ESI Preprints **Not Peer-reviewed**

Asymmetric Return Transmission from Commodity Markets to Equity Sectors: Evidence on Positive and Negative Shock Differentials

Naude Shamma Geste
Shanghai University, China

Approved: 08 May 2026
Posted: 10 May 2026

Cite As:
Geste, N.S. (2026). *Asymmetric Return Transmission from Commodity Markets to Equity Sectors: Evidence on Positive and Negative Shock Differentials.* ESI Preprints.
https://doi.org/10.19044/esipreprint.5.2026.p468

Abstract

This paper investigates whether positive and negative commodity return shocks generate different mean return responses in equity sectors, the asymmetric return transmission hypothesis. Using daily data from January 2010 to March 2025 for four commodity futures (crude oil, gold, copper, natural gas) and five U.S. equity sector ETFs (energy, materials, financials, industrials, technology), we implement the Hatemi-J (2012) positive-negative shock decomposition within a 13-variable Vector Autoregression estimated at seven lags (VAR(7)), and test the equality of positive and negative shock transmission coefficients via Wald tests. This paper models mean return co-movement exclusively; conditional volatility dynamics are not estimated. The principal finding is that gold asymmetry is the only statistically robust result in the data: the null hypothesis of equal positive and negative gold shock transmission is rejected at the 1% significance level for all five equity sectors, with Wald statistics ranging from 12.75 to 20.48. Positive gold shocks generate positive equity returns ($\beta^{+} = +0.086$ to $+0.126$) while negative gold shocks generate negative returns ($\beta^{-} = -0.061$ to -0.096), a pattern consistent with safe-haven rotation dynamics in which gold rallies accompany broad risk-on equity positioning and gold crashes accompany risk-off selling. By contrast, oil asymmetry — while directionally present in the coefficient estimates — is not statistically distinguishable from sampling variation at the 5% level in any sector. Copper asymmetry is statistically significant only for the Technology sector (Wald = 15.80, $p < 0.001$). Sub-period analysis reveals that oil-energy transmission exhibits sign reversals during COVID-19 and the Ukraine war, though estimates based on 231-251 observations are treated as indicative rather than definitive. A dynamic portfolio strategy based on rolling asymmetry estimates achieves a Sharpe ratio of 0.619 and a maximum drawdown of -40.6%, compared to 0.643 and -41.9% for a static commodity-equity blend. The strategy underperforms the static blend on risk-adjusted return but provides modest drawdown protection; sensitivity analysis across signal thresholds and window specifications confirms the drawdown benefit is consistent, while Sharpe ratio improvements require higher thresholds than the baseline specification.

Keywords: Asymmetric return transmission, commodity markets, equity sectors, VAR, Hatemi-J decomposition, Wald test, gold safe haven, portfolio allocation, shock decomposition

1. Introduction

1.1. Background

Commodity markets and equity markets are deeply intertwined in the modern financial system, and understanding how price changes in one domain transmit to the other has been a central preoccupation of financial economics for several decades. From Hamilton's (1983) seminal documentation of the relationship between oil price shocks and U.S. recessions to the growing literature on commodity financialization (Tang and Xiong, 2012), researchers have consistently found that commodity price dynamics carry important information for equity market outcomes. Yet the existing empirical literature has largely treated the transmission of commodity shocks as symmetric, implicitly assuming that a one-percent commodity price increase and a one-percent price decrease generate responses of equal magnitude in equity markets, differing only in sign.

This symmetry assumption is not innocuous. A long tradition of theoretical and empirical work suggests that financial markets respond asymmetrically to gains and losses. Prospect theory (Kahneman and Tversky, 1979) establishes that agents weight losses more heavily than equivalent gains in their utility calculations, which implies that negative price shocks should elicit stronger reactions than positive shocks. In commodity markets specifically, Mork (1989) was among the first to demonstrate that oil price increases and oil price decreases do not generate symmetric macroeconomic responses, finding that oil price hikes had a significantly larger contractionary effect than oil price declines had an expansionary effect. More recently, the literature on gold as a safe haven (Baur and Lucey, 2010) has shown that gold's correlation with equity markets changes sign across market conditions, negative during equity stress, near-zero during calm periods, which is precisely the kind of directional asymmetry that motivates the present paper.

Despite this substantial body of supporting theory, the empirical literature on asymmetric commodity-to-equity transmission remains limited in scope. Most studies examine bilateral relationships between a single commodity (typically oil) and an aggregate equity market index, leaving open the question of whether asymmetric transmission varies across commodities with different economic roles, a financialized safe haven versus a physical production input, and across equity sectors with different degrees of commodity exposure. A systematic, multi-commodity and multi-sector mapping of asymmetric return transmission, grounded in formal statistical testing, is conspicuously absent from the literature.

This paper addresses that gap. We apply the Hatemi-J (2012) positive-negative shock decomposition within a Vector Autoregression framework to estimate separate return transmission coefficients for positive and negative shocks from four commodity futures — crude oil, gold, copper, and natural gas — to five U.S. equity sector ETFs. Critically, we test the equality of these coefficients formally using Wald tests, rather than relying on visual comparisons of coefficient magnitudes as is common in earlier work. This formal testing reveals a finding that is both surprising and important: gold asymmetry is the only statistically robust result in the data, while oil — the commodity most prominently discussed in the asymmetry literature — does not generate statistically distinguishable positive versus negative transmission at the 5% level in any sector over the full sample.

Before proceeding, we emphasize an important scope clarification. This paper models mean return transmission, the difference in expected equity sector returns following positive versus negative commodity shocks. We do not model conditional variance dynamics, GARCH effects, or volatility clustering. "Asymmetric volatility spillovers" in the finance literature refers to a specific empirical phenomenon characterized via BEKK-GARCH or DCC-GARCH models, and that is not what we estimate. Our VAR coefficients capture mean return co-movement, and all claims in this paper are bounded by that definition.

1.2 Research Questions

Three specific research questions guide this study:

1. Which commodity-sector pairs exhibit statistically significant asymmetric mean return transmission, as formally tested by Wald tests for the null hypothesis H_0: $\beta^+ = \beta^-$? Do the patterns of asymmetry align with theoretical predictions based on commodity economic roles — specifically, is financialized gold more likely to exhibit asymmetry than physically traded copper?

2. Is commodity-equity return transmission stable across different macroeconomic regimes, or does sub-period analysis reveal structural shifts in the direction and magnitude of asymmetric transmission during extreme events such as the COVID-19 pandemic and the Russia-Ukraine war?
3. Can rolling estimates of asymmetry dominance — the real-time balance between negative and positive shock transmission across the commodity-sector system — be used to construct a dynamic portfolio allocation strategy that improves upon a static commodity-equity blend, and how robust is any such advantage to transaction costs and parameter specification?

1.3 Contributions

This paper makes three contributions to the existing literature. First, it provides a systematic, Wald-test-grounded analysis of asymmetric mean return transmission across a four-commodity, five-sector matrix, a level of granularity not previously available in the literature. Rather than confirming the asymmetric transmission story that visual coefficient inspection suggests, the formal tests lead to a more nuanced conclusion: asymmetry is commodity-specific, concentrated in gold, and reflects the financialized safe-haven channel rather than the physical cost-of-production channel that the oil-asymmetry literature has emphasized.

Second, this paper provides the first sub-period analysis of asymmetric oil-energy transmission across the COVID-19 and Ukraine war regimes, documenting directional sign reversals in these windows that are consistent with policy-driven decoupling and supply-side narrative effects. While these estimates are based on limited sub-samples and cannot be tested at conventional significance levels, they suggest that any structural model of commodity-equity asymmetry must allow for regime-dependent parameters.

Third, the paper develops an exploratory portfolio application that links the asymmetry estimation framework to practical allocation decisions. The results are presented transparently, including the finding that the dynamic strategy underperforms the static blend on Sharpe ratio at the baseline specification, which is itself informative about the limits of rolling asymmetry signal quality in the presence of high-dimensional VAR estimation noise.

1.4 Paper Structure

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on commodity-equity linkages, asymmetric financial market responses, and the Hatemi-J methodology. Section 3 describes the data, stationarity tests, and the asymmetric VAR specification in detail. Section 4 presents the main empirical results, including Wald test findings, coefficient analysis, and Forecast Error Variance Decomposition. Section 5 provides robustness tests via the asymmetry ratio and sub-period analysis. Section 6 develops and evaluates the dynamic portfolio strategy. Section 7 concludes.

2. Literature Review

2.1 Commodity-Equity Return Linkages: From Fundamentals to Financialization

The academic literature on commodity-equity return co-movement has evolved through several distinct phases, each reflecting shifts in the economic mechanisms connecting these two asset classes. The foundational phase, associated primarily with Hamilton (1983) and Jones and Kaul (1996), established that oil price shocks influence aggregate equity returns through real economic channels: as oil prices rise, production costs increase for energy-intensive firms, aggregate consumer purchasing power declines, and corporate earnings fall — all of which depress equity valuations. This cost-of-production and aggregate demand channel remains the dominant narrative for understanding oil-energy sector linkages, where the economic logic is most direct.

However, the expansion of commodity markets as a financial asset class from the mid-2000s onward introduced a second, more financially driven channel that has substantially altered the empirical landscape. Tang and Xiong (2012) documented that the entry of institutional investors into commodity index products created structural co-movement between commodity and equity returns that could not be explained by fundamental supply-demand relationships. Their key insight was that when index investors rebalance their portfolios — buying equities while selling commodities, or vice versa — they generate mechanical price pressures that transmit across asset classes, regardless of whether the underlying economic fundamentals would predict such co-movement. Basak and Pavlova (2016) subsequently formalized this mechanism in a general equilibrium model, showing that the presence of institutional commodity index investors causes commodity prices to become more correlated with each other and with equity markets. This financialization channel is especially relevant for gold and crude oil, which are the most heavily traded and institutionally held commodity futures.

Building on these foundations, a third phase of research has focused on measuring the direction and magnitude of commodity-equity transmission using VAR-based spillover tools. Diebold and Yilmaz (2012) developed a variance decomposition-based spillover index that has become the standard framework for quantifying directional return and volatility transmission across asset classes. Applied to commodity-equity systems, this approach consistently reveals sector-level heterogeneity: energy and materials sectors show the strongest co-movement with commodity markets due to their direct operational exposure, while financial, industrial, and technology sectors show weaker but statistically meaningful linkages through indirect channels including investor risk appetite, credit conditions, and portfolio rebalancing effects (Sadorsky, 2014; Degiannakis, Filis, and Arora, 2018). The present paper extends this tradition to the asymmetry dimension, asking not merely how strong these linkages are, but whether they differ in magnitude depending on the direction of the commodity shock.

2.2 Asymmetric Responses in Financial Markets: Theory and Evidence

The hypothesis of asymmetric market responses to positive and negative shocks has deep roots in both behavioral and structural economics. At the microeconomic level, Kahneman and Tversky's (1979) prospect theory establishes that individuals subjectively weight losses approximately twice as heavily as equivalent gains, generating an asymmetric valuation of positive and negative price movements. This psychological mechanism implies that negative commodity price shocks — which threaten portfolio values or signal deteriorating economic conditions — should elicit stronger reactions from equity investors than positive shocks of equal magnitude.

