte Carlo simulation to isolate the measurement problem and quantify the sensitivity of FEVD-based attribution to contemporaneous correlation and ordering assumptions. The second stage estimates rolling-window Granger causality to identify time-varying predictive relations in the conditional mean. The third stage summarizes these dynamics with the Granger Leadership Index. The fourth stage applies Gaussian mixture modeling to the joint distribution of bilateral rolling F-statistics to classify predictive regimes. Finally, the paper examines regime persistence through transition analysis and validates the results using weekly frequency data, a stylized equal-weight ex-Petrobras benchmark, GARCH-filtered residuals, within-dataset proxy controls, and out-of-sample forecasting. These robustness checks are important in the Brazilian setting because they help distinguish conditional-mean predictive precedence from mechanical benchmark effects, volatility clustering, and within-market common drivers.

The Brazilian case is therefore useful not only because it is economically relevant, but because it directly sharpens the paper’s identification strategy. If Petrobras appears to predict the Ibovespa only because it is a large index constituent, then the apparent firm-to-index relation should weaken once Petrobras is removed from the stylized benchmark or once daily mechanical transmission becomes less important at weekly frequency. If predictive asymmetry survives these exercises, the evidence becomes less consistent with a purely arithmetic explanation. However, the design does not eliminate all competing interpretations: timing differences between ADR and B3 closing prices, external common shocks such as Brent crude, BRL/USD, VIX, sovereign-risk spreads, and domestic monetary conditions, and multivariate volatility dynamics remain relevant limitations. The paper therefore interprets the evidence as reduced-form predictive precedence rather than structural price discovery.

The contribution of the paper is both methodological and economic. Methodologically, it shows when variance decomposition is informative and when it may become misleading in highly correlated systems. Economically, it shows that the Petrobras–Ibovespa relation is not well characterized as a static hierarchy, but as a regime-dependent predictive system shaped by firm-specific information, commodity conditions, index composition, and broader

macro-financial repricing. More broadly, the paper argues that the interaction between a dominant commodity firm and its market benchmark should be analyzed as a dynamic firm-index system in which conditional-mean predictive content varies across market states.

The remainder of the paper is organized as follows. Section 2 presents the empirical methodology, including the VAR framework, generalized and Cholesky FEVDs, rolling Granger causality tests, the Granger Leadership Index, Gaussian mixture regime classification, and the out-of-sample forecasting design. Section 3 describes the data, sample construction, return transformations, and the rationale for using ADR-based Petrobras prices. Section 4 reports the empirical results, including full-sample predictability, FEVD identification tests, Monte Carlo evidence, rolling-regime dynamics, robustness checks, regime persistence, and forecasting validation. Section 5 discusses the economic interpretation of the findings, emphasizing the distinction between variance connectedness, reduced-form predictive precedence, and structural price discovery. Section 6 concludes by summarizing the main contributions, limitations, and avenues for future research.

## 2. Methodology

This section presents the empirical strategy used to study directional predictability and price discovery between Petrobras and the Ibovespa. The methodology combines four complementary components: a bivariate VAR model to characterize dynamic interactions, generalized forecast error variance decomposition to measure connectedness, rolling Granger causality tests to track time variation in predictive leadership, and Gaussian mixture modeling to classify recurrent leadership regimes. Reported estimates and diagnostic outcomes are presented later in the data and testing sections.

### 2.1. VAR Framework

Let  $r_t = (r_{PBR,t}, r_{IBVSP,t})'$  denote the vector of daily log returns for Petrobras and the Ibovespa at time  $t$ . The dynamic relation between both series is modeled through a bivariate VAR( $p$ ):

$$r_t = c + A_1 r_{t-1} + A_2 r_{t-2} + \cdots + A_p r_{t-p} + \varepsilon_t \quad (1)$$

where  $r_t$  is the  $2 \times 1$  vector of returns;  $c$  is a  $2 \times 1$  vector of intercepts;  $A_1, \dots, A_p$  are  $2 \times 2$  coefficient matrices capturing lagged interactions;  $p$  is the lag length; and  $\varepsilon_t$  is a  $2 \times 1$  innovation vector with zero mean and covariance matrix  $\Sigma$ . This specification provides the basic reduced-form representation used throughout the paper to assess dynamic dependence, predictive content, and variance transmission.

### 2.2. Generalized Forecast Error Variance Decomposition

To evaluate how much of the forecast error variance of one asset is associated with shocks to the other, the paper uses the order-invariant generalized FEVD proposed by Koop et al. (1996) and Pesaran and Shin (1998):

$$\phi_{ij}(h) = \frac{\sigma_{jj}^{-1} \sum_{s=0}^{h-1} (e_i' \Phi_s \Sigma e_j)^2}{\sum_{s=0}^{h-1} (e_i' \Phi_s \Sigma \Phi_s' e_i)} \quad (2)$$

where  $\phi_{ij}(h)$  is the share of the  $h$ -step-ahead forecast error variance of the variable  $i$  explained by shocks to the variable  $j$ ;  $h$  is the forecast horizon;  $\sigma_{jj}$  is the  $j$ -th diagonal element of  $\Sigma$ ;  $e_i$  and  $e_j$  are selection vectors identifying variables  $i$  and  $j$ ; and  $\Phi_s$  is the moving-average coefficient matrix at horizon  $s$ . This measure is useful because it does not depend on variable ordering, unlike Cholesky orthogonalization. The total connectedness index is then computed as the share of cross-variable variance transmission in the system, following the framework of Diebold and Yilmaz (2012, 2014).

### 2.3. Rolling Granger Causality and Leadership Measure

To capture time variation in predictive relations, the paper estimates rolling-window Granger causality tests on successive subsamples of fixed length, following the rolling-window approach of Nakhli et al. (2022). For each window, bilateral F-statistics are obtained for Petrobras  $\rightarrow$  Ibovespa and Ibovespa  $\rightarrow$  Petrobras. Based on these statistics, the paper defines a continuous Granger Leadership Index:

$$LI_{F,t} = F_{PBR \rightarrow IBVSP,t} - F_{IBVSP \rightarrow PBR,t} \quad (3)$$

where  $LI_{F,t}$  is the leadership index in a rolling window  $t$ ;  $F_{PBR \rightarrow IBVSP,t}$  is the rolling F-statistic testing whether Petrobras Granger-causes the Ibovespa; and  $F_{IBVSP \rightarrow PBR,t}$  is the rolling F-statistic for the reverse direction. Positive values of  $LI_{F,t}$  indicate stronger predictive leadership from Petrobras to the Ibovespa, while negative values indicate stronger predictive leadership from the Ibovespa to Petrobras. Because all windows use the same size and degrees of freedom, the F-statistics are directly comparable over time.

### 2.4. Regime Classification

To classify recurrent leadership patterns, the paper applies Gaussian mixture modeling (Kalliovirta et al., 2016) to standardized bilateral rolling F-statistics. Let  $x_t = (F_{PBR,t}, F_{IBVSP,t})'$ . The mixture density is:

$$p(x_t) = \sum_{k=1}^K \pi_k \mathcal{N}(x_t; \mu_k, \Sigma_k) \quad (4)$$

where  $p(x_t)$  is the unconditional density of observation  $x_t$ ;  $K$  is the number of latent regimes;  $\pi_k$  is the mixing weight of the regime  $k$ , with  $\sum_{k=1}^K \pi_k = 1$ ;  $\mu_k$  is the mean vector of the regime  $k$ ;  $\Sigma_k$  is the covariance matrix of the regime  $k$ ; and  $\mathcal{N}(\cdot)$  denotes the multivariate normal density. This framework is used as a descriptive unsupervised taxonomy of predictive configurations. It does not estimate a structural latent-state data-generating process; rather, it summarizes recurrent patterns ranging from relatively balanced interaction to concentrated predictive dominance by one of the two assets.

