

The Fisher–Whittle Filter: Detection of Deterministic Periodic Components in Economic Time Series^{*}

Nicolás Ronderos

Department of Economics, Pontificia Universidad Javeriana

Calle 40 No. 6-23

Email: nronderos@javeriana.edu.co

Orcid: 0000-0002-3447-445

05/2026

Abstract

Economic time series often carry deterministic periodic components of institutional origin: seasonal patterns, billing and settlement cycles, reserve-maintenance rhythms, and the calendar of scheduled policy meetings. Distinguishing these fixed-frequency components from genuinely stochastic dynamics is a measurement problem that conventional spectral methods do not resolve, because the classical periodicity test loses calibration once the noise background is serially correlated. This paper develops a test of periodicity for economic time series that remains valid under such backgrounds. The procedure whitens the periodogram with a maximum-likelihood ARMA spectrum, applies a modification of Fisher's g-test to the whitened ordinates, and selects significant frequencies through Hochberg's step-up procedure with family-wise error control — strict protection against false positives being the appropriate standard when each detected frequency is read as evidence of a specific institutional rhythm. Three propositions establish the test's asymptotic size, its local power, and its control of the family-wise error rate. A taxonomy of thirty U.S. macroeconomic series shows that the procedure detects periodicity precisely where institutional calendars are known to operate — unadjusted money aggregates, reserve balances, commercial paper, initial claims, and the daily two-year Treasury yield — and nowhere else. A complementary cross-section of ninety-five daily financial return series, for which weak-form efficiency predicts an essentially flat spectrum, returns detections at a rate statistically indistinguishable from noise, with interpretable structure confined to the commodity futures whose seasonal and reporting calendars are independently documented. An application to the Lucas–Barro natural-rate hypothesis shows how the detected component yields a measure of anticipated money growth that does not depend on the choice of conditioning variables.

Keywords: periodicity detection; spectral analysis; Fisher's g-test; family-wise error rate; seasonality and calendar effects.

JEL classification: C12; C22; C32; C52; E32.

^{*} The author received no financial support for the research, authorship, or publication of this article. The author declares that there are no conflicts of interest related to this research. All remaining errors are the author's own. Replication material are provided in the supplementary material. The data used in the empirical applications are included there together with the scripts required to reproduce the tables and figures. The code implementing the Fisher-Whittle filter is available through the Eviews Add-ins repository. Once installed, the procedure can be executed either through the point-and-click interface or directly from the command line using `{%series}.fwfilter`.

1. Introduction

Economic time series frequently contain deterministic periodic components whose frequency is fixed by an external calendar rather than generated by stochastic dynamics. Seasonal patterns in unadjusted aggregates, monthly reserve-maintenance cycles in bank balance sheets, scheduled rollover in money-market instruments, and the biweekly rhythm of Federal Open Market Committee operations all leave discrete spectral signatures of institutional origin. Detecting these components is a measurement problem in its own right, distinct from trend-cycle decomposition: the question is not how to partition variance into trend and cycle, but whether spectral mass concentrated at a specific frequency exceeds what a stochastic background alone would produce. Whether economic series harbour deterministic structure at all has a long history in the dynamical-systems literature (Brock and Hsieh, 1991; Hsieh, 1993; LeBaron, 1994); that work sought low-dimensional chaotic dynamics and found little that replicated, whereas the deterministic structure at issue here is of transparent institutional origin and fixed frequency, and the question is simply whether the evidence for it at a given frequency exceeds the stochastic background.

The classical instrument for this question is Fisher's (1929) g-test, which compares the largest periodogram ordinate to the sum of all ordinates. Its exact null distribution, however, presumes white noise, and Whittle (1952) showed that the test is badly miscalibrated when the background is serially correlated — as economic time series invariably are. Detecting periodicity in economic data therefore requires a test that accounts for the colour of the noise. The combination of parametric prewhitening with Fisher-type testing has been developed in other fields — for gene expression by Wichert, Fokianos, and Strimmer (2004) and Ahdesmäki et al. (2005, 2007), and for stellar photometry by Reegen (2007) and Vaughan (2010) — but with backgrounds and multiplicity corrections suited to those settings rather than to economics.

This paper develops a periodicity test for stationary economic time series. A maximum-likelihood ARMA model supplies the spectral background; dividing the periodogram by the estimated spectrum yields whitened ordinates that are asymptotically independent and exponentially distributed under the null, regardless of the colour of the underlying process. A modification of Fisher's g-test is applied to these ordinates, and the multiplicity inherent in scanning the full frequency grid is resolved by Hochberg's (1988) step-up procedure, which controls the family-wise error rate. We choose family-wise control rather than the false-discovery-rate control common in high-dimensional screening because the economic setting is different in kind: each frequency declared significant is interpreted as a specific institutional rhythm, so a false positive carries an interpretive cost that warrants strict protection rather than a tolerated error fraction. The detected frequencies can be inverted back to the time domain to recover the deterministic component itself, when that component is the object of interest.

Reliable periodicity detection in economic data has several uses. It separates calendar and seasonal structure from genuinely stochastic variation, providing a model-free check on the adequacy of seasonal adjustment. It identifies regulatory and settlement rhythms in financial data that would bias inference if treated as innovations. It supplies a measure of predictable variation that requires no choice of conditioning information set — relevant to the rational-expectations tradition, in which only the unpredictable component of a policy variable should have real effects (Muth, 1961; Lucas, 1972; Sargent and Wallace, 1975), and which we take up in the natural-rate application of Section 7.1. And the presence of a detected component in a series of supposedly stochastic origin is itself a diagnostic of misspecification. In financial markets, where weak-form efficiency implies that returns carry no own-past predictability, the test doubles as a specificity check: a detector that found periodicity where none should exist would be of little value, and the cross-section of Section 7.3 confirms that ours does not.