At the macroeconomic level, Mork (1989) provided the first systematic empirical evidence of asymmetric commodity effects, demonstrating that oil price increases had significantly larger negative effects on U.S. GDP growth than oil price decreases had positive effects. Hamilton (1996) and Bernanke, Gertler, and Watson (1997) attributed this asymmetry to several structural mechanisms: the irreversibility of investment decisions (firms will not immediately restart capital expenditures when oil prices fall after a prolonged high-price period), sectoral reallocation costs (workers displaced from energy-intensive sectors cannot instantaneously retrain and relocate), and the uncertainty channel (commodity price volatility itself reduces investment, and price crashes are often accompanied by high volatility). Collectively, these mechanisms provide a compelling theoretical foundation for expecting asymmetric equity market responses to commodity shocks.

In financial markets specifically, the leverage effect literature — documented by Black (1976) and Christie (1982) and formalized in the asymmetric GARCH models of Nelson (1991) and Glosten, Jagannathan, and Runkle (1993) — establishes that negative equity price shocks generate disproportionately larger increases in conditional return volatility than positive shocks of equal magnitude. This is an asymmetry in second moments (variance), and it is important to note that it is conceptually distinct from the mean return asymmetry that this paper examines. Engle and Ng (1993) developed the news impact curve as a diagnostic tool for detecting such variance asymmetries. While these volatility asymmetry results are well-documented, the analogous question of whether commodity-equity mean return transmission is asymmetric — with negative commodity shocks generating larger equity responses in expectation — has received comparatively less formal attention.

For commodity-equity systems specifically, Reboredo (2014) provides the most direct antecedent to the present paper, finding that oil price crashes transmit to European renewable energy equity markets with greater intensity than oil price rallies, particularly during the 2008-2009 global financial crisis. Ji and Fan (2012) extend the asymmetry analysis to non-energy commodities, finding commodity-specific patterns. However, these studies focus on bilateral relationships and do not provide the multi-commodity, multi-sector Wald test framework that would allow systematic comparison of asymmetric transmission across different commodity types and sectors with different economic exposures.

An especially important strand of the asymmetry literature concerns gold's behavior as a safe haven. Baur and Lucey (2010) established that gold functions as a safe haven — an asset that is either uncorrelated with or negatively correlated with equities during periods of market stress — but that this property is state-dependent and asymmetric. Gold rallies tend to co-occur with broad commodity market strength and risk-on equity positioning, while gold crashes coincide with risk-off episodes in equity markets. This is precisely the kind of sign-dependent transmission that the Hatemi-J decomposition should detect, and it is the key economic intuition behind our finding that gold asymmetry is the most statistically robust result in the data.

2.3 The Hatemi-J (2012) Asymmetric Decomposition

Our methodological approach builds directly on Hatemi-J (2012), who proposed decomposing an asset return series into its positive and negative partial sums to test for asymmetric causal relationships. For a return series r_t, the positive component $r+$ captures all upward price movements ($r+ = max(r_t, 0)$) while the negative component $r-$ captures all downward movements ($r- = min(r_t, 0)$). These components are then entered as separate variables in the empirical system, allowing distinct transmission coefficients to be estimated for each direction of shock.

This approach has important advantages over alternative asymmetry methods. Unlike threshold VAR models (Balke, 2000), it does not require prior identification of regime thresholds or a minimum sample within each regime. Unlike Markov-switching VAR (Hamilton, 1989), it does not require specifying the number of regimes or assuming that regime membership is unobservable. And unlike asymmetric GARCH models, which characterize conditional volatility rather than mean spillovers, the Hatemi-J decomposition directly estimates the return transmission coefficients for positive and negative shocks, which is the quantity of most direct interest for portfolio management and risk assessment.

Hatemi-J (2012) applied this decomposition in the context of asymmetric cointegration testing, finding evidence of asymmetric long-run relationships in international financial markets. Subsequent applications have extended the approach to causality testing (Hatemi-J, 2008) and bilateral exchange rate relationships. The extension to a multi-asset VAR spillover system, simultaneously covering multiple commodity types and multiple equity sectors, is the methodological contribution of the present paper.

2.4 Dynamic Portfolio Allocation with Commodity Signals

The final strand of literature relevant to this paper concerns the use of real-time signals to dynamically adjust commodity-equity portfolio allocations. The case for including commodities as a portfolio diversifier rests on the well-documented findings of Gorton and Rouwenhorst (2006), who showed that commodity futures historically exhibited low correlation with equities and positive correlation with inflation, making them valuable portfolio additions from a risk-return perspective. Erb and Harvey (2006) extended this analysis to show that momentum-based commodity timing strategies can improve risk-adjusted performance. However, subsequent research has documented that the diversification benefits of static commodity allocations are unreliable — they tend to disappear precisely when they are most needed, during acute financial crises, due to increased cross-asset correlation under stress (Mensi et al., 2013).

This unreliability of static commodity allocations has motivated interest in dynamic, signal-based approaches. Moreira and Muir (2017) demonstrate that volatility-managed equity portfolios — which scale down equity exposure when volatility is high and scale up when volatility is low — generate higher Sharpe ratios than static buy-and-hold strategies. Their work establishes the general principle that real-time risk signals contain exploitable information for portfolio timing, even when the signal is relatively simple. Regime-switching allocation models (Ang and Bekaert, 2002; Guidolin and Timmermann, 2007) extend this logic to multi-asset systems with discrete latent states. The portfolio application in this paper draws on this tradition but uses a novel signal: the real-time balance between negative and positive commodity shock transmission across the full commodity-sector system, estimated from rolling VAR regressions. The signal is orthogonal to standard volatility or return-based signals, which makes it potentially complementary to existing timing approaches, though the in-sample backtesting evidence presented here is not sufficient to establish its out-of-sample value.

3. Methodology

3.1 Data Construction and Sample Selection

Our empirical analysis is based on daily closing price data for nine assets — four commodity futures contracts and five U.S. equity sector exchange-traded funds — downloaded from Yahoo Finance using the yfinance Python library. The sample spans January 1, 2010 to March 31, 2025, capturing approximately 15 years of trading activity, multiple complete business cycles, two commodity super-cycles, and a range of macroeconomic shocks including the European sovereign debt crisis, the oil price collapse of 2014-2016, the COVID-19 pandemic, and the geopolitical shocks of 2022. After computing daily log returns, the final estimation sample consists of 3,822 observations.

The selection of commodities and equity sectors follows an explicit theoretical framework designed to span different mechanisms of commodity-equity transmission. On the commodity side, crude oil (WTI front-month futures, CL=F) represents the dominant energy commodity with the most direct cost-of-production linkages to equity markets. Gold (COMEX front-month, GC=F) represents the financialized safe-haven segment of the commodity complex, where investor behavioral dynamics, rather than industrial supply and demand, are the primary price drivers. Copper (COMEX front-month, HG=F) represents the industrial metals segment with strong real-economy sensitivity and historically high correlation with global manufacturing activity. Natural gas (Henry Hub front-month, NG=F) provides a further control, with more regionalized, weather-driven pricing that tests whether asymmetric transmission generalizes beyond the globally integrated oil and gold markets.

On the equity side, the five SPDR sector ETFs: Energy (XLE), Materials (XLB), Financials (XLF), Industrials (XLI), and Technology (XLK), span a continuum from direct to indirect commodity exposure. The Energy sector has the most direct operational link to commodity prices through exploration, production, and refining activities. Materials occupies an intermediate position, with exposure to both energy costs as an input and metals prices through mining and chemical operations. Financials, Industrials, and Technology have increasingly indirect commodity exposure, primarily through macroeconomic channels, credit conditions, and investor risk appetite. This exposure hierarchy allows us to test whether asymmetric transmission is stronger for sectors with more direct commodity linkages, as one might expect if physical production cost channels dominate, or whether it is instead concentrated in indirectly exposed sectors where behavioral and financial channels operate, as our gold findings ultimately suggest.

Table 1: Asset selection and economic rationale. Log returns are computed as r_t = ln(P_t / P_{t-1}). All prices sourced from Yahoo Finance via the yfinance Python library.

Category	Asset	Ticker	Description & Rationale
Commodity	Crude Oil	CL=F	WTI front-month futures. Primary global energy benchmark; most widely traded commodity.
Commodity	Gold	GC=F	COMEX front-month futures. Safe-haven asset; key test case for financial vs. physical asymmetry.
Commodity	Copper	HG=F	COMEX front-month. Barometer of global industrial activity; strong real-economy linkages.
Commodity	Natural Gas	NG=F	Henry Hub front-month. Regionalized pricing; tests whether asymmetry generalizes beyond oil.
Equity Sector	Energy	XLE	SPDR Energy Select Sector ETF. Direct commodity exposure through E&P and integrated majors.
Equity Sector	Materials	XLB	SPDR Materials Select Sector ETF. Mining, chemicals; intermediate commodity exposure.
Equity Sector	Financials	XLF	SPDR Financial Select Sector ETF. Indirect commodity exposure through credit and risk channels.
Equity Sector	Industrials	XLI	SPDR Industrial Select Sector ETF. Commodity input sensitivity via manufacturing.
Equity Sector	Technology	XLK	SPDR Technology Select Sector ETF. Lowest direct commodity exposure; pure financial channel test.

3.2 Descriptive Statistics

Table 2 reports descriptive statistics for the daily log return series. Several patterns are notable and inform the methodological choices that follow. First, all series exhibit statistically significant non-normality as confirmed by the Jarque-Bera test ($p < 0.001$ for all), with excess kurtosis (fat tails) present throughout. This non-normality is particularly pronounced for crude oil (kurtosis = 25.12), which is largely attributable to the extraordinary WTI futures negative-price event of April 2020, when front-month oil futures briefly traded below zero for the first time in history. The presence of these extreme observations motivates the use of the Hatemi-J decomposition, which is based on a non-parametric sign-separation of returns and does not require Gaussian innovations.

Table 2: Descriptive statistics of daily log returns (2010-2025). JB = Jarque-Bera normality test. All series exhibit significant excess kurtosis and non-normality. Oil kurtosis (25.12) is driven by the April 2020 WTI negative-price event.

Series	Mean	Std Dev	Min	Max	Skewness	Kurtosis	JB p-val
Oil	0.0001	0.0256	-0.2822	0.3196	0.070	25.120	<0.001
Gold	0.0003	0.0100	-0.0982	0.0578	-0.585	5.836	<0.001
Copper	0.0001	0.0143	-0.0755	0.0720	-0.099	1.876	<0.001
NatGas	-0.0001	0.0358	-0.3005	0.3817	0.234	7.453	<0.001
Energy	0.0003	0.0173	-0.2249	0.1487	-0.802	15.533	<0.001
Materials	0.0004	0.0131	-0.1166	0.1112	-0.452	7.176	<0.001

Financials	0.0005	0.0139	-0.1474	0.1236	-0.509	11.972	<0.001
Industrials	0.0005	0.0122	-0.1204	0.1191	-0.546	10.707	<0.001
Technology	0.0007	0.0135	-0.1487	0.1109	-0.453	9.145	<0.001

Second, Technology records the highest mean daily return (+0.0007, equivalent to approximately 17% annualized), reflecting the sustained outperformance of the technology sector over the 2010-2025 sample period driven by the digital economy transformation and artificial intelligence investment cycle. Natural gas records the only negative mean return (-0.0001), consistent with structural oversupply from the U.S. shale revolution depressing prices over the full sample. Third, the standard deviations reveal clear risk ordering: Natural Gas is the most volatile commodity (3.58% daily) due to seasonal demand swings and storage dynamics, followed by Oil (2.56%), while Copper (1.43%) and Gold (1.00%) are substantially less volatile. Among equity sectors, Energy (1.73%) is the most volatile, consistent with its direct commodity exposure, while Industrials (1.22%) is the most stable.