## 3. Data

The dataset consists of 5139 daily closing prices for seven assets—Petrobras, Vale, Itaú Unibanco, Bradesco, Ambev, Gerdau, and the Ibovespa—covering 3 January 2005 through 16 April 2026. The core analysis focuses on Petrobras and the Ibovespa, while the remaining assets are used for comparison and robustness purposes. Daily log returns are computed as

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \quad (5)$$

where  $r_t$  is the return at the time  $t$ ,  $P_t$  is the closing price at the time  $t$ , and  $P_{t-1}$  is the closing price at the time  $t - 1$ . This produces  $N = 5138$  return observations. Summary descriptive statistics for these series are reported in Table 1.

Prices are adjusted closing prices for NYSE-listed ADRs obtained from Yahoo Finance, all denominated in U.S. dollars and cross-validated against Bloomberg for 2018–2025. The use of USD-denominated ADRs rather than local BRL-denominated B3 prices is intentional because ADRs incorporate exchange-rate variation and are closer to the information set of international investors. However, this data choice does not provide a clean same-exchange price-discovery test. Because Petrobras ADRs trade in the United States and may close after the Brazilian cash market, daily ADR returns may incorporate information released after the B3 close. Therefore, the sample construction mitigates some comparability concerns but does not fully eliminate timing-related lead-lag effects. The results should accordingly be


interpreted as reduced-form conditional-mean predictive precedence rather than definitive same-exchange price discovery. This choice is nevertheless conservative for the paper's identification strategy because it increases Petrobras volatility and slightly lowers its contemporaneous correlation with the Ibovespa. The descriptive implications of this sample construction are reflected in Table 1. The Ambev outlier on 9 June 2005 reflects an extraordinary dividend adjustment and does not affect the bivariate Petrobras–Ibovespa analysis.

**Table 1.** Descriptive statistics (log returns, daily).

| Variable  | Mean (%) | Std (%) | Min (%) | Max (%) | Skew.  | Kurt.  | JB          | Ann.V |
|-----------|----------|---------|---------|---------|--------|--------|-------------|-------|
| Petrobras | 0.045    | 3.255   | −37.09  | 26.44   | −0.579 | 10.185 | 20,764 ***  | 51    |
| Vale      | 0.039    | 3.031   | −28.28  | 25.03   | −0.244 | 7.235  | 10,411 ***  | 48    |
| Itau      | 0.047    | 2.806   | −24.40  | 31.71   | 0.048  | 8.23   | 13,760 ***  | 44    |
| Bradesco  | 0.030    | 2.926   | −24.56  | 31.67   | −0.126 | 8.90   | 16,095 ***  | 46    |
| Ambev     | 0.025    | 2.333   | −48.96  | 17.02   | −2.254 | 45.490 | 430,741 *** | 37    |
| Gerdau    | 0.021    | 3.254   | −28.40  | 25.44   | −0.266 | 5.576  | 6256 ***    | 51    |
| Ibovespa  | 0.040    | 1.655   | −15.99  | 13.68   | −0.418 | 9.448  | 17,649 ***  | 26    |

Note. N = 5138 daily log-return observations covering the period from 4 January 2005 to 16 April 2026. The ADRs are listed on the NYSE and denominated in U.S. dollars, whereas the Ibovespa is expressed in index points. JB denotes the Jarque–Bera test. \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$ . Annualized volatility (Ann. Vol.) is calculated as the daily standard deviation multiplied by  $\sqrt{252}$ . The Ambev outlier (−48.96% on 9 June 2005) reflects an extraordinary dividend adjustment affecting only Ambev; on the same day, the PBR return was +0.43%, while the Ibovespa return was −0.89%. The Granger-causality results remain robust after excluding the three most extreme observations from March 2020:  $PBR \rightarrow IBVSP$ ,  $F(1) = 8.85$ ,  $p = 0.003$ .

Figure 1 provides a descriptive overview of the Petrobras–Ibovespa system over the full sample. Panel A plots normalized equity prices, with the series rebased to 100 on 3 January 2005, while shaded intervals indicate the main leadership regimes identified in the empirical classification. Panel B reports daily log returns. This figure is useful as a visual summary of the long-run evolution of the two series and of the broad temporal distribution of predictive regimes, but it should be treated as descriptive background rather than as evidence on persistence itself.



**Figure 1.** Normalized equity prices and daily log returns (2005–2026). Note. (Panel A): prices normalized to base 100 (3 January 2005). Shaded: PBR dominates (dark red) and Ibovespa dominates (dark blue). (Panel B): daily log returns.

#### 4. Tests and Empirical Results

Table 2 yields mixed but coherent lag-order evidence: BIC selects VAR(1), HQIC selects VAR(4), and AIC reaches its minimum at VAR(7), although the difference between VAR(6) and VAR(7) is negligible (0.00002). The return series are stationary by both ADF and KPSS criteria. We therefore use the richer AIC-guided specification in the baseline analysis, while confirming that the central directional asymmetry is robust to VAR(1). Residual diagnostics are satisfactory: the Portmanteau test does not reject residual autocorrelation at conventional levels, ARCH effects remain present as expected in daily returns, and the main Petrobras-to-Ibovespa predictive effect survives GARCH-filtered estimation.

**Table 2.** VAR lag order selection criteria.

| Lag | AIC        | BIC        | HQIC       | Selected |
|-----|------------|------------|------------|----------|
| 1   | −15.9781   | −15.9704 * | −15.9754   | BIC      |
| 2   | −15.9793   | −15.9666   | −15.9748   |          |
| 3   | −15.9835   | −15.9656   | −15.9772   |          |
| 4   | −15.9846   | −15.9617   | −15.9766 * | HQIC     |
| 5   | −15.9852   | −15.9571   | −15.9754   |          |
| 6   | −15.9872   | −15.9540   | −15.9756   |          |
| 7   | −15.9872 * | −15.9488   | −15.9737   | AIC *    |
| 8   | −15.9862   | −15.9429   | −15.9710   |          |
| 9   | −15.9852   | −15.9367   | −15.9682   |          |
| 10  | −15.9857   | −15.9321   | −15.9660   |          |

Note. \* AIC(7) and AIC(6) near-identical ( $\Delta = 0.00002$ ). AIC is preferred over BIC for superior out-of-sample forecast accuracy (Burnham & Anderson, 2002). Granger directional asymmetry is fully robust to VAR(1); see Table 3. All specifications are stable.

**Table 3.** Full-sample Granger causality tests.

| Direction      | Lag | F-Stat | p-Value | Sig. |
|----------------|-----|--------|---------|------|
| PBR → Ibovespa | 1   | 9.808  | 0.0019  | ***  |
|                | 2   | 6.697  | 0.0013  | ***  |
|                | 3   | 7.614  | 0.0001  | ***  |
|                | 4   | 5.853  | 0.0001  | ***  |
|                | 5   | 4.678  | 0.0003  | ***  |
| Ibovespa → PBR | 1   | 4.361  | 0.0369  | *    |
|                | 2   | 6.945  | 0.0010  | ***  |
|                | 3   | 6.911  | 0.0003  | ***  |
|                | 4   | 5.178  | 0.0010  | ***  |
|                | 5   | 4.168  | 0.0022  | ***  |

Note. The null hypothesis ( $H_0$ ) is that there is no Granger causality. Results are based on the SSR  $F$ -test using a VAR(6) model with  $N = 5138$  observations. The robustness analysis using the BIC-optimal VAR(1) specification yields similar results: PBR → IBVSP,  $F(1) = 9.81$ ,  $p = 0.002$ ; and IBVSP → PBR,  $F(1) = 4.36$ ,  $p = 0.037$ . The directional asymmetry is preserved ( $AR = 2.25$ ). \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , and \*  $p < 0.10$ .