The contribution is fourfold: a periodicity test combining ARMA-based spectral estimation, a modified Fisher g-test, and Hochberg multiplicity correction; three asymptotic results on size, local power, and family-wise error control; an extension to richer parametric backgrounds, including ARFIMA for long memory, documented in Appendix B; and an empirical mapping of where deterministic structure is, and

is not, present: a taxonomy of thirty U.S. macroeconomic series, complemented by a cross-section of ninety-five daily financial return series that supplies a stringent specificity check in a setting where efficient-markets theory predicts an essentially flat spectrum. The remainder of the paper proceeds as follows. Section 2 reviews the relevant literatures, Section 3 develops the framework, Section 4 states the main results with proofs in Appendix A, Section 5 describes implementation, Section 6 reports Monte Carlo evidence, Section 7 presents three empirical applications — a test of the natural-rate hypothesis, a macroeconomic taxonomy, and a financial cross-section — and Section 8 concludes.

2. Related Literature

The procedure proposed in this paper sits at the intersection of three distinct literatures, each of which we review briefly.

2.1 Trend–cycle filtering as background

Filtering procedures that decompose macroeconomic series into trend and cycle components — the Hodrick–Prescott (1980, 1997), Baxter–King (1999), Christiano–Fitzgerald (2003), and Hamilton (2018) approaches, together with the model-based unobserved-components route of Harvey and Trimbur (2003) — define the cyclical component by mechanical operation or by an estimated state-space model rather than by inference on individual frequencies. The business-cycle band these methods target traces to the durations catalogued by Burns and Mitchell (1946), and the procedures are known to distort integrated processes: detrending a difference-stationary series induces spurious periodicity (Nelson and Kang, 1981; King and Rebelo, 1993; Cogley and Nason, 1995; Harvey and Jaeger, 1993), the Hamilton filter has its own documented artefacts (Schuler, 2021), and the stylised facts recovered from a series depend materially on which filter is applied (Canova, 1998). These procedures address a question different from the one taken up here — they aim to estimate a smooth trend and a residual cycle, rather than to identify specific frequencies as deterministic — but they share with the present paper an interest in the frequency-domain content of economic time series.

2.2 Frequency-domain filtering

An alternative implementation operates directly on the Fourier coefficients. With $A_j = \sum_{t=1}^N X_t \cos(\omega_j t)$ and $B_j = \sum_{t=1}^N X_t \sin(\omega_j t)$, a band-pass filter selecting an index set $\mathcal{S} \subset \{1, \dots, m\}$ reconstructs the cyclical component as

$$C_t = \frac{2}{N} \sum_{j \in \mathcal{S}} A_j \cos(\omega_j t) + B_j \sin(\omega_j t)$$

Pollock (2007, 2014, 2018) has been the most consistent advocate of this approach within econometrics, demonstrating through the IDEOLOG software that frequency-domain filtering can be implemented as a transparent operation on the discrete Fourier transform. Iacobucci and Noullez (2005) develop a related selective filter for short series, and Hinich, Foster, and Wild (2009) document the cycle-extraction properties of discrete Fourier transform filters. These contributions establish the operational viability of Fourier-based reconstruction; they leave open the question of which Fourier ordinates to retain. The present paper proposes a data-driven answer: the index set \mathcal{S} is determined by a multiplicity-controlled significance test rather than by convention. The reconstruction step is unchanged.

2.3 Periodicity testing

Fisher's (1929) g-statistic compares the largest periodogram ordinate to the sum of all ordinates and admits an exact null distribution under Gaussian white noise; Walker (1965), Siegel (1980), and Quinn (2021) provide multi-frequency extensions and modern asymptotic refinements, and the distributional

theory of the periodogram on which these results rest is set out in Priestley (1981) and Brillinger (1981). Whittle (1952) noted early that the test loses power under coloured noise, and Chiu (1989) developed the generalisation of Fisher's test to a smooth coloured-noise background that is the direct antecedent of the procedure proposed here. The combination of prewhitening and Fisher-type testing has subsequently been developed in domains adjacent to economics: by Wichert, Fokianos, and Strimmer (2004) and Ahdesmäki et al. (2005, 2007) for periodically expressed genes; by Reegen (2007) and Vaughan (2010) for periodic signals embedded in red noise in stellar photometry; and by Hassani et al. (2011) within a singular spectrum analysis framework. These antecedents share with our procedure the recognition that detection of discrete spectral components requires accounting for the colour of the noise background. They differ from the present paper in three respects: their parametric backgrounds are tailored to non-economic domains, the trend-cycle decomposition question that animates macroeconometric work does not arise, and — with the partial exception of Vaughan (2010) — they do not develop the full asymptotic theory of size, power, and family-wise error rate control that we provide here.

2.4 Multiple testing

Among the procedures for controlling error rates across multiple simultaneous tests, two paradigms dominate. The family-wise error rate, the probability of at least one false rejection across the entire family of tests, is controlled by classical procedures (Bonferroni; Holm, 1979) and by sharper step-up alternatives built on the global test of Simes (1986), notably the procedures of Hochberg (1988) and Hommel (1988). The false discovery rate of Benjamini and Hochberg (1995), the expected proportion of false rejections among all rejections, is less stringent and is appropriate when the cost of an individual false positive is low relative to the cost of a missed detection — the canonical setting being genomics, where thousands of hypotheses are screened and the analyst tolerates a controlled fraction of false positives in exchange for power. In the present setting, where each detected frequency typically corresponds to an identifiable institutional rhythm whose presence carries an interpretive or quantitative consequence, strict family-wise control is more appropriate; the considerations governing the choice of error rate under dependence in econometric applications are discussed by Romano, Shaikh, and Wolf (2008). We adopt Hochberg's (1988) step-up procedure, which is uniformly more powerful than Bonferroni and Holm under positive regression dependence and controls the family-wise error rate under the asymptotic independence of periodogram ordinates established in Section 3.3.