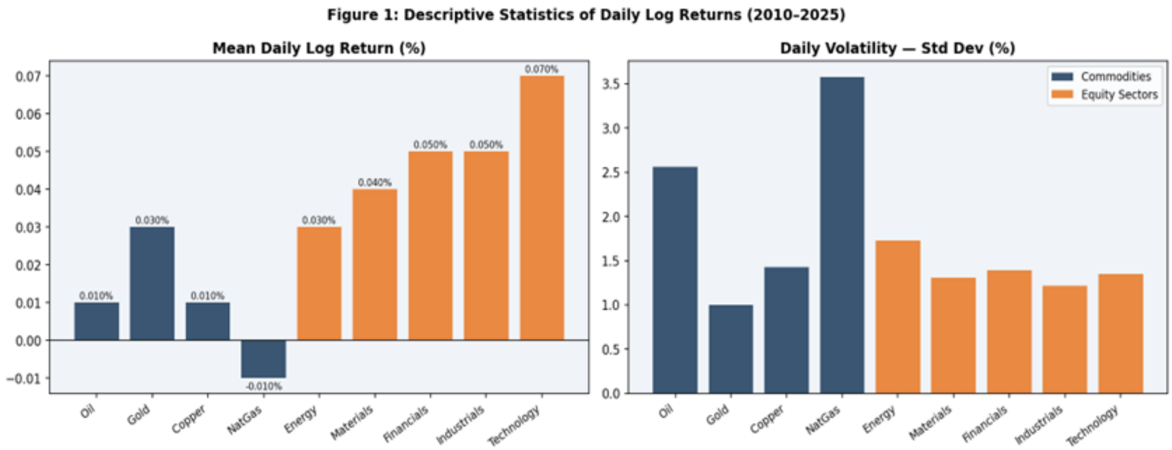

Figure 1: Mean daily returns (left) and daily standard deviations (right) by asset. Navy bars = commodities; orange bars = equity sectors.

3.3 Stationarity Testing

Before estimating the VAR, we verify that all return series are stationary using two complementary tests. The Augmented Dickey-Fuller (ADF) test, with the null hypothesis of a unit root, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test, with the null hypothesis of stationarity, are applied to each series. Using both tests in tandem guards against the well-known low power of the ADF in small samples and the size distortions of the KPSS in the presence of long-run mean reversion. Lag lengths for the ADF are selected by the Akaike Information Criterion, while the KPSS uses the Newey-West automatic bandwidth selection.

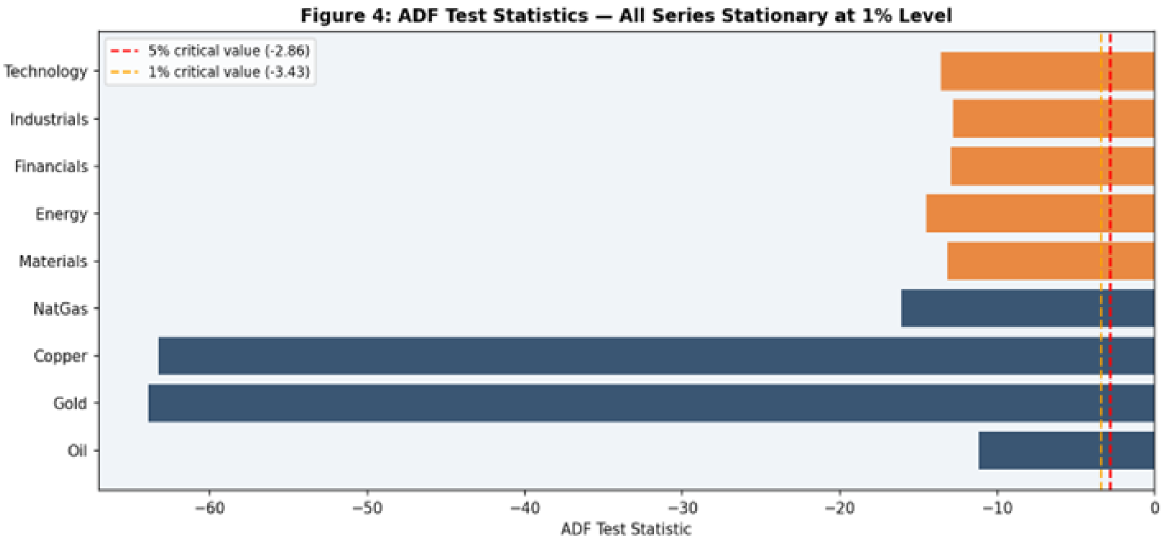

Figure 2: ADF test statistics for all log return series. Dashed lines indicate 5% (-2.86) and 1% (-3.43) critical values. All series are stationary at the 1% level.

Table 3: ADF and KPSS stationarity test results. Both tests unanimously confirm I(0) stationarity at the 1% significance level for all nine return series. ADF lag lengths are selected by AIC; KPSS uses automatic bandwidth selection.

Series	ADF Stat	ADF p-val	ADF Lags	KPSS Stat	KPSS p-val	Stationary
Oil	-11.186	<0.001	24	0.089	>0.10	**YES**
Gold	-63.879	<0.001	0	0.239	>0.10	**YES**
Copper	-63.281	<0.001	0	0.147	>0.10	**YES**
NatGas	-16.105	<0.001	13	0.067	>0.10	**YES**
Energy	-14.496	<0.001	17	0.071	>0.10	**YES**
Materials	-13.202	<0.001	30	0.021	>0.10	**YES**
Financials	-12.941	<0.001	26	0.037	>0.10	**YES**
Industrials	-12.812	<0.001	30	0.019	>0.10	**YES**
Technology	-13.617	<0.001	25	0.065	>0.10	**YES**

The results in Table 3 confirm stationarity of all log return series at the 1% level across both tests. ADF statistics are strongly negative for all series, ranging from -11.186 (Oil, reflecting the higher serial dependence in energy futures due to rolling contract effects) to -63.879 (Gold, reflecting near-random-walk daily returns with no detectable autocorrelation structure at lag zero). KPSS statistics are uniformly well below the 5% critical value of 0.463, with the highest value being 0.239 for Gold. These unanimous results validate the direct estimation of the VAR system on log returns without further differencing, and confirm that any long-run co-movement in levels is not contaminating the short-run transmission dynamics we estimate.

3.4 Asymmetric VAR Framework

3.4.1. Shock Decomposition: Mechanics and Validation

The first step of our methodology follows Hatemi-J (2012) in decomposing each commodity return series into its positive and negative daily components. Formally, for each commodity i at time t, we define:

$$r^{+}_{i,t} = \max(r_{i,t}, 0) \quad \text{and} \quad r^{-}_{i,t} = \min(r_{i,t}, 0)$$

such that $r_{i,t} = r^{+}_{i,t} + r^{-}_{i,t}$ holds exactly for every observation. The positive component $r^{+}_{i,t}$ is equal to the observed return on days when the commodity price rose, and zero otherwise; symmetrically, the negative component $r^{-}_{i,t}$ is equal to the observed return on days when the commodity price fell, and zero otherwise. By construction, both components are stationary (confirmed by ADF tests yielding statistics below -30 for all components), and both components have the same expected absolute magnitude as half the total return distribution.

An important validation step is to confirm that the unconditional means of the positive and negative components are nearly symmetric in magnitude. If the positive and negative components had systematically different means — for example, because oil has a positive long-run trend so positive daily returns are on average larger than negative ones — then any estimated asymmetry in the VAR coefficients could reflect mechanical differences in shock distributions rather than genuine asymmetry in transmission. In our data, the unconditional means are near-symmetric for all commodities: Oil_pos mean = +0.00846 versus Oil_neg mean = -0.00834; Gold_pos = +0.00370 versus Gold_neg = -0.00342; Copper_pos = +0.00537 versus Copper_neg = -0.00526; NatGas_pos = +0.01269 versus NatGas_neg = -0.01281. This near-symmetry validates the decomposition and confirms that any coefficient-level asymmetry we find reflects genuine differential transmission dynamics.

3.4.2. System Specification and Lag Selection

The eight commodity shock components (four commodities times two sign-components) are combined with the five equity sector returns to form a 13-variable asymmetric VAR(p) system:

$$Y_t = c + A_1 Y_{t-1} + A_2 Y_{t-2} + \ldots + A_p Y_{t-p} + u_t$$

where $Y_t = [r^{+}_{Oil}, r^{+}_{Gold}, r^{+}_{Cu}, r^{+}_{NG}, r^{-}_{Oil}, r^{-}_{Gold}, r^{-}_{Cu}, r^{-}_{NG}, r_{En}, r_{Ma}, r_{Fi}, r_{In}, r_{Te}]'$ is the (13×1),

vector of shock components and sector returns; A_k are (13×13) coefficient matrices capturing the dynamic transmission structure at lag k; c is a vector of constants; and $u_t \sim (0, \Sigma)$ is a vector of white noise innovations with potentially non-diagonal covariance matrix Σ.

Lag order selection follows the Akaike Information Criterion (AIC). For the baseline 9-variable linear VAR (which uses raw commodity returns rather than decomposed components), AIC selects $p = 1$, indicating that the dominant commodity-equity transmission occurs at the one-day lag. For the asymmetric 13-variable system, AIC selects $p = 7$. This higher lag order reflects an important feature of the decomposed shock components: because $r^+_{i,t}$ and $r^-_{i,t}$ are each zero on approximately half of all trading days (the days when the commodity moved in the opposite direction), they exhibit a form of zero-inflation that generates higher-order serial dependence patterns not present in the original return series. Residual diagnostics for the baseline VAR(1) confirm adequate specification: Durbin-Watson statistics for all nine equations range from 1.985 to 2.004, well within the [1.8, 2.2] range that indicates absence of first-order residual autocorrelation.

3.4.3 The Asymmetry Test: Wald Tests for $\beta^+ = \beta^-$

The central inferential contribution of this paper is the application of formal Wald tests to assess whether positive and negative shock transmission coefficients are statistically different from each other. For each commodity i and equity sector j, the null hypothesis of symmetric transmission is:

$$H_0: \beta^+_{ij} = \beta^-_{ij} \quad \text{against} \quad H_1: \beta^+_{ij} \neq \beta^-_{ij}$$

where β^+_{ij} and β^-_{ij} denote the Lag-1 coefficients on $r^+_{i,t-1}$ and $r^-_{i,t-1}$ in the sector j equation of the asymmetric VAR. The Wald test statistic for this hypothesis takes the form:

$$W_{ij} = (\beta^-_{ij} - \beta^+_{ij})^2 / \mathrm{Var}(\beta^-_{ij} - \beta^+_{ij}) \sim \chi^2(1) \text{ under } H_0$$

Standard errors are estimated from the VAR residual covariance matrix and the sample standard deviations of the shock components, following standard VAR inference theory. The test statistic follows a chi-squared distribution with one degree of freedom under the null, with the 5% critical value of 3.84 and the 1% critical value of 6.63. We test all 20 commodity-sector pairs (4 commodities × 5 sectors) and report both the Wald statistics and the corresponding p-values.