##### 4.1. Full-Sample Granger Causality Tests

Table 3 reports the full-sample pairwise Granger-causality tests and provides evidence of bidirectional predictive dependence. The null hypothesis of no Granger causality is rejected for both Petrobras → Ibovespa and Ibovespa → Petrobras across lags 1–5, indicating that the two series are dynamically linked in the conditional mean. However, the predictive asymmetry is tilted toward Petrobras. At the one-lag horizon, the evidence of causality from Petrobras to the Ibovespa is stronger ( $F = 9.808$ ,  $p = 0.0019$ ) than that observed in the reverse direction ( $F = 4.361$ ,  $p = 0.0369$ ), yielding an asymmetry ratio of 2.25. At longer lag orders, this directional tilt remains evident overall, although it is not uniform across all specifications. The result is robust to the BIC-optimal VAR(1) specification, under

which the Petrobras-to-Ibovespa direction continues to exhibit stronger statistical evidence of predictive precedence than the reverse direction.

#### 4.2. FEVD Identification Test and Monte Carlo Evidence

Table 4 reports the static GFEVD and shows a remarkably stable connectedness pattern across forecast horizons: the cross-variance share from the Ibovespa to Petrobras remains near 37.6–37.7%, while the total connectedness index is correspondingly flat. Taken in isolation, this near symmetry could be read as evidence of balanced informational interaction. Table 5 shows why that inference is not warranted. In the Monte Carlo design, as contemporaneous correlation increases, row-normalized GFEVD attribution converges mechanically toward symmetry, whereas Cholesky-based attribution diverges sharply with variable ordering. At the observed correlation ( $\rho = 0.7758$ ), the GFEVD cross-series share is 37.64%, while the Cholesky decomposition assigns 60.30% under one ordering. The implication is that Table 4 identifies a pattern of variance connectedness, but not informational symmetry; directional informational asymmetry must instead be inferred from first-moment predictability, as in the Granger causality results.

**Table 4.** Generalized FEVD (Pesaran & Shin, 1998): static connectedness.

| h  | PBR ← PBR (%) | PBR ← Ibovespa (%) |
|----|---------------|--------------------|
| 1  | 62.41         | 37.59              |
| 2  | 62.39         | 37.61              |
| 3  | 62.41         | 37.59              |
| 5  | 62.43         | 37.57              |
| 10 | 62.32         | 37.68              |
| 12 | 62.32         | 37.68              |
| 52 | 62.32         | 37.68              |

Note. Order-invariant GFEVD (Koop et al., 1996; Pesaran & Shin, 1998). TCI bootstrapped 95% CI: [36.44%, 38.74%], SE = 0.60% (B = 500). GFEVD symmetry is mechanically implied under row-normalization with similar forecast error variances (Monte Carlo confirmation in Table 5). Informational asymmetry resides in first-moment Granger dynamics (Table 3, Section 4.1). Cholesky FEVD (h = 10): 59.81%/0.66% vs. 0.49%/60.07% reversed: 59.32 pp ordering artifact.


**Table 5.** Monte Carlo simulation: GFEVD vs. Cholesky FEVD under varying correlation.

| Rho               | GFEVD Cross-Series (%) | GFEVD Std (%) | Cholesky [A, B] (%)      | Ordering Artifact (pp) |
|-------------------|------------------------|---------------|--------------------------|------------------------|
| 0.2               | 4.11                   | 1.08          | 4.30                     | 0.17                   |
| 0.4               | 14.01                  | 1.52          | 16.30                    | 2.24                   |
| 0.6               | 26.61                  | 1.27          | 36.24                    | 9.55                   |
| 0.7758 (observed) | 37.64                  | 0.74          | 60.30 [CI: 58.55, 61.87] | 22.62                  |
| 0.9               | 44.78                  | 0.32          | 81.00                    | 36.26                  |

Note. The Monte Carlo simulation uses 1000 replications for each value of  $\rho$ , with a fixed seed to ensure reproducibility. The data-generating process is a bivariate VAR(1) with coefficient matrix  $A = [[0.05, 0.02], [0.01, 0.03]]$ , unit-variance innovations, and  $T = 5138$  observations. The generalized forecast-error variance decomposition (GFEVD) is computed at  $h = 10$  following Pesaran and Shin (1998) and is row-normalized. The Cholesky FEVD reports the contribution  $PBR \rightarrow IBVSP$  under the ordering [PBR, IBVSP]. The ordering artifact is defined as the difference between the Cholesky FEVD and the GFEVD. This difference increases monotonically with  $\rho$ , supporting H1. The previously used reversal-from-50% measure was non-monotonic because the Cholesky FEVD crosses the 50% threshold at approximately  $\rho = 0.70$ ; it is therefore replaced in Table 5 by the monotonic ordering-artifact measure. The 95% confidence intervals are computed across Monte Carlo replications. As  $\rho$  increases, the GFEVD mechanically converges toward 50%, whereas the Cholesky FEVD diverges monotonically from the GFEVD. The observed system ( $\rho = 0.7758$ ) is consistent with this pattern, yielding an ordering artifact of 22.62 percentage points and a Cholesky FEVD estimate with a 95% confidence interval of [58.55%, 61.87%]. Informational asymmetry is instead identified through first-moment Granger predictive precedence, as discussed in Section 4.1.

#### 4.3. Rolling Causality Tests

The time-varying nature of predictive leadership is illustrated in Figure 2. Panel A reports the rolling bilateral Granger  $F$ -statistics, Panel B shows the corresponding Granger Leadership Index  $LI_F$  by regime, and Panel C presents the discrete regime sequence obtained from the GMM classification. Together, these panels provide a compact visualization of how predictive dominance alternates over time between Petrobras-led, Ibovespa-led, and more balanced configurations. Because the figure directly summarizes the rolling estimation and the regime-classification outputs, it belongs with the dynamic results rather than with the persistence subsection.



**Figure 2.** Rolling Granger  $F$ -statistics, Granger Leadership Index ( $LI_F$ ), and regime sequence. Panel (A): Rolling Granger  $F$ -statistics. Panel (B):  $LI_F$  classified according to the five-regime GMM solution ( $k = 5$ ). Panel (C): Discrete regime sequence. Regime colors represent empirically estimated clusters in the predictive  $F$ -statistic space, interpreted in terms of concentrated versus dispersed information structures. Each rolling window contains  $W = 252$  observations, with a step size of 21 observations and  $T_w = 233$  total windows.

#### 4.4. Regime-Classification Tests

The regime-classification analysis is reported after the rolling causality exercise because it is constructed directly from the bilateral rolling  $F$ -statistics. Table 6 shows that the coarser  $k = 3$  partition delivers tighter geometric separation according to standard internal clustering metrics, but the preferred specification is  $k = 5$ . That choice is driven by economic resolution rather than geometric compactness alone: BIC favors  $k = 5$ , seed stability remains high, and the additional states recover heterogeneity in predictive configurations that is obscured under the coarser partition. Table 6 shows that this refinement is highly

structured rather than arbitrary. The  $k = 3$  Stackelberg-Ibovespa regime maps entirely into the  $k = 5$  Ibovespa-dominant state, while the  $k = 3$  Nash regime maps primarily into Ibovespa-leads and Neutral states. By contrast, the  $k = 3$  Stackelberg-PBR regime is split across PBR-dominant, PBR-leads, and Neutral states, indicating that Petrobras-led episodes are empirically more heterogeneous than the coarser taxonomy suggests. The high cross-classification statistic ( $\chi^2 = 346.97$ ,  $p < 0.001$ ; Cramér's  $V = 0.916$ ) confirms that the  $k = 5$  partition is strongly nested within the  $k = 3$  structure while adding substantively meaningful variation.