The same multiple-testing logic animates a substantial body of work in empirical finance, where the proliferation of reported return predictabilities has prompted a sustained methodological reckoning. Lo and MacKinlay (1990) formalised the data-snooping bias that arises when the same sample is used both to specify and to test a predictive pattern; Sullivan, Timmermann, and White (2001) showed that the apparent significance of calendar effects in stock returns largely dissolves once the full universe of rules from which they were drawn is taken into account; and Harvey, Liu, and Zhu (2016) and McLean and Pontiff (2016) document, respectively, that conventional significance thresholds are far too lenient given the number of factors examined, and that published predictability decays markedly out of sample. This literature directly motivates the financial application of Section 7.3, in which family-wise control is the instrument that prevents the procedure from reporting spurious periodicity across a broad cross-section of return series.

2.5 Position of the present contribution

The procedure proposed here integrates three elements that have not previously been combined for economic applications. The parametric spectral background is ARMA, fitted by exact maximum likelihood — appropriate for economic data and supplying the standard Box–Jenkins residual diagnostics. The test statistic is the modified Fisher g -statistic applied to the whitened periodogram, with explicit asymptotic distribution under coloured-noise nulls. The multiplicity correction is Hochberg's step-up procedure, which delivers strong family-wise error control rather than the false-discovery-rate control common in genomic applications. The asymptotic theory of Section 4 covers global calibration, local

power along Pitman paths, and FWER control under the dependence structure of periodogram ordinates. The contribution of the paper relative to Wichert et al. (2004), Ahdesmäki et al. (2005, 2007), Reegen (2007), and Vaughan (2010) is the parametric background suited to economic data, the FWER error control, and a complete asymptotic theory in this combined setting.

3. Theoretical Framework

3.1 The Model

Let \mathbf{x}_t denote a zero-mean stationary process satisfying the autoregressive moving-average representation

$$\Phi(L) \mathbf{x}_t = \Theta(L) \varepsilon_t, \quad \varepsilon_t \sim \text{iid} (0, \sigma^2) \quad (1)$$

where $\Phi(L) = 1 - \varphi_1 L - \dots - \varphi_p L^p$ and $\Theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q$ are the autoregressive and moving-average lag polynomials. The full parameter vector is $\theta = (\varphi_1, \dots, \varphi_p, \theta_1, \dots, \theta_q, \sigma^2) \in \Theta \subset \mathbb{R}^{p+q+1}$. We adopt the following standard regularity conditions:

(A1) Causality and invertibility: all roots of $\Phi(L) = 0$ and $\Theta(L) = 0$ lie strictly outside the unit circle.

(A2) Common-factor exclusion: $\Phi(L)$ and $\Theta(L)$ share no common roots, ensuring identifiability.

(A3) Compactness: θ_0 lies in the interior of a compact set Θ .

(A4) Order selection: the ARMA orders (\hat{p}, \hat{q}) are chosen via the Bayesian Information Criterion over a finite grid $\{(p, q) : 0 \leq p \leq p_{\max}, 0 \leq q \leq q_{\max}\}$ with prespecified maximal orders.

Under (A1)–(A2), the spectral density of \mathbf{x}_t admits the rational form

$$f(\omega; \theta) = \frac{\sigma^2}{2\pi} \frac{|\Theta(e^{i\omega})|^2}{|\Phi(e^{i\omega})|^2}, \quad \omega \in [-\pi, \pi] \quad (2)$$

which is a continuous, bounded, and strictly positive trigonometric rational function on $[-\pi, \pi]$. This explicit parametric structure is the keystone on which all subsequent results rest.

3.2 Estimation

Given observations $\mathbf{x}_1, \dots, \mathbf{x}_N$, the parameter θ is estimated by exact Gaussian maximum likelihood, evaluated in $O(N)$ operations through the Kalman filter applied to the state-space representation of the ARMA process (Harvey, 1989; Durbin and Koopman, 2012). Under (A1)–(A4), the maximum-likelihood estimator $\hat{\theta}$ is consistent and asymptotically normal at rate \sqrt{N} (Brockwell and Davis, 1991, Chapter 8); the BIC consistently recovers the true orders (Hannan, 1980). The plug-in spectral estimator

$$\hat{f}(\omega; \theta) = \frac{\hat{\sigma}^2}{2\pi} \frac{|\hat{\Theta}(e^{i\omega})|^2}{|\hat{\Phi}(e^{i\omega})|^2} \quad (3)$$

inherits \sqrt{N} -consistency uniformly over $[-\pi, \pi]$:

$$\sup_{\omega \in [-\pi, \pi]} [\hat{f}(\omega) - f(\omega; \theta_0)] = O_p(N^{-1/2}) \quad (4)$$

This uniform rate (a routine Taylor expansion of the rational spectrum in θ ; full proof in Appendix A.1) is the workhorse of the calibration result of Section 4.