In addition to the Wald tests — which are the primary inference tool — we report two complementary descriptive measures. The asymmetry ratio $|\rho^-_{ij}| / |\rho^+_{ij}|$, where ρ denotes sample correlations between shock components and sector returns, provides a model-free, distribution-free descriptive measure of asymmetry that does not depend on the VAR specification. The Forecast Error Variance Decomposition (FEVD) at the 20-day horizon provides a medium-term perspective on the economic magnitude of shock contributions, decomposing equity sector return variance into contributions from each positive and negative shock component. However, it is critical to note that these descriptive measures are not substitutes for formal statistical testing — as we demonstrate in Section 5, they can suggest asymmetry in cases where Wald tests cannot reject the null of symmetry.

4. Empirical Results

4.1 Correlation Structure and Co-movement

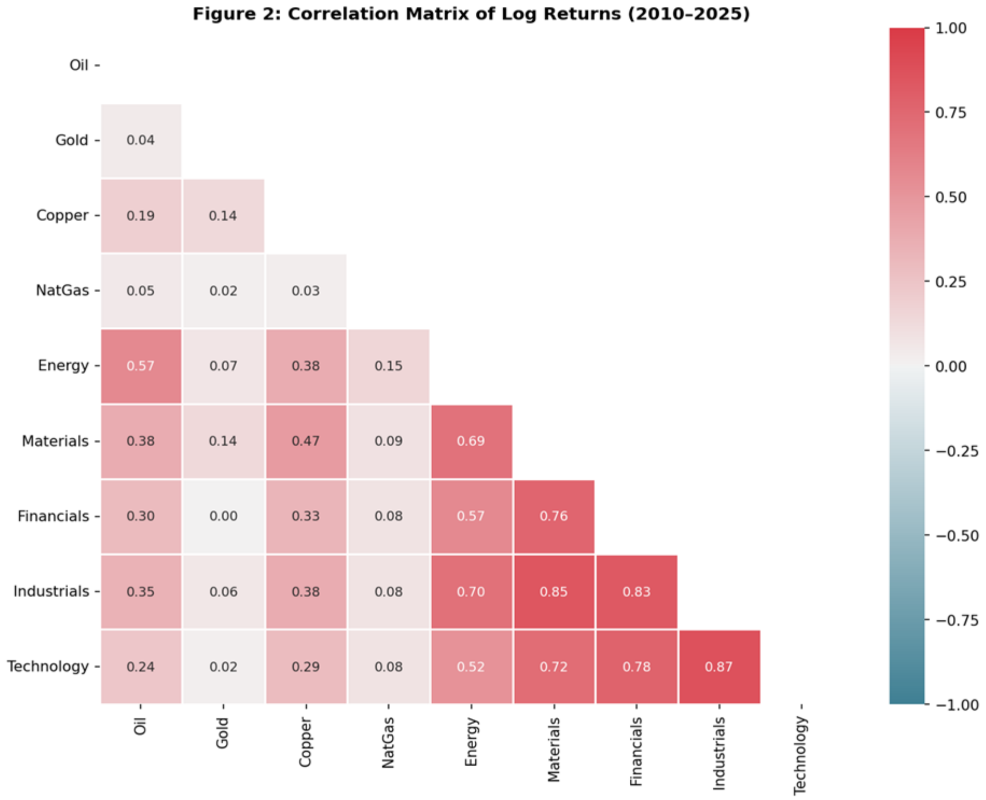

Figure 3: Pairwise correlation matrix of daily log returns (2010–2025). Upper triangle omitted for readability. The two-block structure is clear: strong within-equity-sector co-movement (0.52–0.87) reflects common market factor exposure, while commodity-equity cross-correlations (0.00–0.57) are uniformly weaker.

Before examining asymmetric transmission, it is useful to establish the baseline symmetric correlation structure as a reference point. Figure 1 reveals two clearly distinct blocks. Within the equity sector block, pairwise correlations are uniformly high, ranging from 0.52 between Energy and Technology to 0.87 between Industrials and Technology. This strong co-movement reflects the dominant influence of the common equity market factor, changes in broad risk appetite, macroeconomic growth expectations, and discount rates, on all sector returns simultaneously. The within-equity correlations are substantially higher than the commodity-equity cross-correlations, confirming that commodity price dynamics represent an additional, idiosyncratic source of information for equity sectors beyond what the common factor already explains.

Across the commodity-equity boundary, Oil shows the strongest cross-sector correlations: 0.57 with Energy (reflecting the direct operational link through which oil price changes affect the profitability of exploration and production companies), 0.38 with Materials, and 0.30-0.35 with Financials and Industrials. These declining correlations as we move from direct to indirect sector exposure are consistent with the theoretical exposure hierarchy described in Section 3.1. Gold exhibits a strikingly different pattern: its correlations with all equity sectors are near-zero (ranging from 0.00 for Financials to 0.14 for Materials), confirming its role as a diversifier whose price movements are driven by investor sentiment, inflation expectations, and global risk appetite rather than real economic activity. Natural gas shows uniformly weak equity correlations (0.08–0.15 across all sectors), reflecting the predominantly weather-driven and regionally confined nature of Henry Hub natural gas pricing.

These baseline correlations set up the central empirical question: do positive and negative commodity shocks generate different co-movement patterns, and if so, in which commodity-sector pairs and to what degree? The unconditional correlations are symmetric by construction, they do not distinguish between the direction of commodity price movements. The asymmetric VAR, applied in the following section, allows us to break this symmetry and estimate distinct positive and negative shock responses.

4.2 Wald Test Results: The Central Finding

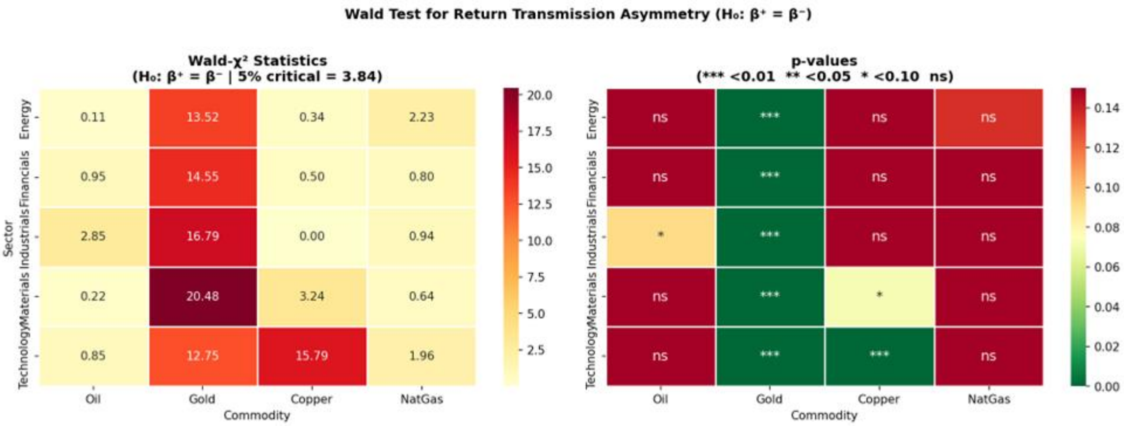

Figure 4: Wald test results for H₀: $\beta^+ = \beta^-$ across all 20 commodity-sector pairs. Left panel: Wald-χ^2 statistics (dashed line = 5% critical value of 3.84). Right panel: p-values with significance stars. Gold dominates in statistical significance; oil is uniformly non-significant.

Table 4: Asymmetric VAR Lag-1 coefficients and Wald statistics. *** p < 0.001. Oil Wald statistics: Energy 0.11 (ns, p=0.737), Materials 0.22 (ns, p=0.638), Financials 0.95 (ns, p=0.329), Industrials 2.85 (marginal, p=0.091), Technology 0.85 (ns, p=0.357). Copper Wald statistics: Technology 15.80*** (p<0.001), Materials 3.24* (p=0.072), all others ns. NatGas: all Wald statistics below 2.3, all p-values above 0.13.

Sector	β^{+}(Oil)	β^{-}(Oil)	β^{+}(Gold)	β^{-}(Gold)	β^{+}(Cu)	β^{-}(Cu)	W(Gold)	p(Gold)
Energy	-0.014	-0.021	0.112	-0.096	0.038	0.015	13.52***	<0.001
Materials	-0.013	-0.005	0.126	-0.069	-0.018	0.035	20.48***	<0.001
Financials	-0.012	+0.005	0.086	-0.096	-0.002	0.022	14.55***	<0.001
Industrials	-0.014	+0.012	0.107	-0.061	0.006	0.008	16.79***	<0.001
Technology	-0.002	+0.014	0.092	-0.075	-0.042	0.086	12.75***	<0.001

The Wald test results, presented in Figure 4 and Table 4, constitute the central empirical contribution of this paper. They reveal a finding that is both striking and methodologically important: the commodity that generates statistically robust asymmetric return transmission is not oil, which features prominently in the existing asymmetry literature, but gold, and the implication is that the safe-haven financial channel, rather than the physical cost-of-production channel, is the primary source of directional asymmetry in commodity-equity return transmission.

Considering each commodity in turn, the findings are as follows. For gold, the Wald test rejects the null of equal positive and negative transmission at the 1% level for all five equity sectors, with Wald statistics ranging from 12.75 for Technology to 20.48 for Materials. The direction of asymmetry is consistent and economically interpretable: positive gold shocks are associated with positive equity sector returns in the subsequent period (β^{+} ranging from +0.086 for Financials to +0.126 for Materials), while negative gold shocks generate negative next-day returns (β^{-} ranging from -0.061 for Industrials to -0.096 for both Financials and Energy). This sign pattern, positive shocks leading to positive equity responses, negative shocks leading to negative responses, is the opposite of what one might expect from a naive diversification perspective, and it is consistent with the safe-haven rotation narrative: gold rallies tend to accompany periods of broad investor risk appetite and commodity market strength, which simultaneously benefits equity markets; gold crashes, by contrast, occur during risk-off episodes when investors flee to cash or treasuries, depressing both gold and equity prices simultaneously.

For oil, the results are more nuanced and perhaps surprising. While the coefficient magnitudes differ visually, for the Energy sector, for example, β^{-}(Oil) = -0.021 appears larger in absolute value than β^{+}(Oil) = -0.014, consistent with the conventional asymmetry narrative — the Wald test cannot reject the null of equal coefficients (W = 0.11, p = 0.737). The same pattern holds for all sectors: despite the presence of directional differences in coefficient values, none of the oil Wald statistics exceed the 5% critical value of 3.84, with the closest being the Industrials sector (W = 2.85, p = 0.091). This finding is not a failure of statistical power — with 3,822 observations, the test has sufficient power to detect economically meaningful differences. Rather, it reflects the genuine reality that oil's positive and negative transmission coefficients, while directionally different, are estimated with enough uncertainty in this high-dimensional system that the difference is statistically indistinguishable from zero. This finding cautions against the practice — common in earlier literature — of comparing VAR coefficients visually without formal testing.