**Table 6.** Supplementary validation: clustering metrics, stability, and out-of-sample forecasting. Panel A. GMM clustering validation metrics.

| $k$             | Silhouette | Davies-Bouldin | Calinski-Harabasz | BIC    |
|-----------------|------------|----------------|-------------------|--------|
| 3               | 0.4027     | 0.7937         | 271.16            | 975.10 |
| 5 (BIC-optimal) | 0.2463     | 1.1717         | 202.20            | 921.90 |
| 6               | 0.2281     | 1.2452         | 193.70            | 932.96 |

Note. Clustering is performed in the standardized ( $F_{PBR}$ ,  $F_{IBVSP}$ ) space using  $N = 233$  rolling windows. Higher Silhouette and Calinski-Harabasz values indicate better cluster separation, whereas lower Davies-Bouldin values indicate greater within-cluster compactness. The  $k = 3$  solution exhibits stronger geometric separation according to all three internal validation criteria. Nevertheless,  $k = 5$  is retained as the preferred specification for three reasons. First, the BIC favors the five-regime model, indicating that the coarser  $k = 3$  partition provides an insufficient fit to the distribution of the Granger  $F$ -statistics. Second, seed stability remains high, with a mean agreement rate of 93.2% and 78 of 100 seeds exceeding 95% agreement. Third, the additional regimes recover economically interpretable heterogeneity that is obscured under  $k = 3$ , including the distinction between concentrated and moderate PBR-lead episodes and the identification of a structurally persistent IBVSP-dominant state. The objective is therefore not to maximize geometric compactness alone, but to recover economically meaningful heterogeneity without sacrificing classification stability. GMM classification uncertainty is low: the mean observation-level entropy is 0.273 on the logarithmic scale, relative to a maximum of  $\ln(5) = 1.609$ , yielding a normalized entropy of 0.169, where 0 indicates perfect separation and 1 indicates maximum classification uncertainty. Mean posterior classification probabilities across regimes range from 0.836 to 0.998, confirming high-certainty assignments. These entropy-based diagnostics provide additional support for the  $k = 5$  classification beyond the internal geometric validation criteria.

The cross-classification of  $k = 3$  and  $k = 5$  regimes confirms that the finer partition is structured rather than arbitrary. The  $k = 3$  Stackelberg-Ibovespa regime ( $N = 12$ ) maps entirely to the  $k = 5$  Ibovespa-dominant state. The  $k = 3$  Nash regime ( $N = 102$ ) maps primarily to Ibovespa-leads (80%) and Neutral (19%) states. The  $k = 3$  Stackelberg-PBR regime ( $N = 104$ ) distributes across PBR-leads (34%), Neutral (42%), and PBR-dominant (19%) states—the internal heterogeneity that the finer partition recovers. The cross-classification statistic confirms strong nestedness ( $\chi^2 = 346.97$ ,  $df = 8$ ,  $p < 0.001$ ; Cramér's  $V = 0.916$ ), while the additional  $k = 5$  states add substantive economic content rather than statistical ornament.

#### 4.5. Persistence and Robustness Tests

This section reports robustness checks designed to challenge the main identification result. First, an equal-weighted ex-Petrobras benchmark is constructed as  $r_{ex,t} = \frac{1}{5}(r_{Vale,t} + r_{Itaú,t} + r_{Bradesco,t} + r_{Ambev,t} + r_{Gerdau,t})$ , where each component is the daily log return of the corresponding NYSE-listed ADR, denominated in U.S. dollars and rebalanced daily. The resulting benchmark has a contemporaneous correlation of 0.774 with Petrobras, which is nearly identical to the Petrobras-Ibovespa correlation of 0.776. As additionally reported in the note to Table 7, the lag-1 Granger-causality test indicates predictive precedence from PBR to the ex-PBR equal-weighted benchmark,  $F = 4.590$ ,  $p = 0.032$ , whereas the reverse direction is not statistically significant,  $F = 0.161$ ,  $p = 0.688$ , yielding an asymmetry ratio of 28.5. This result indicates that Petrobras' predictive precedence is not solely attributable to its direct inclusion in the Ibovespa.

Table 7 further evaluates robustness to temporal aggregation. At the weekly frequency, the Petrobras-to-Ibovespa effect remains stronger,  $F = 10.79$ ,  $p = 0.001$ , than the reverse

direction,  $F = 5.43$ ,  $p = 0.020$ , producing an asymmetry ratio of 1.99. Because weekly returns reduce the influence of daily index rebalancing and short-horizon mechanical adjustment, this result weakens the interpretation that the observed predictive asymmetry is driven exclusively by index composition.

**Table 7.** Robustness to data frequency: weekly returns.

| Specification | N Obs. | PBR → IBVSP<br>(Lag 1)          | IBVSP → PBR<br>(Lag 1)        | AR (Lag 1) | Mechanical Causality    |
|---------------|--------|---------------------------------|-------------------------------|------------|-------------------------|
| Daily (main)  | 5138   | $F = 9.81$ *** ( $p = 0.002$ )  | $F = 4.36$ ** ( $p = 0.037$ ) | 2.25       | Present at lag 1        |
| Weekly        | 1028   | $F = 10.79$ *** ( $p = 0.001$ ) | $F = 5.43$ ** ( $p = 0.020$ ) | 1.99       | Absent (no rebalancing) |

Note. Weekly log returns are calculated as the sum of daily log returns within each calendar week. The ex-Petrobras equal-weighted benchmark is defined as  $r_{ex,t} = \frac{1}{5}(r_{Vale,t} + r_{Itaú,t} + r_{Bradesco,t} + r_{Ambev,t} + r_{Gerdau,t})$ . It is rebalanced daily and constructed using U.S.-dollar-denominated ADR returns. The lag-1 Granger-causality test indicates predictive precedence from PBR to the ex-PBR equal-weighted benchmark: PBR → ex-PBR EW,  $F = 4.590$ ,  $p = 0.032$ ; the reverse direction is not statistically significant,  $F = 0.161$ ,  $p = 0.688$ . The resulting asymmetry ratio is  $AR = 28.5$ . Thus, Petrobras' predictive precedence persists in this stylized equal-weighted non-Petrobras benchmark exercise. Although the benchmark mitigates concerns related to Petrobras' direct weight in the Ibovespa, it should not be interpreted as either a value-weighted or sector-neutral reconstruction of the index. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , and \*  $p < 0.10$ .

Table 8 assesses whether the main result is attributable to conditional heteroskedasticity by applying Granger-causality tests to standardized residuals obtained from separate GARCH(1,1) models. Petrobras retains significant lag-1 predictive content for the Ibovespa,  $F = 4.12$ ,  $p = 0.042$ , whereas the reverse direction is not statistically significant at any of the evaluated lags. These findings indicate that the Petrobras-to-Ibovespa predictive-precedence result is not solely driven by univariate ARCH effects, although the specification does not explicitly model dynamic conditional correlations or cross-market volatility spillovers.