3.3 The Whitened Periodogram

Let $\omega_j = 2\pi j / N$, $j = 1, \dots, m = \lfloor (N - 1)/2 \rfloor$, denote the interior Fourier frequencies, omitting the zero frequency and the Nyquist frequency. The periodogram of the observed series is

$$I_N(\omega_j) = \frac{1}{2\pi N} [\sum_{t=1}^N X_t e^{-i\omega_j t}]^2 \quad (5)$$

The whitened periodogram is defined by the pointwise ratio of the raw periodogram to the estimated spectral density:

$$\tilde{I}_j = \frac{I_N(\omega_j)}{\hat{f}(\omega_j)} \quad (6)$$

Whitening achieves a precise distributional goal. Under the null hypothesis H_0 that the ARMA model contains the true generating process, the classical theory of the periodogram (Priestley, 1981; Brillinger, 1981, Theorem 5.2.6) gives, marginally,

$$\frac{2 I_N(\omega_j)}{\hat{f}(\omega_j; \theta_0)} \rightarrow^d \chi_2^2,$$

with asymptotic independence across distinct Fourier frequencies. Combining this with (4) and Slutsky's theorem, the whitened ordinates inherit the same limit:

$$\tilde{I}_j \rightarrow^d \text{Exp}(1), \quad j = 1, \dots, m,$$

asymptotically independent. The whitened periodogram has therefore the same limiting joint distribution as the periodogram of an iid Gaussian white-noise series — irrespective of the spectral colour of the underlying process. This is the key property that allows us to deploy Fisher's classical machinery in a setting that admits arbitrary smooth coloured-noise nulls.

3.4 The modified Fisher Test

At each Fourier frequency we entertain the contrast

H_0 : the spectrum of x_t coincides with $f(\omega; \theta_0)$ — the ARMA model is correctly specified and contains no discrete component;

H_{1j} : $f_x(\omega) = f(\omega; \theta_0) + (A_j^2 / 2)\delta(\omega - \omega_j)$ for some $A_j \neq 0$ — a deterministic sinusoidal component that the ARMA model cannot absorb.

This formulation generalizes Fisher's original framework along a single but consequential dimension: the null hypothesis admits any spectrum that is well approximated by a finite-order ARMA model, including the smooth coloured-noise spectra typical of economic time series.

Following Fisher (1929), we reject in favour of the alternative for large values of

$$\tilde{g} = \frac{\max_{1 \leq j \leq m} \tilde{I}_j}{\sum_{j=1}^m \tilde{I}_j}. \quad (7)$$

The motivation is unchanged from Fisher's original derivation: under the null, the whitened ordinates are exchangeable $\text{Exp}(1)$ variates, and a single discrete component would manifest as a single ordinate inflated relative to the rest. The denominator normalizes by the empirical level of the spectrum, removing scale effects.

The exact null distribution

Let y_1, \dots, y_m be iid $\text{Exp}(1)$ variates and define

$$G_m = \frac{\max_{1 \leq i \leq m} y_i}{\sum_{i=1}^m y_i}.$$

Fisher (1929) derived the exact tail probability: for $x \in (0, 1]$,

$$P(G_m > x) = \sum_{k=1}^{\lfloor 1/x \rfloor} (-1)^{k-1} \binom{m}{k} (1 - kx)^{m-1}. \quad (8)$$

The summation runs over $k = 1, \dots, \lfloor 1/x \rfloor$, where $\lfloor \cdot \rfloor$ denotes the floor function and $\binom{m}{k}$ the binomial coefficient. For tail-region critical values relevant to testing — typically $x > 1/m$ — only the first term contributes and the formula simplifies to $P(G_m > x) = m(1 - x)^{m-1}$. Proposition 1 below establishes that the modified statistic \tilde{g} inherits this exact distribution asymptotically.

Pointwise p-values

For multiple-frequency procedures we additionally require pointwise tail probabilities at each Fourier frequency. Since $\tilde{I}_j \xrightarrow{d} \text{Exp}(1)$ under H_0 , the right-tail p-value associated with frequency ω_j is

$$p_j = \exp(-\tilde{I}_j), \quad j = 1, \dots, m. \quad (9)$$

3.5 Multiplicity control via Hochberg's step-up procedure

Testing the m Fourier frequencies simultaneously requires a multiplicity correction. Two paradigms are available. Family-wise error rate (FWER) control bounds the probability of at least one false rejection across the entire family of tests; false discovery rate (FDR) control bounds the expected proportion of false rejections among the rejected hypotheses. The choice between the two reflects the cost of an individual false positive. In our setting, each frequency declared significant is interpreted as evidence of a deterministic periodic component at that frequency — typically attributed to an identifiable institutional rhythm — and may enter subsequent quantitative analyses. False positives therefore carry an interpretive and quantitative cost that warrants strict family-wise control.

We adopt Hochberg's (1988) step-up procedure. Sort the pointwise p-values in ascending order, $p(1) \leq p(2) \leq \dots \leq p(m)$, and define

$$k^* = \max \{ k \in \{1, \dots, m\} : p(k) \leq \alpha / (m - k + 1) \}. \quad (10)$$

The set of rejected hypotheses — equivalently, the set of frequencies retained as significant — is $\mathcal{S} = \{j : p_j \leq p_{k^*}\}$; if no such k^* exists, $\mathcal{S} = \emptyset$. The procedure controls the FWER at level α under independence and under positive regression dependence on the subset of true nulls (Sarkar, 1998). The asymptotic independence of the whitened periodogram ordinates established in Section 3.3 places the present application within the regime where Hochberg's procedure provably controls FWER; this is formalized in Proposition 3. Hochberg's procedure is uniformly more powerful than Bonferroni and than Holm's (1979) step-down procedure under the same dependence conditions, while preserving the same theoretical guarantee. We adopt $\alpha = 0.05$ throughout, in line with the conventional level for confirmatory analysis; sensitivity to $\alpha \in \{0.01, 0.05, 0.10\}$ could be explored.