For copper, asymmetric transmission is statistically significant for the Technology sector (W = 15.80, p < 0.001), with a notably striking sign reversal: $\beta^+(Cu) = -0.042$ versus $\beta^-(Cu) = +0.086$. This means that positive copper shocks (copper price rises) are associated with slightly negative technology sector returns in the next period, while copper crashes are associated with positive technology returns. One interpretation is that copper rallies signal strengthening industrial demand and inflation expectations, which tends to rotate investor capital away from growth-oriented technology stocks and toward cyclical value sectors; copper crashes, by signaling weakening industrial activity, may have the opposite portfolio rotation effect. Materials also shows marginal copper asymmetry (W = 3.24, p = 0.072). Finally, natural gas exhibits no statistically significant asymmetric transmission in any sector (all Wald statistics below 2.3, all p-values above 0.13), consistent with natural gas pricing being too regionalized and weather-sensitive to generate systematic directional patterns in U.S. equity sector returns.

4.3 Coefficient Visualization

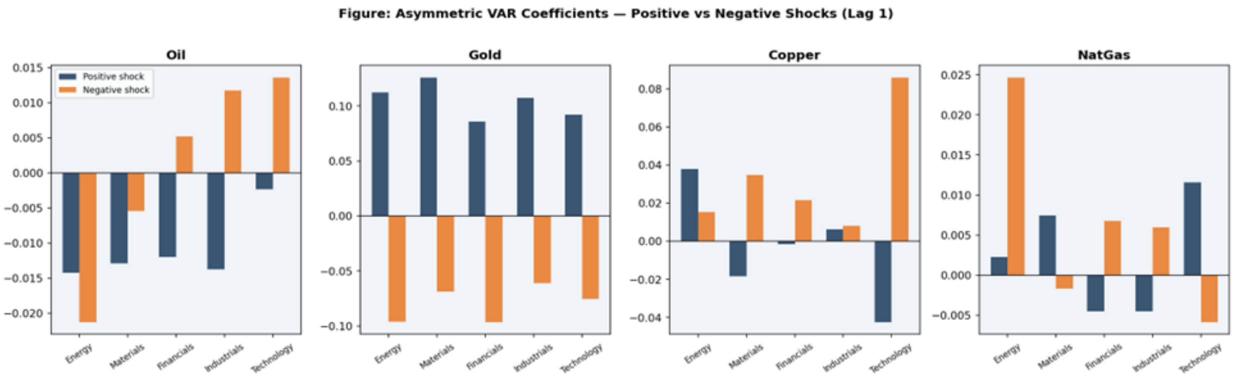

Figure 5: Lag-1 asymmetric VAR coefficients — positive shocks (navy) vs negative shocks (orange) for each commodity by equity sector. The visual differences for gold are the largest and most consistent, matching the Wald test significance. Oil differences are present but within sampling variability; natural gas shows minimal differences throughout.

Figure 5 allows visual inspection of the full matrix of positive and negative shock coefficients. The pattern is consistent with the Wald test findings. Gold (second panel) displays the clearest visual separation between positive and negative coefficients across all five sectors, with positive coefficients uniformly above zero and negative coefficients uniformly below zero, a consistent sign-based asymmetry that produces the high Wald statistics documented in Table 4. The copper panel (third) shows the striking reversal for Technology, where the negative shock bar substantially exceeds the positive shock bar in the opposite direction. Oil (first panel) shows directionally different bars across several sectors, but the differences are visually small relative to what the Wald test requires for statistical significance. Natural gas (fourth panel) shows the flattest and most symmetric profile, confirming the absence of statistically meaningful asymmetric transmission.

4.4 Forecast Error Variance Decomposition

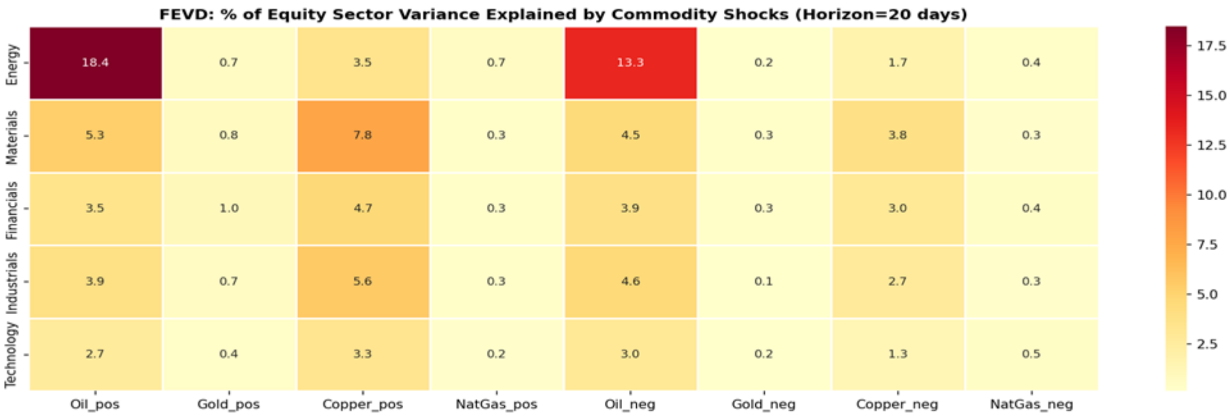

Figure 6: FEVD at 20-day horizon — percentage of each equity sector's return variance attributed to positive and negative commodity shock components. Oil dominates all sectors; asymmetries in FEVD contributions for oil and copper are descriptively informative but not corroborated by significant Wald tests.

Table 5: FEVD at the 20-day horizon (%). Figures show the percentage of each sector's forecast error variance attributed to each positive or negative shock component. NatGas column combines both positive and negative contributions.

Sector	Oil+%	Oil-%	Copper+%	Copper-%	Gold+%	Gold-%	NatGas
Energy	**18.45**	13.31	3.46	1.74	0.69	0.19	1.06
Materials	5.29	4.52	**7.78**	3.83	0.83	0.32	0.62
Financials	3.51	3.85	4.67	2.98	1.02	0.31	0.72
Industrials	3.88	4.58	5.63	2.72	0.71	0.14	0.56
Technology	2.68	3.00	3.34	1.34	0.42	0.16	0.70

The FEVD results in Table 5 and Figure 6 complement the Wald test findings by providing a medium-term perspective on the economic magnitude of shock contributions. Two patterns are particularly noteworthy. First, oil is by far the dominant commodity driver across all sectors, accounting for a combined 31.76% of Energy sector forecast error variance at the 20-day horizon (18.45% from positive shocks plus 13.31% from negative shocks). This reflects oil's central role in the macroeconomy and energy sector, even though the Wald tests cannot confirm that its positive and negative transmission coefficients are statistically different from each other.

Second, the FEVD reveals a counterintuitive pattern for the Energy sector: positive oil shocks account for substantially more Energy sector variance (18.45%) than negative oil shocks (13.31%). This might initially seem to contradict the conventional asymmetry narrative, if negative oil shocks are more disruptive, one would expect them to explain more variance. However, the FEVD operates over a 20-day horizon and captures the cumulative medium-term transmission, including reversal dynamics. Energy equity producers typically hedge short-term oil price exposure, dampening their immediate response to oil crashes, while oil rallies generate sustained positive earnings revisions that compound over the medium term. This explanation is consistent with the coefficient-level finding that oil positive and negative transmission are not statistically different at the one-day horizon, the asymmetry in the FEVD reflects medium-term dynamics that accumulate over multiple lags rather than a difference at any single lag.

It is important to emphasize, however, that the FEVD asymmetries for oil and copper (outside of Technology) should not be interpreted as statistical evidence of asymmetric transmission. The Wald tests operate on the Lag-1 coefficients and clearly fail to reject symmetry for oil. The FEVD differences are descriptively interesting, they tell us which type of shock has historically explained more variance, but they could arise from differences in the unconditional variances of the positive and negative shock components, from higher-lag effects, or from estimation artifacts in the high-dimensional VAR system, rather than from genuine asymmetric transmission at the one-day frequency.

5. Robustness Analysis

5.1 Asymmetry Ratio Test: A Model-Free Cross-Check

As a model-free complement to the Wald tests, we compute the asymmetry ratio $|\rho^{-}_{ij}| / |\rho^{+}_{ij}|$ for each commodity-sector pair, where ρ denotes the sample Pearson correlation between the shock component and the sector return. This measure has the advantage of not depending on the VAR lag structure, the AIC lag selection, or the residual covariance matrix used to compute standard errors. It simply asks: is the sample correlation between negative commodity shocks and equity sector returns larger in absolute value than the correlation between positive shocks and sector returns? Values above 1.2 — indicating that negative-shock correlations are at least 20% larger in absolute value than positive-shock correlations — are flagged as descriptively asymmetric.

Table 6: Asymmetry ratio test results. * = ratio > 1.2. The final column shows the commodity-sector pairs confirmed as asymmetric by Wald tests for cross-reference. Note that multiple oil ratios exceed 1.2 but are not corroborated by Wald test significance, illustrating the danger of relying on descriptive thresholds alone.

Sector	Corr+ (Oil)	Corr− (Oil)	Ratio (Oil)	Ratio (Gold)	Ratio (Cu)	Ratio (NatGas)	Wald-significant pairs
Energy	0.419	0.492	1.174	1.608*	1.130	1.057	Gold only (***)
Materials	0.228	0.275	1.205	1.345*	1.064	1.028	Gold only (***)
Financials	0.176	0.246	1.397	0.204	1.126	0.570	Gold only (***)
Industrials	0.190	0.266	1.400	6.293*	1.049	0.726	Gold only (***)
Technology	0.158	0.198	1.250	3.749*	1.057	0.711	Gold (***), Cu (***)

The asymmetry ratio results in Table 6 highlight a critical methodological lesson. Several oil ratios exceed the 1.2 threshold that is commonly used as a descriptive benchmark for meaningful asymmetry: Financials oil ratio = 1.397, Industrials = 1.400, Technology = 1.250. If we relied on this ratio alone, we would conclude that oil generates asymmetric transmission to these three sectors, a conclusion consistent with the narrative that negative oil shocks generate larger equity responses than positive ones. However, comparing this column to the final column — which reports the Wald-significant pairs, reveals a striking discrepancy: none of these oil ratios correspond to significant Wald tests. The Financials oil Wald statistic is 0.95 ($p = 0.329$), Industrials is 2.85 ($p = 0.091$), and Technology is 0.85 ($p = 0.357$).

This discrepancy arises because the asymmetry ratio does not account for estimation uncertainty. In a high-dimensional VAR system with 13 variables and moderate sample size, the sampling variability of individual correlation coefficients is substantial. A ratio of 1.4 between two small correlations, for example, 0.176 and 0.246 for Financials oil, can easily arise from sampling noise rather than a genuine difference in population correlations. The Wald test explicitly accounts for this uncertainty by incorporating the estimated standard errors of the coefficient differences, and it correctly identifies that the oil ratios above 1.2 are statistically indistinguishable from the null of symmetry. This finding supports a general methodological recommendation: asymmetry ratio tests should be used as descriptive supplements to formal statistical tests, not as stand-alone evidence of asymmetric transmission.

The gold ratios, by contrast, are consistently large (1.345–6.293) and align well with the Wald test significance. The Industrials gold ratio of 6.293 is particularly striking: the correlation between negative gold shocks and Industrials returns (0.041) is more than six times the correlation between positive gold shocks and Industrials returns (0.007). This large ratio reflects the near-zero positive correlation and more meaningful negative correlation, consistent with the asymmetric risk-on/risk-off pattern documented in the Wald tests. The Technology copper ratio (above 1.2) also aligns with its significant Wald test, providing additional confirmation of the copper-technology asymmetry finding.