**Table 8.** Robustness to ARCH effects: Granger causality on GARCH(1,1)-filtered residuals.

| Lag | PBR → IBVSP F<br>(GARCH Resid) | $p$    | Sig. | IBVSP → PBR F<br>(GARCH Resid) | $p$    | Sig. | AR   |
|-----|--------------------------------|--------|------|--------------------------------|--------|------|------|
| 1   | 4.120                          | 0.0423 | **   | 0.473                          | 0.4915 | ns   | 8.71 |
| 2   | 2.673                          | 0.0693 | *    | 1.151                          | 0.3166 | ns   | 2.32 |
| 3   | 1.805                          | 0.1437 | ns   | 0.912                          | 0.4353 | ns   | 1.98 |

Note. GARCH(1,1) models are estimated separately for each univariate return series using maximum likelihood under a normal error distribution, with returns expressed in percentage terms. Granger causality tests are then applied to the standardized residuals, which have approximately zero mean and unit variance. The null hypothesis is no Granger causality, and inference is based on SSR F-tests. The contemporaneous correlation between the GARCH(1,1) standardized residuals is 0.744 in the full sample, while rolling correlations range from 0.488 to 0.891, with a mean of 0.747 and a standard deviation of 0.084. This pattern indicates time-varying co-volatility that is not fully captured by separate univariate GARCH filters. Petrobras-to-Ibovespa predictive precedence remains significant at lag 1 after GARCH filtering,  $F_{PBR \rightarrow IBVSP} = 4.12$ ,  $p = 0.042$ , whereas the reverse direction is not statistically significant at any lag. These results suggest that the Petrobras-to-Ibovespa predictive-precedence result is not solely driven by univariate ARCH effects. However, because the specification does not model dynamic conditional correlations or cross-asset volatility spillovers, a full DCC-GARCH or BEKK framework would provide a stronger volatility-based robustness test and is left for future research. The asymmetry ratio is defined as  $AR = F_{PBR \rightarrow IBVSP} / F_{IBVSP \rightarrow PBR}$ . \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ , ns = not significant.

Table 9 extends the robustness analysis by incorporating within-dataset proxy controls for commodity-market and domestic macro-financial conditions. At the daily frequency, conditioning on Vale as a commodity-sector proxy reduces, but does not eliminate, the Petrobras-to-Ibovespa effect,  $F = 3.964$ ,  $p = 0.047$ . However, when Vale and Itaú Unibanco are included jointly, the effect becomes statistically insignificant,  $F = 0.622$ ,  $p = 0.430$ , indicating that common commodity and macro-financial channels account for a substantial share of short-horizon conditional-mean predictability. At the weekly frequency, by contrast,

Petrobras' predictive precedence remains significant under all control specifications. The Petrobras-to-Ibovespa statistic ranges from  $F = 8.582$ ,  $p = 0.004$ , to  $F = 9.304$ ,  $p = 0.002$ , while the asymmetry ratio reaches 12.49 when Vale and Itaú are included jointly.

**Table 9.** Multivariate Granger robustness.

| Specification                       | PBR → IBVSP F | <i>p</i> | IBVSP → PBR F | <i>p</i> | AR    |
|-------------------------------------|---------------|----------|---------------|----------|-------|
| PANEL A. DAILY FREQUENCY (lag = 1)  |               |          |               |          |       |
| (i) Bivariate baseline              | 9.808 ***     | 0.002    | 4.361 **      | 0.037    | 2.25  |
| (ii) + Vale (commodity proxy)       | 3.964 **      | 0.047    | 5.704 **      | 0.017    | 0.69  |
| (iii) + Itaú (banking/macro proxy)  | 2.838 *       | 0.092    | 14.660 ***    | 0.000    | 0.19  |
| (iv) + Vale + Itaú (joint)          | 0.622 ns      | 0.430    | 15.190 ***    | 0.000    | 0.04  |
| PANEL B. WEEKLY FREQUENCY (lag = 1) |               |          |               |          |       |
| (i) Bivariate baseline              | 9.382 ***     | 0.002    | 4.719 **      | 0.030    | 1.99  |
| (ii) + Vale (commodity proxy)       | 8.682 ***     | 0.003    | 3.845 *       | 0.050    | 2.26  |
| (iii) + Itaú (banking/macro proxy)  | 9.304 ***     | 0.002    | 0.737 ns      | 0.391    | 12.63 |
| (iv) + Vale + Itaú (joint)          | 8.582 ***     | 0.004    | 0.687 ns      | 0.407    | 12.49 |

Note. Granger-causality  $F$ -tests are estimated at lag 1. Vale's NYSE-listed ADR serves as a commodity-sector proxy, whereas Itaú Unibanco's NYSE-listed ADR serves as a banking and domestic macro-financial proxy. External controls—including Brent crude oil prices, the BRL/USD exchange rate, the VIX, and sovereign CDS spreads—are unavailable in the current dataset and should be incorporated in future research. The asymmetry ratio is defined as  $AR = F_{PBR \rightarrow IBVSP} / F_{IBVSP \rightarrow PBR}$ . \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ , and ns = not significant.

#### 4.6. Regime Persistence

Table 10 shows that regime dynamics are strongly persistent. In all five cases, observed self-transition probabilities exceed stationary frequencies, implying excess persistence beyond unconditional incidence. To address the concern that overlapping windows inflate persistence, block-bootstrap 95% confidence intervals ( $B = 1000$ , block size = 12 windows) confirm that self-persistence substantially exceeds stationary frequency in all regimes. For the three most populated: PBR leads [ $N = 52$ ] self-persist = 0.635, CI = [0.438, 0.714], stationary = 0.223; Neutral [ $N = 81$ ] self-persist = 0.704, CI = [0.549, 0.776], stationary = 0.348; IBVSP leads [ $N = 62$ ] self-persist = 0.590, CI = [0.333, 0.701], stationary = 0.266. PBR dom. [ $N = 23$ ] and IBVSP dom. [ $N = 15$ ] show excess persistence of 0.640 and 0.802 above their stationary frequencies, respectively. The most persistent state is Ibovespa dominates (self-transition probability 0.867, stationary frequency 0.064) (0.867 versus 0.064), while IBVSP leads and PBR dom. are the least persistent relative to their stationary frequencies. This pattern indicates that the estimated leadership regimes are recurrent rather than transitory, and that macro-systemic repricing episodes are harder to exit than firm-specific predictive states.

**Table 10.** Regime persistence: empirical transition probabilities (GMM  $k = 5$ ).

| From/To                  | PBR Dom.     | PBR Leads    | Neutral      | IBVSP Leads  | IBVSP Dom.   |
|--------------------------|--------------|--------------|--------------|--------------|--------------|
| PBR dom. [ $N = 23$ ]    | <b>0.739</b> | 0.043        | 0.043        | 0.174        | 0.000        |
| PBR leads [ $N = 52$ ]   | 0.058        | <b>0.635</b> | 0.212        | 0.096        | 0.000        |
| Neutral [ $N = 81$ ]     | 0.000        | 0.123        | <b>0.704</b> | 0.173        | 0.000        |
| IBVSP leads [ $N = 62$ ] | 0.049        | 0.131        | 0.197        | <b>0.590</b> | 0.033        |
| IBVSP dom. [ $N = 15$ ]  | 0.000        | 0.000        | 0.000        | 0.133        | <b>0.867</b> |

Note. The table reports empirical first-order Markov transition probabilities. Bold diagonal entries indicate regime self-persistence. Under first-order independence, expected self-persistence equals the unconditional regime probability,  $\pi_r$ , where  $\pi_r$  is 0.099 for PBR dominant, 0.223 for PBR leads, 0.348 for Neutral, 0.266 for IBVSP leads, and 0.064 for IBVSP dominant. Actual self-persistence substantially exceeds  $\pi_r$  in every regime, with excess persistence ranging from 0.324 for IBVSP leads to 0.802 for IBVSP dominant. The IBVSP-dominant regime is the most persistent, with a self-transition probability of 0.867, a pattern consistent with the high inertia of macro-systemic repricing episodes.