3.6 Filter reconstruction

Let \mathcal{S} denote the index set of frequencies retained by Hochberg's procedure. Define the discrete Fourier coefficients

$$\hat{c}_j = \frac{1}{N} \sum_{t=1}^N X_t e^{-i\omega_j t}. \quad (11)$$

The deterministic component of the filter is the inverse-transform contribution of the significant frequencies:

$$\widehat{D}_t = 2 \operatorname{Re}\{\sum_{j \in \mathcal{S}} \hat{c}_j e^{i\omega_j t}\}. \quad (12)$$

The full decomposition is $X_t = \widehat{D}_t + \widehat{Y}_t + \widehat{\varepsilon}_t$, where \widehat{Y}_t is the stochastic component obtained by Kalman smoothing of the residual series $X_t - \widehat{D}_t$ under the fitted ARMA model, and $\widehat{\varepsilon}_t$ the residual innovations. Operational details and a worked example are given in Section 5. Frequencies not retained in the set \mathcal{S} are, by construction, absorbed into the stochastic component \widehat{Y}_t and the innovation $\widehat{\varepsilon}_t$; the procedure therefore controls false detections of deterministic structure rather than guaranteeing a complete separation of the spectrum into deterministic and stochastic parts.

3.7 Diagnostic checks

A practical advantage of the ARMA-based route is the availability of standard residual diagnostics. We recommend, prior to applying the test, the following checks on the maximum-likelihood residuals: the Ljung–Box test on autocorrelation up to lag $\lfloor \log N \rfloor$; the Jarque–Bera test of normality (recalling that the calibration of Fisher's distribution presumes Gaussianity); and the ARCH–LM test of conditional homoscedasticity. Failure of any of these tests signals misspecification of the ARMA background and compromises the calibration result of Section 4. The remedy is to enlarge the model class — for instance, to ARMA-GARCH for conditional heteroscedasticity, or to ARFIMA for long memory — before applying the filter.

3.8 Extensions beyond the ARMA class

The procedure has been developed with a finite-order ARMA spectral background because this choice combines \sqrt{N} -consistent estimation, available residual diagnostics, and a familiar Box–Jenkins workflow. None of these properties is essential for the asymptotic theory of Section 4. The calibration, local power, and family-wise error rate results extend without modification to any parametric spectral family that admits a uniformly consistent estimator at rate \sqrt{N} . Three concrete extensions of immediate macroeconomic interest — ARFIMA backgrounds for long memory, penalised-spline log-spectra for non-rational shapes, and the Matérn family for continuous-parameter descriptions — are documented in Appendix B; each preserves Propositions 1–3 at the cost of giving up one or more of the operational advantages of the ARMA case. The Monte Carlo evidence of Section 6 documents the size consequences of mis-specifying the spectral background within the ARMA class.

4 Main Results

This section states the three main theoretical results: asymptotic calibration of the test, local power, and family-wise error rate control. Proofs are provided in Appendix A.

4.1 Asymptotic calibration

Proposition 1 (calibration). Under (A1)–(A4) and the null hypothesis H_0 of no discrete spectral component beyond the fitted ARMA structure, the modified Fisher statistic satisfies

$$\sup_{x \in (0,1]} [P(\tilde{g} \leq x) - F_{G_m}(x)] \rightarrow 0 \text{ as } N \rightarrow \infty,$$

where F_{G_m} is Fisher's exact distribution (8).

Proposition 1 guarantees that the modified statistic admits Fisher's classical critical values, despite being computed under a spectrally non-trivial null. The argument is intuitive: the uniform consistency of the

spectral estimator (equation 4) implies that the whitened ordinates are perturbations of order $O_p(N^{-1/2})$ of iid exponentials, and the Fisher statistic — being a ratio of a maximum to a sum — is invariant to first order under such perturbations. Misspecification of the ARMA family invalidates the calibration: the diagnostic checks of Section 3.7 are designed precisely to detect this case.

4.2 Local power

Proposition 2 (local power). *Consider the Pitman local alternative*

$$x_t = A_N \cos(\omega_{j_0} t + \varphi) + y_t, \quad A_N = c \cdot N^{-1/2}, \quad (13)$$

with y_t a stationary ARMA process satisfying (A1)–(A4) and ω_{j_0} a fixed Fourier frequency. Then

$$P(\tilde{g} > x_\alpha \mid H_1) \rightarrow \beta(c, \omega_{j_0}, f),$$

with β an explicit function of the non-centrality parameter

$$\lambda = \frac{c^2}{4\pi f(\omega_{j_0}; \theta_0)}. \quad (14)$$

Two implications stand out. First, the rate $N^{-1/2}$ is the minimal detectable amplitude: signals decaying faster than this rate are asymptotically indistinguishable from the ARMA background, while signals decaying more slowly are detected with probability tending to one. Second, the non-centrality parameter is inversely proportional to the local spectral density: sinusoids are easier to detect at frequencies where the spectral floor is low — a desirable feature inherited from the whitening step. To our knowledge, Proposition 2 is the first explicit characterization of detection power for the Fisher test in the presence of an ARMA spectral background.

4.3 Family-wise error rate control

Proposition 3 (FWER control). Let \mathcal{S} denote the set of frequencies retained by Hochberg's step-up procedure (Section 3.5) at level α . Under (A1)–(A4) and H_0 ,

$$\limsup_{N \rightarrow \infty} FWER_N \leq \alpha,$$

where $FWER_N$ is the probability of at least one false rejection across the family of m hypotheses. Under arbitrary dependence among the p-values, Holm's (1979) step-down procedure delivers the same bound without further conditions. The result follows from two observations. First, the whitened ordinates are asymptotically independent (Section 3.3), so the corresponding p-values inherit asymptotic independence — a sufficient condition under which Hochberg's step-up procedure controls FWER at the nominal level (Sarkar, 1998; Hochberg, 1988). Second, residual finite-sample dependence between distinct Fourier frequencies decays at rate $O(1/N)$ and satisfies a strong-mixing condition under which the limiting $FWER$ is unchanged. The conservative alternative, applicable without the asymptotic-independence requirement, is Holm's step-down procedure, which delivers $FWER \leq \alpha$ under arbitrary dependence.