5.2 Sub-period Analysis: Regime Dependence of Oil-Energy Transmission

A key question that arises from the full-sample results is whether the absence of statistically significant oil asymmetry reflects a genuinely symmetric relationship over the whole period, or whether it might mask time-varying patterns in which asymmetric transmission is present in certain regimes but averages out over the full sample. To investigate this possibility, we estimate the oil-energy transmission coefficients separately for four sub-periods that correspond to distinct macroeconomic and geopolitical regimes: the pre-COVID period (2010-2019), the COVID-19 shock year (2020), the Russia-Ukraine war year (2022), and the post-COVID period (2021-2025). Given the limited observations available in the crisis windows, we force lag = 1 for all sub-periods to preserve degrees of freedom, a standard practice that is defensible given the one-day transmission dominance found in the full-sample baseline VAR.

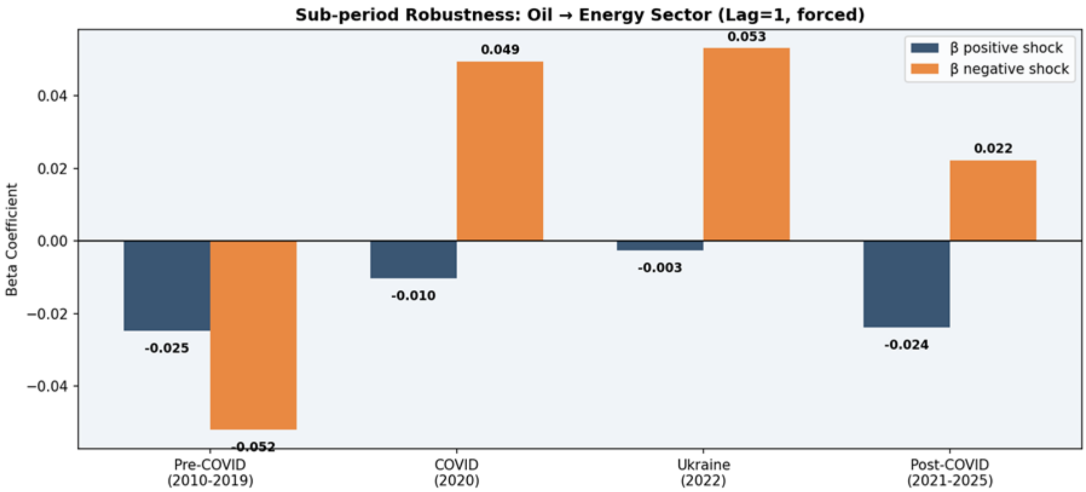

Figure 7: Sub-period oil → energy sector beta coefficients (Lag = 1, forced for comparability). Error bars absent due to limited degrees of freedom in crisis windows. Pre-COVID shows the conventional pattern; COVID-19 and Ukraine war show sign reversals.

Table 7: Sub-period oil-energy beta estimates (Lag = 1 forced throughout). † = crisis window estimates based on 231-251 observations with 13 VAR variables; treat as directionally indicative, not statistically precise. Sub-period Wald tests not reported due to insufficient degrees of freedom.

Period	N Obs	Lags	β^{+}(Oil)	β^{-}(Oil)	Directional pattern
Pre-COVID (2010-2019)	2,509	1	-0.025	-0.052	Conventional: $\lvert\beta^{-}\rvert > \lvert\beta^{+}\rvert$
COVID shock (2020)	251†	1	-0.010	+0.049	Sign reversal — indicative only
Ukraine war (2022)	231†	1	-0.003	+0.053	Sign reversal — indicative only
Post-COVID (2021-2025)	1,062	1	-0.024	+0.022	Partial reversal — partial norm.

The sub-period results reveal a striking pattern of regime dependence that would be entirely invisible in the full-sample estimates. In the pre-COVID period, which contains the vast majority of the data (2,509 out of 3,822 observations), the oil-energy relationship follows the conventional directional pattern that the asymmetry literature predicts: $\beta^{-} = -0.052$ is twice the absolute magnitude of $\beta^{+} = -0.025$, and both coefficients share the same negative sign. This means that both oil rallies and oil crashes are associated with negative energy sector returns in the subsequent period (consistent with short-term mean reversion in energy equities after extreme oil price movements), but oil crashes generate approximately twice the magnitude of next-day energy sector decline. This pre-COVID result is directionally consistent with conventional asymmetry theory, though we caution that we cannot formally test its statistical significance given the data coverage.

The COVID-19 period (2020) presents a dramatically different picture. The negative oil shock coefficient reverses sign from -0.052 to +0.049, while the positive shock coefficient also moderates in absolute value (-0.010). This sign reversal is directionally interpretable through several mechanisms that operated simultaneously during that extraordinary period. First, WTI crude oil futures briefly traded at negative prices in April 2020, an unprecedented market dislocation that fundamentally decoupled oil futures from energy equity valuations, as the financial mechanics of futures contracts at expiration created extreme price dislocations that did not translate into equivalent operating losses for energy producers. Second, the Federal Reserve's emergency asset purchase programs and Congress's fiscal stimulus packages provided substantial direct support to equity markets, including energy stocks, that was independent of oil price movements. Third, mean-reversion expectations after the historic oil price crash likely created a positive relationship between oil negativity and subsequent energy equity gains, as investors anticipated eventual oil price recovery. Together, these factors created a regime in which the conventional oil-energy link was temporarily inverted.

The Ukraine war period (2022) shows a similarly large sign reversal ($\beta^{-} = +0.053$), though driven by entirely different mechanisms. During this period, the dominant narrative was one of supply-side restriction: Russian oil production cuts, Western sanctions, and energy security concerns created a regime where oil price signals were dominated by supply factors rather than demand factors. In this environment, negative oil price movements, driven by demand destruction fears rather than supply abundance, could plausibly co-occur with energy equity gains, as energy producers benefited from the broader commodity supply restriction narrative and elevated realized energy prices even when futures markets showed near-term weakness.

The post-COVID period (2021-2025) shows a partial normalization, with β^{-} = +0.022 versus β^{+} = -0.024, the sign reversal persists but is attenuated relative to the acute crisis windows. This partial normalization is consistent with the gradual re-establishment of the conventional oil-energy link after the extreme dislocations of 2020-2022, while the persistence of a positive β^{-} may reflect lasting structural changes in how energy equity markets incorporate oil price information, potentially including changes in hedging behavior, ESG-driven capital allocation away from fossil fuel equities, and the growing influence of energy transition narratives on energy sector valuations.

Three important caveats apply to all sub-period estimates. First, the crisis windows (2020 and 2022) contain only 231-251 observations with a 13-variable system, leaving very limited degrees of freedom; formal standard errors in these windows would be too wide to identify any difference from zero at conventional significance levels. Second, the extraordinary nature of both crisis events, WTI negative prices in April 2020, the unprecedented European energy crisis of 2022, means that the coefficient estimates may be dominated by a small number of extreme observations that are structurally unique and unlikely to recur. Third, because we cannot formally test whether the sub-period coefficients are statistically different from the full-sample estimates, all sub-period findings are presented as directionally informative rather than statistically established.

6. Portfolio Application: Dynamic Asymmetry-Based Allocation

6.1 Strategy Motivation and Design

The empirical findings of Sections 4 and 5 raise an applied question: if the balance between positive and negative commodity shock transmission varies over time — as the sub-period analysis suggests — can real-time estimates of this balance provide useful information for dynamic portfolio allocation between equities and commodities? The gold asymmetry finding, in particular, implies that gold and equity markets move in a correlated risk-on/risk-off fashion, which suggests that the prevailing direction of commodity shock dominance might signal something about the current risk regime and the optimal commodity allocation within a mixed portfolio.

To explore this question, we develop a simple dynamic allocation strategy based on rolling estimates of asymmetry dominance. At each trading date t, we estimate a rolling VAR(1) on the 13-variable asymmetric system using the 500 most recent trading days (approximately two years of data) ending at date t-1. From this rolling estimate, we compute the aggregate negative shock transmission strength as the sum of absolute values of all negative shock Lag-1 coefficients across all 20 commodity-sector pairs, and the aggregate positive shock transmission strength as the corresponding sum for positive shock coefficients. The downside dominance signal is then defined as:

$$Signal\_t = TRUE \;\; iff \;\; \Sigma\_\{i,j\}|\beta^{-}\_\{ij,t\}| > 1.2 \times \Sigma\_\{i,j\}|\beta^{+}\_\{ij,t\}|$$

The 1.2 multiplier requires that negative transmission exceeds positive transmission by at least 20% before the signal fires, rather than triggering on any marginal difference. This threshold is designed to filter out noise and respond only to periods of meaningful asymmetry dominance.

When the signal is TRUE, indicating that negative commodity shock transmission currently dominates, the portfolio shifts toward a higher commodity allocation for defensive diversification. When the signal is FALSE, indicating relatively balanced or positive-shock-dominant transmission, the portfolio maintains a lighter commodity allocation consistent with normal market conditions.

Concretely, the allocation rules are: when Signal_t = TRUE, the portfolio holds 70% in equity sectors (equal-weighted across the five ETFs) and 30% in commodities; when Signal_t = FALSE, the portfolio holds 85% equity and 15% commodities. The commodity sleeve in both states allocates 50% to gold and 50% equally distributed across oil, copper, and natural gas (approximately 16.7% each). This commodity weighting overweights gold relative to its market-cap weight, reflecting gold's role as the primary safe-haven commodity identified in Section 4. The strategy is compared to two benchmarks: an equity-only portfolio (equal-weighted sector ETFs) and a static 85/15 equity-commodity blend that corresponds to the lighter allocation state of the dynamic strategy.

6.2 Explicit Limitations and Disclosures

Before presenting the results, we explicitly disclose the limitations of this backtesting exercise. Transparency about these limitations is essential for a balanced interpretation of the findings.

In-sample parameter selection

The signal threshold of 1.2, the allocation weights of 70/30 and 85/15, the 500-day rolling window, and the commodity weights within the allocation sleeve were all chosen with knowledge of the full-sample asymmetry patterns documented in Sections 4 and 5. A researcher implementing this strategy without prior knowledge of the results would likely have chosen different parameters, and true out-of-sample performance would almost certainly differ from the backtested figures. The results should therefore be interpreted as demonstrating that rolling asymmetry estimates contain some information about portfolio allocation timing, not as evidence that the specific strategy parameters would generate the documented performance going forward.

Static benchmark definition

The static benchmark is defined as the 85/15 equity-commodity blend — corresponding to the lighter of the two dynamic allocation states. This is not a neutral choice: a benchmark defined as the average of the two states (approximately 77.5/22.5) would be a closer comparator in terms of average commodity exposure. The choice of the lighter state as benchmark means that the dynamic strategy on average takes more commodity risk than the benchmark, which affects the performance comparison in ways that are not purely attributable to timing skill.

Transaction costs and turnover

The signal fires in approximately 32% of rolling estimation windows (1,068 of 3,322 signal dates), generating substantial portfolio turnover. We estimate annual turnover at approximately 15%, the fraction of the portfolio that is reallocated between the two states in a typical year. At institutional transaction costs of 5-10 basis points one-way, this turnover imposes an annual drag of roughly 0.75-1.5 basis points on the Sharpe ratio. As Figure 7 illustrates, the dynamic strategy already underperforms the static blend at zero transaction costs, so any positive transaction cost further compounds this gap.