#### 4.7. Regime Dynamics and Macroeconomic Context

The regime labels used in this section are economically informed descriptions of empirically estimated predictive patterns, not structural estimates of strategic interaction. Information architectures denote the configuration through which predictive signals are concentrated, diffused, or transmitted between the two assets within a given rolling window—a descriptive term, not a formal game-theoretic construct. From this perspective, the five-regime taxonomy points to a persistent alternation between concentrated and dispersed predictive structures. Periods in which one asset exhibits a marked predictive advantage may be interpreted, heuristically, as Stackelberg-like episodes of concentrated signal processing, whereas periods in which neither asset displays sustained forecasting dominance are more consistent with Nash-like configurations of relatively symmetric information diffusion.

This alternation is not random. Table 10 shows that all five regimes exhibit self-persistence well above their stationary frequencies, implying that regime realizations are more persistent than would be expected from unconditional incidence alone. Excess persistence ranges from 0.324 in the IBVSP leads regime to 0.802 in the Ibovespa-dominant regime. The latter is by far the most persistent state, with a self-transition probability of 0.867 relative to a stationary frequency of 0.064, indicating that macro-systemic repricing episodes are substantially harder to exit than other predictive states. More generally, the evidence indicates that the estimated regimes are recurrent informational states rather than short-lived statistical fluctuations. Additional analysis suggests that regime transitions are not purely endogenous to the return system, motivating future extensions with external covariates such as Brent crude, the VIX, and the BRL/USD exchange rate.


The concentrated-information pole—defined by the combined PBR dominates, and PBR leads regimes, which account for 32.2% of rolling windows (75 out of 233)—clusters around two identifiable episodes of firm-specific informational concentration. The first is the pre-salt discovery period (2005–2006), during which Petrobras held unusually concentrated information about subsalt reserves. The second is the Lava Jato investigation-to-recovery episode (2015–2017), when firm-level governance and political information became central to market pricing. Both episodes are consistent with concentrated firm-specific information of the type emphasized in Kyle-type models and with broader geopolitical and firm-level risk transmission.

By contrast, the dispersed-information pole—captured by the IBVSP leads and IBVSP dominates regimes, which together account for 33.1% of rolling windows (77 out of 233)—clusters in macro-systemic repricing episodes where aggregate signals dominate sectoral or firm-specific ones. These episodes include the post-global-financial-crisis recovery (2009), the Bolsonaro electoral cycle (2018–2019), and the COVID-19 shock and recovery period (2020–2021). Such phases are consistent with environments in which broad macro-financial shocks dominate idiosyncratic information and induce system-wide repricing.

The Neutral regime, which represents 34.8% of rolling windows (81 out of 233), can be interpreted as the equilibrium state of the system. In this regime, neither asset holds a sustained predictive advantage, and signals diffuse relatively symmetrically across the firm-index pair. In heuristic terms, this configuration is more consistent with Nash-like informational interaction. This interpretation reinforces the view that predictive leadership in the Petrobras–Ibovespa system is episodic rather than permanent.


Additional insight emerges from the  $k = 6$  sensitivity analysis. As shown in Figure 3, the higher-dimensional classification recovers a distinct bidirectional crisis regime characterized by simultaneous bilateral significance. This regime clusters mainly in 2009 and 2020, suggesting that crisis periods may generate two-way predictive transmission rather than

clear unilateral dominance. That pattern is consistent with periods of stressed liquidity and simultaneous information arrival across both assets.



**Figure 3.** F-statistic space by informational regime (GMM,  $k = 6$  sensitivity). Note. Observations in  $(F\{PBR \rightarrow IBVSP\}, F\{IBVSP \rightarrow PBR\})$  space. Color: LI\_F (red = PBR leads, blue = IBVSP leads).

The persistence structure of this alternation is itself asymmetric. Table 10, together with Figure 4, indicates that macro-systemic repricing regimes are structurally harder to exit than episodes of concentrated firm-specific information. Once the system enters a regime dominated by broad market-wide signals, it tends to remain there longer than when predictive leadership is concentrated in Petrobras-specific information. This asymmetry is economically meaningful, as it suggests that aggregate macro-financial disturbances generate more durable informational configurations than firm-level events.



**Figure 4.** Regime transition matrix and calendar distribution (GMM,  $k = 6$  sensitivity). Note. (Panel A): empirical transition probability matrix. (Panel B): calendar heatmap.

Taken together, these findings indicate that regime dynamics in the Petrobras–Ibovespa system reflect not only endogenous shifts in predictive leadership but also the interaction between firm-specific information, macro-financial shocks, and changing market conditions. The evidence therefore supports an interpretation of price discovery as a regime-dependent process shaped by evolving information architectures rather than by a fixed hierarchy of informational dominance.

#### 4.8. Robustness Synthesis and Out-of-Sample Forecasting Validation

Robustness exercises materially narrow alternative interpretations. Petrobras carries a substantial benchmark weight, so a purely mechanical-composition explanation is a serious concern. However, the predictive asymmetry survives the stylized equal-weight ex-Petrobras benchmark, weekly aggregation, and GARCH filtering. At weekly frequency, where daily rebalancing mechanics are less likely to dominate, the Petrobras-to-Ibovespa effect remains stronger than the reverse direction. Under GARCH-filtered residuals, Petrobras retains lag-1 predictive content, while the reverse direction becomes statistically insignificant. These results do not establish structural causality or definitive price discovery, but they make a purely arithmetic or univariate volatility-driven explanation less convincing. The evidence is therefore more consistent with reduced-form directional forecasting dominance embedded in a commodity-sensitive firm-index system.

To evaluate whether this predictive precedence has exploitable forecasting value, we conduct a recursive one-step-ahead forecasting exercise. The benchmark model is an AR(1) specification for Ibovespa returns, while the alternative model is an ADL(1) specification that augments the AR(1) benchmark with lagged Petrobras returns. The sample is split into a 75% in-sample estimation period and a 25% out-of-sample evaluation period. Forecast accuracy is assessed using RMSE, MAE, the RMSE ratio, out-of-sample  $R^2$ , and the Diebold–Mariano test with the [Harvey et al. \(1997\)](#) small-sample correction.

## 5. Discussion

The results yield four main findings. First, the measurement problem is substantive rather than cosmetic. In a highly correlated firm-index system such as Petrobras–Ibovespa, variance connectedness and informational leadership are not the same object. The Monte Carlo evidence shows that, as contemporaneous correlation rises, row-normalized GFEVD attribution converges mechanically toward symmetry, whereas Cholesky-based attribution becomes increasingly sensitive to ordering. At the observed correlation, the GFEVD cross-series share is 37.64%, while the Cholesky decomposition assigns 60.30% (95% CI = [58.55%, 61.87%] across 1000 replications) under the [PBR, IBVSP] ordering, yielding an ordering artifact of 22.62 pp relative to the GFEVD attribution. The ordering artifact column in Table 5 reports (Cholesky – GFEVD), which is monotonically increasing in  $\rho$  and replaces the previously considered  $|\text{Cholesky} - 50\%|$  metric, which is non-monotone near  $\rho = 0.70$ . This is precisely the type of identification problem foreshadowed by [Kyle \(1985\)](#), [Glosten and Milgrom \(1985\)](#), and [Hasbrouck \(1991, 1995\)](#): price discovery is about the relative speed of information incorporation, not merely about forecast-error variance shares. In that sense, H1 is strongly supported. Cholesky-based informational dominance becomes progressively less reliable as contemporaneous correlation rises, while generalized FEVD symmetry should not be interpreted as evidence of informational symmetry.