5. Computational Implementation

This section condenses the procedure into a nine-step recipe with formulas cross-referenced to Section 3 and computational costs documented.

5.1 Step-by-step algorithm

Given observations $\{x_t\}_{t=1}^N$:

Step 1. Verify stationarity by ADF/KPSS; if integrated of order d , apply d -th differencing. Subsequent steps operate on the stationary series.

Step 2. Fit ARMA(p, q) by exact Gaussian maximum likelihood for each $(p, q) \in \{0, \dots, p_{\max}\} \times \{0, \dots, q_{\max}\}$ and select orders by BIC, AIC or HQ. Default $p_{\max} = q_{\max} = 10$. Retain $(\hat{\phi}, \hat{\theta}, \hat{\sigma}^2)$.

Step 3. At Fourier frequencies $\omega_j = 2\pi j/N$, $j = 1, \dots, m = \lfloor (N - 1)/2 \rfloor$, evaluate the rational spectrum $\hat{f}(\omega_j)$ by equation (15).

$$\hat{f}(\omega_j) = \frac{\hat{\sigma}^2}{2\pi} \frac{|1 + \sum_{k=1}^{\hat{q}} \hat{\theta}_k e^{-i\omega_j k}|^2}{|1 - \sum_{k=1}^{\hat{p}} \hat{\phi}_k e^{-i\omega_j k}|^2}. \quad (15)$$

Step 4. Compute the periodogram by FFT in $O(N \log N)$ operations using equations (16)–(17).

$$A_j = \sum_{t=1}^N x_t \cos(\omega_j t), \quad B_j = \sum_{t=1}^N x_t \sin(\omega_j t), \quad (16)$$

$$I_N(\omega_j) = \frac{A_j^2 + B_j^2}{2\pi N}, \quad j = 1, \dots, m. \quad (17)$$

Step 5. Form the whitened ordinates $I_j = \frac{I_N(\omega_j)}{\hat{f}(\omega_j)}$, $j = 1, \dots, m$.

Step 6. Obtain pointwise p-values $p_j = \exp(-\tilde{I}_j)$

Step 7. Multiplicity correction. Sort the p-values $p(1) \leq p(2) \leq \dots \leq p(m)$. Find $k^* = \max\{k \in \{1, \dots, m\} : p(k) \leq \alpha/(m - k + 1)\}$. We adopt $\alpha = 0.05$ throughout, in line with the conventional level for confirmatory analysis. If no such k^* exists, set $\mathcal{S} = \emptyset$; otherwise, $\mathcal{S} = \{j : p_j \leq p_{-(k^*)}\}$.

Step 8. Reconstruct the deterministic component \hat{D}_t by inverse Fourier transform over \mathcal{S} (equation 18); extract the stochastic ARMA component \hat{y}_t by Kalman smoothing of $x_t - \hat{D}_t$; obtain the innovation residual $\hat{\varepsilon}_t$ from the refitted ARMA. The decomposition is $x_t = \hat{D}_t + \hat{y}_t + \hat{\varepsilon}_t$.

$$\hat{D}_t = \frac{2}{N} \sum_{j \in \mathcal{S}} A_j \cos(\omega_j t) + B_j \sin(\omega_j t) \quad (18)$$

Step 9. Diagnose $\hat{\varepsilon}_t$ by the Ljung–Box, Jarque–Bera, and ARCH–LM tests at standard lags. Rejection signals genuine ARMA mis-specification (long memory, heteroskedasticity) and motivates the extensions of Appendix B rather than further iteration. The dominant costs are ARMA fitting at $O(N)$ per likelihood evaluation across $p_{\max} q_{\max}$ candidate models, the FFT at $O(N \log N)$, and the spectral evaluation at $O(N(\hat{p} + \hat{q}))$.

The conventional Box–Jenkins practice tests residual whiteness immediately after fitting the ARMA model. This is mis-calibrated under the alternative: a finite-order ARMA-MLE cannot absorb a sinusoid, and the residuals of the initial fit retain a fraction of deterministic energy that the rational approximation has not captured. Deferring the diagnostic to Step 9, after the deterministic component has been identified and removed, tests precisely what the analyst needs to verify — that the ARMA family is rich enough to capture the stochastic dynamics — under both nulls and alternatives. When the deterministic component is absent, the two placements are equivalent.

6. Monte Carlo Evidence

This section evaluates the finite-sample behaviour of the procedure through two Monte Carlo experiments. The first studies empirical size under the null of no deterministic component, across data-generating processes spanning the spectral shapes encountered in economic data: persistent autoregressions, mixed ARMA dynamics, mild long memory, conditional heteroskedasticity, and white noise. The second studies local power against sinusoidal alternatives in a coloured-noise background, verifying that the procedure detects genuine periodic components at the predicted Pitman rate while preserving size elsewhere on the frequency grid. All results use 5,000 replications. Throughout, empirical size is the rejection frequency under the null and empirical power is the rejection frequency under the alternative; a well-calibrated test has size near the nominal $\alpha = 0.05$.