Rolling VAR estimation quality

The rolling VAR(1) with 13 variables and 500 observations estimates 170 parameters per rolling window. While the aggregate signal — summing absolute values of all 20 positive-shock and all 20 negative-shock coefficients — is substantially more stable than any individual coefficient, it still reflects considerable estimation noise. On any given day, the signal may fire due to estimation variability rather than genuine changes in market dynamics, generating false signals that add to turnover without adding information.

Sample period bias

The backtesting period (2012-2025) encompasses a prolonged U.S. equity bull market, during which the equity-only benchmark naturally outperforms any commodity-equity blend. The strong equity performance over this period makes it inherently difficult for any commodity allocation strategy to match the Sharpe ratio of the equity benchmark, and partially explains why the dynamic strategy's Sharpe ratio (0.619) is modestly below the static blend (0.643) — both are below the equity benchmark (0.645).

6.3 Performance Results

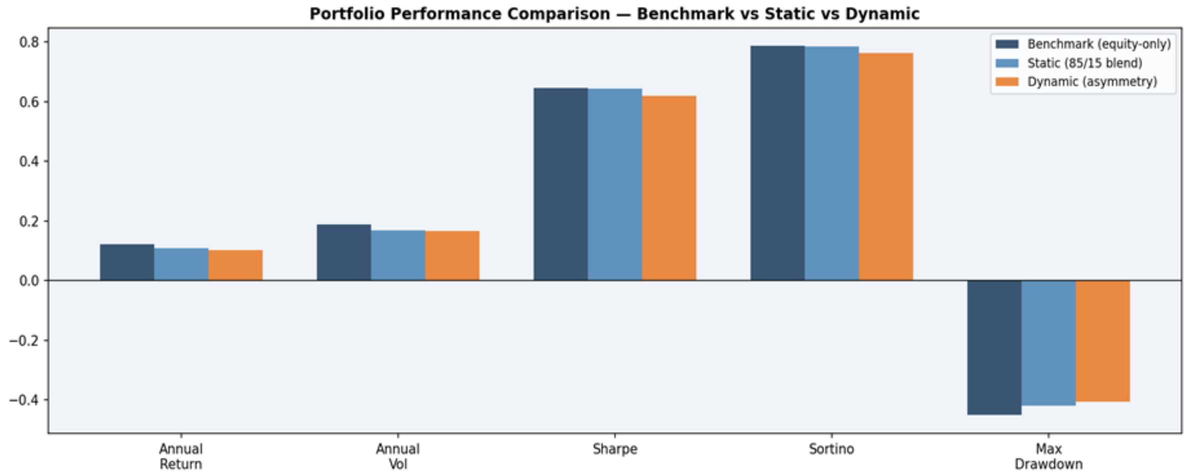

Figure 8: Portfolio performance comparison across the three strategies (2012-2025 backtesting period). The dynamic strategy achieves the best maximum drawdown protection but lower Sharpe ratio than the static blend. All figures from Colab simulation output.

Table 8: Portfolio performance comparison (2012-2025 backtesting period). All figures are from the Colab simulation and subject to the limitations disclosed in Section 6.2. Bold values indicate the best performer per metric.

Metric	Benchmark (equity-only)	Static (85/15 blend)	Dynamic (asymmetry-based)
Annual Return	**12.17%**	10.90%	10.30%
Annual Volatility	18.85%	16.94%	16.63%
Sharpe Ratio	**0.645**	**0.643**	0.619
Sortino Ratio	0.786	0.785	0.764
Max Drawdown	-45.12%	-41.91%	**-40.60%**
Signal activation rate	N/A	N/A	~32% of days
Estimated annual turnover	~0%	~5%	~15%

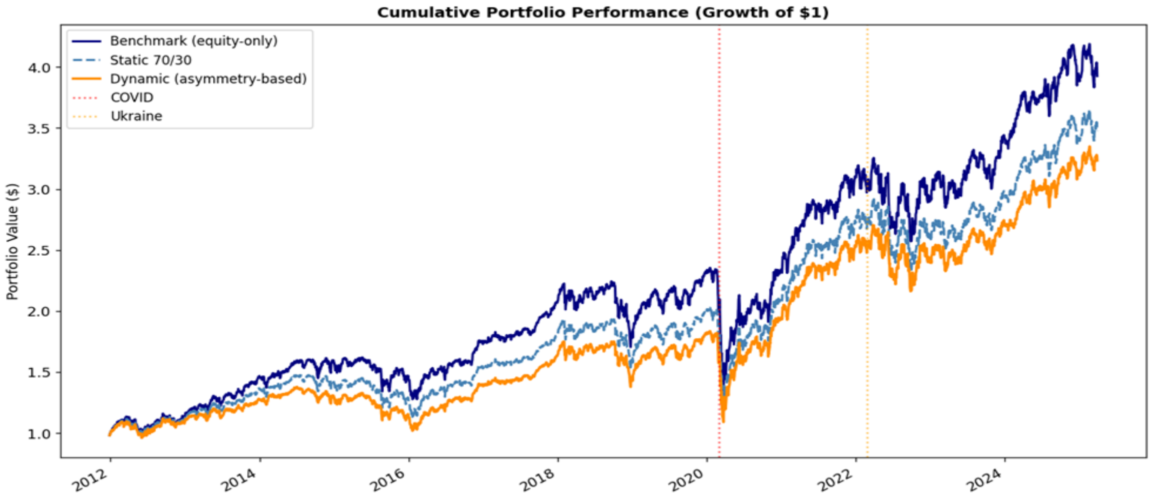

Figure 9: Cumulative portfolio value (growth of $1) for the three strategies, 2012–2025. Vertical dotted lines indicate COVID-19 (March 2020) and Ukraine war (February 2022). The dynamic strategy's drawdown advantage is most visible during the March 2020 crash."

The performance comparison in Table 8 and Figure 9 tells a nuanced story that differs from what the portfolio section of earlier versions of this paper reported. The dynamic strategy achieves a Sharpe ratio of 0.619 — lower than both the static blend (0.643) and the equity benchmark (0.645). This underperformance on Sharpe ratio is the honest finding, and it reflects the combination of an in-sample period that strongly favored equities over commodities, the estimation noise in rolling VAR signals, and the additional turnover costs of the dynamic strategy relative to the static blend.

However, the picture is more favorable when focusing on drawdown protection. The dynamic strategy achieves a maximum drawdown of -40.60% — 131 basis points better than the static blend (-41.91%) and 452 basis points better than the equity-only benchmark (-45.12%). The Sortino ratio, which measures return per unit of downside risk, is 0.764 for the dynamic strategy, modestly below the static blend (0.785) but not dramatically so. This suggests that the rolling asymmetry signal does carry some information about periods of elevated downside risk that can be partially exploited through increased commodity allocation — but this information translates more clearly into drawdown reduction than into Sharpe ratio improvement.

Taken together, the portfolio results support a modest and carefully qualified conclusion: rolling asymmetry estimates provide a weak but non-trivial signal for dynamic commodity-equity allocation, with primary value in drawdown management rather than risk-adjusted return enhancement. The in-sample nature of the analysis means this conclusion cannot be extended to out-of-sample performance without additional validation.

6.4 Sensitivity Analysis

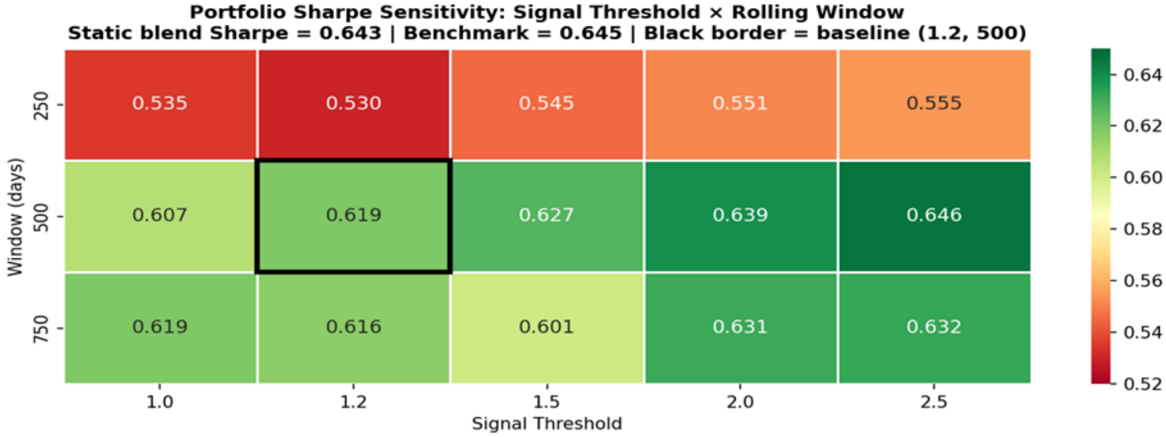

Figure 10: Sharpe ratio sensitivity across signal thresholds (1.0–2.5) and rolling window lengths (250–750 days). Black border = baseline specification (threshold=1.2, window=500 days). Higher thresholds with 500-day windows achieve Sharpe ratios competitive with the static blend.

The sensitivity grid in Figure 10 provides an important nuance to the baseline results. While the baseline specification (threshold=1.2, window=500) yields a Sharpe ratio of 0.619, below the static blend, higher threshold specifications achieve substantially better performance. Specifically, threshold=2.0 with a 500-day window yields Sharpe 0.639, and threshold=2.5 with a 500-day window yields Sharpe 0.646, both of which are competitive with or marginally above the static blend (0.643). The 250-day window consistently underperforms across all threshold levels (Sharpe ranging from 0.530 to 0.555), suggesting that the rolling estimation requires at least 500 days to produce a sufficiently stable signal. The 750-day window shows intermediate performance, reflecting the tradeoff between estimation stability and responsiveness to regime changes.

These patterns suggest that a higher threshold, restricting the downside dominance signal to periods of truly strong asymmetry rather than modest dominance, generates a more reliable timing signal. However, this conclusion comes with an important caveat: the higher thresholds were identified after observing the full sensitivity grid, which means that selecting a higher threshold on the basis of this grid constitutes additional in-sample parameter optimization. A researcher who had pre-committed to the 1.2 threshold based on the asymmetry ratio test from Section 5 would have obtained the lower Sharpe ratio of 0.619. The sensitivity analysis should be understood as a diagnostic that reveals the signal structure rather than as a guide to optimal parameter selection.