Second, directional leadership exists, but not as a fixed hierarchy. Full-sample Granger tests show bilateral predictability, yet the Petrobras-to-Ibovespa effect is stronger on average, especially at short horizons, and remains robust under the parsimonious VAR(1) benchmark. This finding is consistent with the microstructure intuition that the asset closest to firm-specific and commodity-fundamental information should impound that information first

(Kyle, 1985; Glosten & Milgrom, 1985). At the same time, the rolling results show that predictive dominance changes sign across states. Petrobras leads in some episodes, whereas the Ibovespa leads in others. The evidence therefore rejects both constant symmetry and permanent Petrobras dominance. H2 is supported, but only in a state-dependent reduced-form sense: directional predictive leadership emerges intermittently and depends on the dominant source of information, exactly as anticipated in the theoretical framing of the paper.

Third, the leadership process is regime-dependent and economically interpretable. The rolling Granger design, the Granger Leadership Index, and the GMM classification show that predictive dominance alternates across concentrated-information states, neutral states, and macro-dominant states. This result is closely aligned with Back et al. (2000), who show that informational interaction can range from relatively symmetric Nash-like configurations to more concentrated Stackelberg-like configurations, and it also fits the broader literature documenting asymmetric and time-varying oil-equity linkages in emerging and commodity-exposed markets (Basher et al., 2012; Antonakakis et al., 2017; Degiannakis et al., 2018; Shi et al., 2018; Hong et al., 2024). The finer  $k = 5$  partition adds substantive economic content rather than statistical ornament, because it reveals that Petrobras-led episodes are internally heterogeneous, whereas the Ibovespa-dominant state is more sharply defined. H3 is therefore supported: informational leadership is state-dependent rather than constant, and the Petrobras–Ibovespa system alternates between concentrated and more balanced informational configurations.

Fourth, the estimated regimes are persistent rather than transitory. Self-transition probabilities exceed stationary frequencies in all five states, and the Ibovespa-dominant regime is by far the most persistent. This is economically important because it implies that macro-systemic repricing states are harder to exit than firm-specific predictive states. The result is consistent with Lo's (2004) adaptive markets hypothesis, under which informational efficiency and predictive advantage vary through time rather than remaining constant (see also Fama, 1970), and with the idea that aggregate information is slower-moving and more durable than firm-specific information. H4 is therefore strongly supported: the regime structure is recurrent and persistent, not a sequence of short-lived statistical fluctuations.

The historical clustering of the estimated regimes strengthens the economic interpretation of the classification. The concentrated-information pole is associated with the pre-salt discovery period and the Lava Jato episode, when Petrobras-specific fundamentals, governance, and political information plausibly dominated local price formation. By contrast, the dispersed-information pole clusters in the post-GFC recovery, the Bolsonaro electoral cycle, and the COVID-19 shock and recovery, when broad macro-financial repricing dominated sectoral and firm-specific information. The neutral regime, which accounts for roughly one-third of rolling windows, is best interpreted as the benchmark state of the system: neither asset holds a persistent predictive edge, and signals diffuse more symmetrically. This contrast is consistent with the distinction between concentrated private information and market-wide contagion or integration dynamics emphasized by Caldara and Iacoviello (2022) and Bekaert et al. (2005).

From a portfolio management perspective, the evidence suggests that Petrobras should not be treated merely as a high-beta proxy for Brazilian equity risk, but rather as a regime-contingent informational state variable that can inform portfolio formation and risk allocation. In Petrobras-led regimes, firm-specific and commodity-related information appears to be incorporated into Petrobras prices before being fully reflected in the Ibovespa, implying that Petrobras can serve as a tactical signal for short-horizon adjustments in Brazil equity exposure. This signal can be operationalized through the sign, magnitude, and persistence of the rolling Granger Leadership Index, the lagged Petrobras return in the

ADL/VAR specification, and the posterior classification of the market state: a positive and persistent Petrobras-led signal would justify a moderate increase in benchmark exposure, whereas a negative signal would support a reduction in Brazil equity beta or the use of index hedges. Conversely, in Ibovespa-led or macro-dominant regimes, portfolio decisions should rely more heavily on systematic risk indicators—oil prices, exchange rates, sovereign risk, and global volatility—because benchmark-level repricing is driven primarily by aggregate information rather than firm-level signals. However, the out-of-sample forecasting exercise (Table 11) shows that the ADL(1) model does not significantly improve 1-step-ahead Ibovespa forecasts over an AR(1) benchmark (RMSE ratio = 1.002, OOS- $R^2$  = −0.003, DM-HLN = −0.753,  $p$  = 0.451), confirming that predictive precedence is not economically exploitable at a one-step horizon. These implications should therefore be framed as guidance for regime identification, risk overlay design, and tactical allocation rather than as evidence of a mechanically exploitable abnormal-return strategy.

**Table 11.** Out-of-sample ADL(1) forecasting (1-step-ahead, recursive).

| Model               | RMSE    | MAE     | RMSE Ratio | OOS-R2 | DM Stat | $p$ -Value |
|---------------------|---------|---------|------------|--------|---------|------------|
| AR(1) baseline      | 0.01135 | 0.00867 | —          | —      | —       | —          |
| ADL(1) + lagged PBR | 0.01137 | 0.00868 | 1.002      | −0.003 | −0.753  | 0.451      |

Note. Recursive 1-step-ahead forecasting. In-sample  $N$  = 3853 (75%); out-of-sample  $N$  = 1285 (25%). AR(1): Ibovespa on own lag. ADL(1): AR(1) + lagged Petrobras return. Diebold-Mariano test (Harvey et al., 1997), two-sided. OOS- $R^2$  < 0 indicates ADL(1) does not improve on AR(1) MSE benchmark. The absence of exploitable gains is consistent with efficient markets (Grossman & Stiglitz, 1980).

From a policy perspective, the findings imply that the informational role of a dominant commodity-linked firm can extend beyond the issuer itself and shape benchmark-level price discovery. This gives firm-level governance, disclosure quality, and transparency broader market relevance. At the same time, the persistence of macro-dominant states suggests that regulators and market authorities should monitor oil shocks, exchange-rate movements, and sovereign-risk repricing as part of benchmark-level surveillance in commodity-dependent emerging markets. The policy implication is therefore twofold: strengthen issuer-level informational quality and reinforce macro-financial monitoring of cross-market transmission channels. Since the empirical design is reduced-form, these implications should be interpreted as recommendations for surveillance and market design rather than as claims of structural welfare causality.

A final point concerns scope. The framework proposed here is not limited to Brazil. By separating FEVD measurement distortion from the substantive question of directional predictive dominance, and by classifying the latter through persistent regime structures, the paper offers a tractable template for other high-correlation commodity-equity pairs. Replication in systems such as Ecopetrol/COLCAP, YPF/Merval, or Gazprom/MOEX would help determine whether the alternation between concentrated and dispersed informational regimes is a general feature of commodity-dependent equity markets rather than a Brazil-specific outcome. Three limitations should bound interpretation. First, Granger causality identifies predictive priority in the conditional mean, not structurally informed trading in the strict microstructure sense of Kyle (1985); throughout the paper, predictive precedence is therefore the more accurate term. Second, the strategic labels attached to the regimes are heuristic rather than game-theoretically estimated. Third, the within-dataset multivariate controls (Vale, Itau Unibanco) are partial proxies for commodity and macro-financial conditions; Brent crude, BRL/USD, VIX, and CDS spreads are not available in the current dataset and should be incorporated in future work to sharpen identification of common-shock effects. The daily PBR → IBVSP result is attenuated when Vale and Itau are included as controls, while the weekly result survives all within-dataset controls—a

pattern that motivates future extensions with external data. These limitations clarify the findings' scope and motivate the next stage of research.