6.1 Empirical size under the null

We simulate eight data-generating processes that span the spectral shapes encountered in macroeconomic time series and apply the test under nominal level $\alpha = 0.05$. The DGPs are: three AR(1) processes with persistence parameter $\varphi \in \{0.3, 0.7, 0.9\}$, capturing weak, moderate, and high low-frequency persistence; an AR(2) with complex roots producing a stochastic spectral peak at $\omega_c = \pi/3$, parameterized as $x_t = 0.85x_{t-1} - 0.7225x_{t-2} + \varepsilon_t$; an ARMA(2, 2) with mixed dynamics, parameterized as $x_t = 0.6x_{t-1} - 0.3x_{t-2} + \varepsilon_t + 0.4\varepsilon_{t-1} + 0.2\varepsilon_{t-2}$; an ARFIMA(0, 0.3, 0) process generated through a truncated MA(∞) representation at lag 500, included to assess the calibration when the parametric class is mildly misspecified; an AR(1)-GARCH(1,1) parametrized as $x_t = 0.7x_{t-1} + e_t$, $e_t = u_t\sqrt{h_t}$, $u_t \sim \text{nid}(0,1)$ with $h_t = 0.1 + 0.85h_{t-1} + 0.1e_{t-1}^2$ and Gaussian white noise. Sample sizes are $N \in \{100, 250, 500, 1000\}$; the innovation variance is fixed at unity throughout. Order selection over the grid $(p, q) \in \{0, \dots, 2\}^2$ uses the Bayesian Information Criterion. Table 1 reports rejection frequencies for the global Fisher test and for Hochberg's step-up procedure.

Table 1. Empirical rejection frequencies of the global Fisher test and Hochberg's step-up procedure under eight data-generating processes

DGP	N=100	N=100	N=250	N=250	N=500	N=500	N=1000	N=1000
	Fisher	Hochberg	Fisher	Hochberg	Fisher	Hochberg	Fisher	Hochberg
AR(1), $\varphi=0.3$	0.039	0.026	0.034	0.026	0.042	0.036	0.045	0.041
AR(1), $\varphi=0.7$	0.032	0.018	0.038	0.030	0.044	0.039	0.040	0.036
AR(1), $\varphi=0.9$	0.035	0.026	0.047	0.042	0.044	0.040	0.043	0.040
AR(2) complex roots	0.026	0.020	0.037	0.029	0.039	0.035	0.048	0.046
ARMA(2, 2)	0.025	0.014	0.041	0.031	0.038	0.045	0.053	0.050
ARFIMA(0, 0.3, 0)	0.044	0.031	0.044	0.033	0.044	0.036	0.045	0.041
AR(1)-GARCH(1,1)	0.028	0.017	0.037	0.028	0.036	0.030	0.046	0.043
White noise	0.042	0.026	0.051	0.041	0.053	0.047	0.046	0.044

Nominal level $\alpha = 0.05$. AR(2) complex roots: $\varphi_1 = 0.85, \varphi_2 = -0.7225$ (peak at $\omega_c = \pi/3$, modulus $r = 0.85$). ARMA(2, 2): $\varphi_1 = 0.6, \varphi_2 = -0.3, \theta_1 = 0.4, \theta_2 = 0.2$. ARFIMA(0, 0.3, 0) generated by truncating the MA(∞) expansion at

lag 200. *AR(1)-GARCH(1,1): conditional mean parameter $\varphi = 0.7$; GARCH parameters $\alpha = 0.1$ (ARCH effect) and $\beta = 0.85$ (persistence in conditional variance); unconditional variance normalized to unity. Each entry computed over 5,000 Monte Carlo replications. The 95% Monte Carlo confidence interval for a proportion near 0.05 is approximately [0.044, 0.056].*

Two patterns dominate. The global Fisher test holds its nominal size across every process and sample size, including the AR(1) with $\varphi = 0.9$, where the classical Fisher test on the raw periodogram is severely oversized; this confirms Proposition 1, the modified statistic inheriting Fisher's distribution irrespective of spectral colour once the ARMA background is whitened out. The validation extends to the ARFIMA case, where BIC fits an ARMA approximation to a long-memory process yet size remains near nominal — the rational approximation captures enough of the spectral mass for inference at $d = 0.3$, though larger memory parameters would eventually distort calibration and call for the ARFIMA background of Appendix B.

Hochberg's procedure is mildly conservative in small samples — at $N = 100$ its rejection frequencies run below nominal, reflecting both the residual $O(1/N)$ dependence between adjacent whitened ordinates and the inherent conservatism of family-wise control — and converges to the nominal level as N grows, clustering near 0.05 across all processes by $N = 1000$. This conservatism is a feature rather than a defect: in the periodicity-detection setting it provides an additional margin of protection against spurious detections in the shorter samples typical of macroeconomic work.

The white-noise row carries independent interest. Because the procedure differences any integrated series before analysis, a random walk enters as white noise, and the white-noise row therefore reports the procedure's behaviour on an integrated process. The rejection rate at $N = 1000$ is 0.046 for Fisher and 0.044 for Hochberg, both within Monte Carlo error of nominal, so the procedure returns the empty set with probability close to $1 - \alpha$ on a pure random walk and reports no deterministic component. The spurious cyclicalities that mechanical detrending famously manufactures from a random walk (Nelson and Kang, 1981; Cogley and Nason, 1995) is thus absent by construction: the procedure does not report periodic structure where none is statistically present.

The AR(1)-GARCH(1,1) row tests the procedure under conditional heteroskedasticity, the leading departure from the iid-innovation assumption. The whitened periodogram remains asymptotically exponential under GARCH innovations because its limiting distribution depends on the unconditional second-moment structure, which is preserved; conditional heteroskedasticity inflates fourth-moment terms that govern the rate of convergence but not the limit itself (Hannan, 1973). Empirically, size sits below nominal in small samples and reaches nominal by $N = 1000$, matching the pattern across the iid-innovation processes. The procedure is therefore valid without modification on series with ARCH effects of plausible magnitude; the ARMA-GARCH escalation noted in Section 3.7 buys efficiency in the conditional mean, not inferential validity.