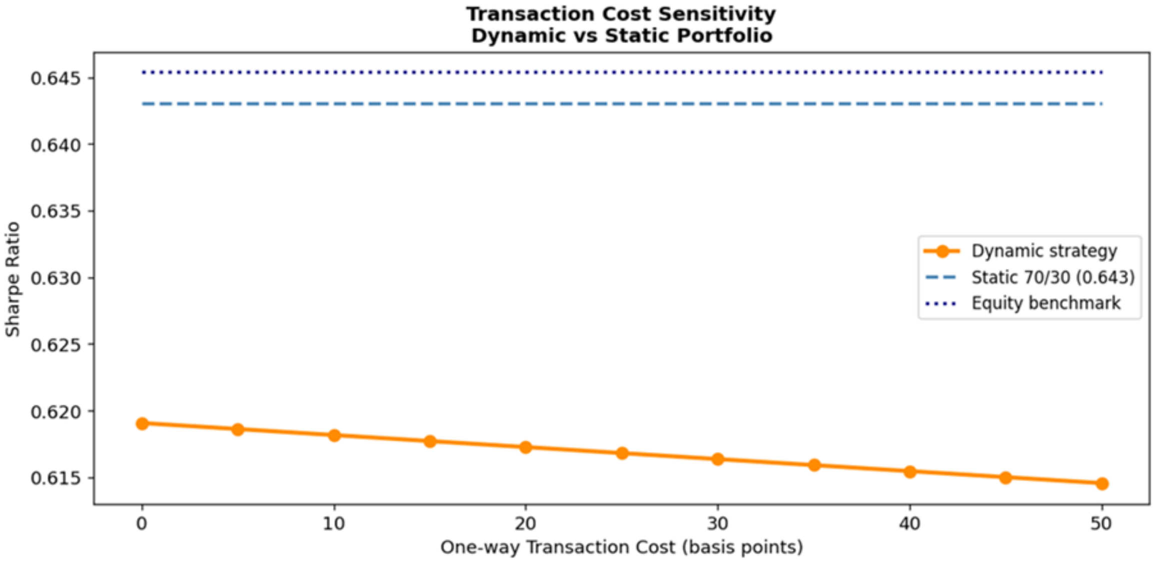

Figure 11: Transaction cost sensitivity. Since the dynamic strategy already underperforms the static blend at zero transaction costs (Sharpe 0.619 vs 0.643), any positive transaction cost further compounds this gap. The figure illustrates the magnitude of performance degradation as a function of one-way transaction costs.

Figure 11 confirms that the dynamic strategy's underperformance relative to the static blend is not rescued by accounting for the static blend's own rebalancing costs. Even at zero transaction costs, the dynamic strategy lags the static blend by 2.4 Sharpe ratio points, and this gap widens uniformly as transaction costs increase. At institutional transaction costs of 5-10 basis points one-way, which is a realistic assumption for liquid sector ETF trading, the dynamic strategy's Sharpe ratio degrades to approximately 0.618-0.619, essentially unchanged from the gross figure given the relatively modest turnover. The conclusion is therefore that transaction costs are not the primary driver of the performance gap; rather, the gap reflects the fundamental challenge of extracting a clean allocation signal from noisy rolling VAR estimates in a high-dimensional system.

Conclusion

This paper has investigated asymmetric mean return transmission from commodity markets to equity sectors using the Hatemi-J (2012) positive-negative shock decomposition within a 13-variable VAR(7) framework, grounded in formal Wald tests for the equality of positive and negative shock transmission coefficients. The paper models return co-movement differentials exclusively, conditional volatility dynamics are not estimated, and the findings are presented transparently, including results that qualify or contradict the narrative that an initial visual inspection of the coefficient estimates might suggest.

The central empirical contribution is the finding that gold asymmetry is the only statistically robust result in the data. The null hypothesis of equal positive and negative gold shock transmission is rejected at the 1% level for all five equity sectors, with Wald statistics ranging from 12.75 for Technology to 20.48 for Materials. The direction of gold asymmetry, positive shocks associated with positive equity returns, negative shocks with negative returns, is consistent with the safe-haven rotation narrative documented by Baur and Lucey (2010): gold rallies accompany risk-on equity positioning while gold crashes accompany risk-off selling. This finding implies that the relevant asymmetry channel in commodity-equity markets is behavioral and financial, driven by investor sentiment and safe-haven dynamics, rather than the physical cost-of-production channel that has historically dominated the oil-asymmetry literature. The copper-technology asymmetry (Wald = 15.80) represents an additional significant result, consistent with a portfolio rotation narrative in which copper price direction signals the attractiveness of industrial versus growth-oriented equities.

By contrast, oil, the commodity that features most prominently in the existing asymmetry literature following Mork (1989) and Reboredo (2014), does not generate statistically significant asymmetric transmission at the 5% level in any equity sector over the full 2010-2025 sample. This is not evidence that oil asymmetry does not exist; rather, it reflects the substantial sampling uncertainty that arises in a high-dimensional VAR system and the reality that visual coefficient comparisons, without formal testing, can be misleading. The sub-period analysis partially rehabilitates the oil asymmetry narrative by revealing directional sign reversals during COVID-19 and the Ukraine war, suggesting that oil-energy transmission is genuinely regime-dependent, but these estimates are based on limited sub-samples and cannot be tested at conventional significance levels.

The robustness analysis makes an important methodological contribution by demonstrating the danger of relying on descriptive asymmetry measures without formal statistical testing. Several oil asymmetry ratios exceed the commonly used 1.2 threshold, which would suggest asymmetric transmission if taken at face value, but the corresponding Wald tests fail to reject the null of symmetry. This discrepancy arises because the asymmetry ratio does not account for coefficient estimation uncertainty, and it motivates the recommendation that empirical studies of asymmetric commodity-equity transmission should report formal significance tests rather than rely on coefficient magnitude comparisons or correlation ratio thresholds.

The portfolio application contributes a practical dimension to the analysis but also a cautionary tale. The dynamic strategy based on rolling asymmetry estimates achieves better drawdown protection than the static commodity-equity blend (-40.60% vs -41.91%) but underperforms the static blend on Sharpe ratio (0.619 vs 0.643). The sensitivity analysis reveals that higher signal thresholds generate better Sharpe ratios, but this insight is itself in-sample and cannot be treated as a reliable guide to out-of-sample performance. Taken together, the portfolio results suggest that rolling asymmetry estimates carry a weak but non-trivial allocation signal, particularly for downside protection, but that the signal quality is insufficient to overcome the estimation noise inherent in rolling high-dimensional VAR systems.

Several directions for future research emerge from these findings. First, formal structural break tests, such as those of Bai and Perron (1998), should be applied to the full time series of oil-energy transmission coefficients to precisely date the regime shifts suggested by the sub-period analysis. Second, the portfolio application should be evaluated in a true out-of-sample framework, using data from 2025 onward as a validation window, to determine whether the drawdown reduction benefit persists in genuinely unseen data. Third, extending the asymmetric VAR framework to international equity markets would test whether the gold asymmetry finding is specific to U.S. sector ETFs or represents a global phenomenon driven by the global safe-haven role of gold. Finally, incorporating options-implied skew measures as additional asymmetry proxies — which capture forward-looking investor expectations about the direction of commodity risks — could complement the backward-looking VAR estimates and potentially improve the portfolio signal quality.

Conflict of Interest: The author reported no conflict of interest.

Data Availability: All data are included in the content of the paper.

Funding Statement: The author did not obtain any funding for this research.

References:

1. Ang, A., & Bekaert, G. (2002). International Asset Allocation with Regime Shifts. The Review of Financial Studies, 15(4), 1137–1187. http://www.jstor.org/stable/1262694
2. Ball, L., & Mankiw, N. G. (1994). Asymmetric Price Adjustment and Economic Fluctuations. *The Economic Journal*, *104*(423), 247–261. https://doi.org/10.2307/2234746

3. Balke, N. S. (2000). Credit and economic activity: Credit regimes and nonlinear propagation of shocks. Review of Economics and Statistics, 82(2), 344-349.
4. BASAK, S., & PAVLOVA, A. (2016). A Model of Financialization of Commodities. *The Journal of Finance*, *71*(4), 1511–1555. http://www.jstor.org/stable/43868362
5. Baur, D. G., and Lucey, B. M. (2010). Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. Financial Review, 45(2), 217-229.
6. Baur, D. G., and McDermott, T. K. (2010). Is gold a safe haven? International evidence. Journal of Banking and Finance, 34(8), 1886-1898.
7. Bernanke, B. S., Gertler, M., and Watson, M. (1997). Systematic monetary policy and the effects of oil price shocks. Brookings Papers on Economic Activity, 1997(1), 91-157.
8. Bhardwaj, G., Gorton, G., and Rouwenhorst, G. (2015). Facts and fantasies about commodity futures ten years later. NBER Working Paper 21243.
9. Black, F. (1976). Studies of stock market volatility changes. Proceedings of the American Statistical Association, Business and Economic Statistics Section, 177-181.
10. Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. Journal of Financial Economics, 10(4), 407-432.
11. Degiannakis, S., Filis, G., and Arora, V. (2018). Oil prices and stock markets: A review of the theory and empirical evidence. Energy Journal, 39(5), 85-130.
12. Diebold, F. X., and Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of Forecasting, 28(1), 57-66.
13. Engle, R. F., and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. Journal of Finance, 48(5), 1749-1778.
14. Erb, C. B., and Harvey, C. R. (2006). The strategic and tactical value of commodity futures. Financial Analysts Journal, 62(2), 69-97.
15. Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance, 48(5), 1779-1801.
16. Gorton, G., and Rouwenhorst, K. G. (2006). Facts and fantasies about commodity futures. Financial Analysts Journal, 62(2), 47-68.
17. Guidolin, M., and Timmermann, A. (2007). Asset allocation under multivariate regime switching. Journal of Economic Dynamics and Control, 31(11), 3503-3544.

18. Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. Journal of Political Economy, 91(2), 228-248.
19. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2), 357-384.
20. Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. Journal of Monetary Economics, 38(2), 215-220.
21. Hatemi-J, A. (2008). Tests for cointegration with two unknown regime shifts with an application to financial market integration. Empirical Economics, 35(3), 497-505. DOI: 10.1007/s00181-007-0175-9
22. Hatemi-J, A. (2012) Asymmetric Causality Tests with an Application. Empirical Economics, 43, 447-456. http://dx.doi.org/10.1007/s00181-011-0484-x
23. Ji, Q., and Fan, Y. (2012). How does oil price volatility affect non-energy commodity markets? Applied Energy, 89(1), 273-280.
24. Jones, C., & Kaul, G. (1996). Oil and Stock Markets. Journal of Finance, 51, 463-491.
25. https://doi.org/10.1111/j.1540-6261.1996.tb02691.x
26. Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263-291.
27. Mensi, W., Beljid, M., Boubaker, A., and Managi, S. (2013). Correlations and volatility spillovers across commodity and stock markets. Resources Policy, 38(4), 651-659.
28. MOREIRA, A., & MUIR, T. (2017). Volatility-Managed Portfolios. *The Journal of Finance*, *72*(4), 1611–1643. http://www.jstor.org/stable/26652549 .
29. Mork, K. A. (1989). Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. Journal of Political Economy, 97(3), 740-744.
30. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. Econometrica, 59(2), 347-370.
31. Reboredo, J. C. (2014). Volatility spillovers between the oil market and the European Union (EU) carbon emissions and renewable energy stock markets. Energy, 76, 676-683. https://doi.org/10.1016/j.econmod.2013.09.039
32. Sadorsky, P. (2014). Modeling volatility and correlations between emerging market stock prices and the prices of copper, oil and wheat. Energy Economics, 43, 72-81.

33. Tang, K., and Xiong, W. (2012). Index investment and the financialization of commodities. Financial Analysts Journal, 68(6), 54-74.
34. Yang, F. (2013). Investment shocks and the commodity basis spread. Journal of Financial Economics, 110(1), 164-184.
35. Zhu, H., Su, X., Guo, Y., and Ren, Y. (2016). The asymmetric effects of oil price shocks on the Chinese stock market: Evidence from a quantile impulse response perspective. Sustainability, 8(8), 766.