## 6. Conclusions

This paper shows that price discovery in the Petrobras–Ibovespa system cannot be characterized either by static variance shares or by a fixed directional hierarchy. Methodologically, the evidence demonstrates that in highly correlated firm-index systems, Cholesky-based FEVDs become increasingly unreliable as measures of variance attribution, whereas generalized FEVD symmetry should not be interpreted as informational symmetry. Conditional-mean predictive precedence—the concept used throughout this paper in preference to the stronger claim of structural informational leadership—is better identified through first-moment Granger-type predictability than through second-moment variance attribution alone.

Economically, the results show bilateral predictability with a recurring directional tilt from Petrobras to the Ibovespa, but that tilt is regime dependent rather than constant. The system alternates across persistent informational states associated with firm-specific concentration, neutral diffusion, and macro-systemic repricing. The persistence of these states indicates that price discovery in a commodity-dependent firm-index pair is neither static nor unidirectional; it is a recurrent, state-dependent process shaped by changing information architectures and by the source of shocks hitting the system.

The Brazilian case is especially informative because it concentrates the paper's central identification challenge: Petrobras is both commodity sensitive and a large benchmark component. The fact that the predictive asymmetry survives the ex-Petrobras benchmark, weekly aggregation, and GARCH filtering supports an informational interpretation over a purely arithmetic one. At the same time, the out-of-sample forecasting exercise (Table 11) confirms that this predictive precedence does not produce exploitable one-step-ahead gains, establishing its consistency with semi-strong-form market efficiency and delimiting the scope of the result to regime identification and informational analysis rather than return prediction. More broadly, the paper suggests that in commodity-dependent equity markets, the interaction between a dominant firm and its benchmark should be understood as a dynamic informational system whose leadership structure depends on the prevailing regime and on whether shocks originate in firm-specific fundamentals, commodity conditions, or macro-financial repricing.

The broader contribution is therefore both methodological and economic. The paper offers a portable framework for separating FEVD identification problems from the substantive question of directional predictive dominance in high-correlation commodity-equity systems. Within that framework, the alternation between concentrated and dispersed informational regimes emerges as a dimension of market efficiency that the price-discovery literature has documented only imprecisely. Future work should: (i) extend the analysis using external covariates—Brent crude, VIX, BRL/USD, and sovereign CDS spreads—which are absent from the current dataset but critical for sharper identification; (ii) incorporate multivariate volatility dynamics through DCC-GARCH or BEKK specifications (Bauwens et al., 2006); (iii) implement non-overlapping-window and Markov-switching VAR alternatives (Hamilton, 1989) for regime persistence estimation; and (iv) test the framework in other commodity-dependent benchmark systems such as Ecopetrol/COLCAP, YPF/Merval, and Gazprom/MOEX, with Ecopetrol/COLCAP being especially tractable given the analogous institutional setting.

**Author Contributions:** Conceptualization, A.P.-y.-S.-D., J.M.C.-V. and E.A.-E.; methodology, A.P.-y.-S.-D. and J.M.C.-V.; software, A.P.-y.-S.-D.; validation, J.M.C.-V. and E.A.-E.; formal analysis, A.P.-y.-S.-D., J.M.C.-V. and E.A.-E.; investigation, A.P.-y.-S.-D. and J.M.C.-V.; resources, E.A.-E.; data curation, A.P.-y.-S.-D., J.M.C.-V. and E.A.-E.; writing—original draft preparation, A.P.-y.-S.-D.; writing—review and editing, A.P.-y.-S.-D., J.M.C.-V. and E.A.-E.; visualization, A.P.-y.-S.-D.; supervision, J.M.C.-V.; project administration, E.A.-E. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The price data analyzed in this study are publicly available from Yahoo Finance and were cross-validated against Bloomberg. The derived datasets and the code generated during the study are available from the corresponding author upon reasonable request.

**Acknowledgments:** During the preparation of this manuscript, the authors used Claude (Anthropic; Claude Opus 4.8) for the purposes of translation from Spanish into English, bibliographic search, reference organization, and language editing. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## References

- Antonakakis, N., Chatziantoniou, I., & Filis, G. (2017). Oil shocks and stock markets: Dynamic connectedness under the prism of recent geopolitical and economic unrest. *International Review of Financial Analysis*, 50, 1–26. [\[CrossRef\]](#)
- Back, K., Cao, C. H., & Willard, G. A. (2000). Imperfect competition among informed traders. *The Journal of Finance*, 55(5), 2117–2155. [\[CrossRef\]](#)
- Basher, S. A., Haug, A. A., & Sadorsky, P. (2012). Oil prices, exchange rates and emerging stock markets. *Energy Economics*, 34(1), 227–240. [\[CrossRef\]](#)
- Bauwens, L., Laurent, S., & Rombouts, J. V. K. (2006). Multivariate GARCH models: A survey. *Journal of Applied Econometrics*, 21(1), 79–109. [\[CrossRef\]](#)
- Bekaert, G., Harvey, C. R., & Ng, A. (2005). Market integration and contagion. *The Journal of Business*, 78(1), 39–70. [\[CrossRef\]](#) [\[PubMed\]](#)
- Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference* (2nd ed.). Springer. [\[CrossRef\]](#)
- Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225. [\[CrossRef\]](#)
- Degiannakis, S., Filis, G., & Arora, V. (2018). Oil prices and stock markets: A review of the theory and empirical evidence. *The Energy Journal*, 39(5), 85–130. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134. [\[CrossRef\]](#)
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417. [\[CrossRef\]](#)
- Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100. [\[CrossRef\]](#)
- Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), 393–408.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384. [\[CrossRef\]](#)
- Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2), 281–291. [\[CrossRef\]](#)
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance*, 46(1), 179–207. [\[CrossRef\]](#)
- Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *The Journal of Finance*, 50(4), 1175–1199. [\[CrossRef\]](#)
- Hong, Y., Zhang, R., & Zhang, F. (2024). Time-varying causality impact of economic policy uncertainty. *International Review of Financial Analysis*, 91, 102991. [\[CrossRef\]](#)

- Kalliovirta, L., Meitz, M., & Saikkonen, P. (2016). Gaussian mixture vector autoregression. *Journal of Econometrics*, 192(2), 485–498. [\[CrossRef\]](#)
- Koop, G., Pesaran, M. H., & Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1), 119–147. [\[CrossRef\]](#)
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335. [\[CrossRef\]](#)
- Lo, A. W. (2004). The adaptive markets hypothesis. *The Journal of Portfolio Management*, 30(5), 15–29. [\[CrossRef\]](#)
- Nakhli, M. S., Dhaoui, A., & Chevallier, J. (2022). Bootstrap rolling-window Granger causality dynamics between momentum and sentiment: Implications for investors. *Annals of Finance*, 18(3), 267–283. [\[CrossRef\]](#)
- Pesaran, M. H., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1), 17–29. [\[CrossRef\]](#)
- Shi, S., Phillips, P. C. B., & Hurn, S. (2018). Change detection and the causal impact of the yield curve. *Journal of Time Series Analysis*, 39(6), 966–987. [\[CrossRef\]](#)

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.