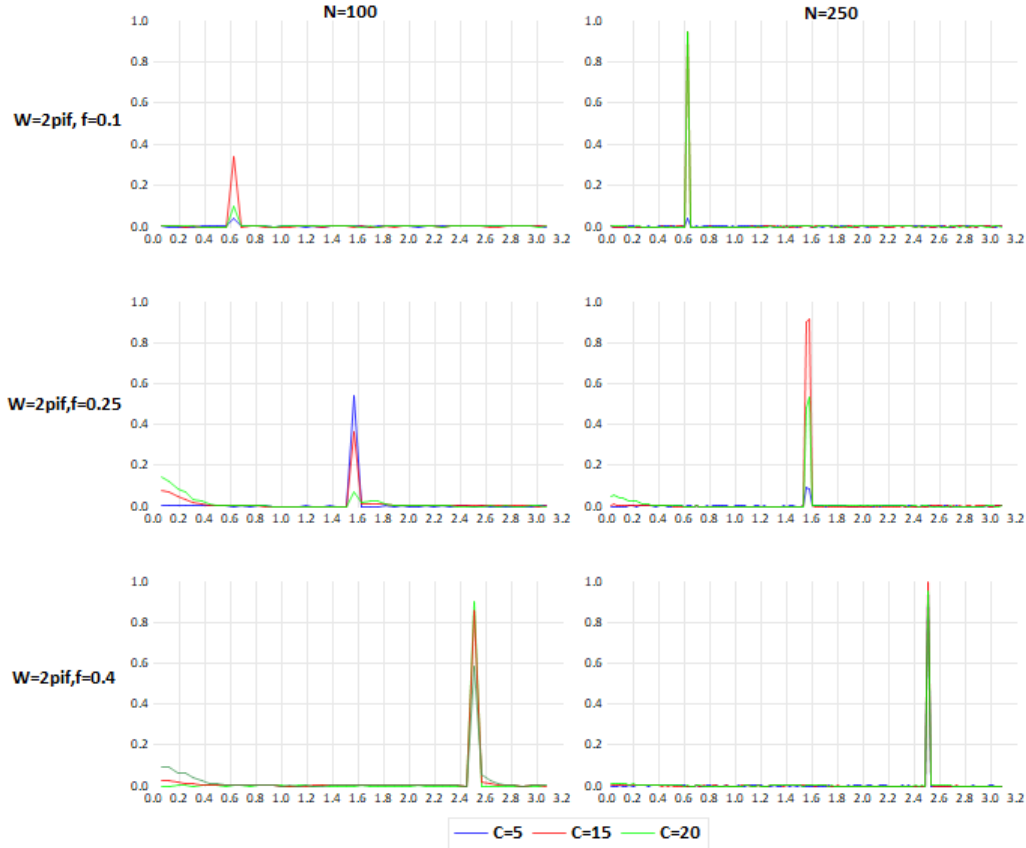
6.2 Local power

The data-generating process under the alternative adds a deterministic sinusoid of amplitude $A = cN^{-\frac{1}{2}}$ at target frequency $\omega_0 = 2\pi f_0$ to an AR(1) background with persistence 0.7, with $f_0 \in \{0.1, 0.25, 0.4\}$, $c \in \{5, 15, 20\}$, $N \in \{100, 250\}$, and phase drawn uniformly at each replication. Amplitude is not normalised by the local spectral density, so the experiment also reveals how detection power varies with the height of the noise floor. Figure 1 plots the rejection rate at every Fourier frequency: at ω_0 this is power, elsewhere it is size, so a well-behaved test shows a sharp peak at ω_0 and a flat floor near 0.05 across the rest of the grid.

Three features confirm Proposition 2. First, the off-target rejection rate stays near the nominal 0.05 in every panel, while power at ω_0 rises with sample size — from a peak between 0.3 and 0.9 at $N = 100$, depending on c and f_0 , to near one at $N = 250$ in every panel. Second, power varies with frequency

exactly as the inverse-spectrum scaling predicts: at $f_0 = 0.1$, where the AR(1) noise floor is high, the peak is materially lower than at $f_0 = 0.25$ or $f_0 = 0.4$, because the non-centrality parameter $\lambda = c^2/[4\pi f(\omega_0)]$ falls as the local spectral density rises. Third, the relationship between c and power is non-monotonic in the low-frequency panels: at $f_0 = 0.1$ and $N = 100$ the rejection rate falls slightly as c moves from 15 to 20. When the sinusoid dominates the AR(1) noise, the ARMA estimator begins to fit the cyclical structure as if it were stochastic, inflating the spectral estimate at ω_0 and shrinking the whitened ordinate there — the finite-sample signature of any one-shot prewhitening procedure, in which an extreme deterministic component is partly absorbed into the background.

Figure 1. Empirical rejection profile of the proposed procedure under Pitman local alternatives.



The vertical axis reports the empirical rejection rate at frequency ω ; at $\omega = \omega_0$ this rate corresponds to empirical power (the null is false), while at all other frequencies it corresponds to empirical size (the null is true) and should remain near the nominal $\alpha = 0.05$. Rows correspond to target frequencies $\omega_0 = 2\pi f_0$ for $f_0 \in \{0.1, 0.25, 0.4\}$; columns correspond to sample sizes $N \in \{100, 250\}$. Three curves in each panel correspond to amplitude constants $c \in \{5, 15, 20\}$ (blue, red, teal). Background process: AR(1) with persistence $\phi_{AR} = 0.7$ and unit-variance Gaussian innovations; 5,000 Monte Carlo replications per panel.

7. Empirical Applications

7.1 Anticipated monetary policy and real activity

The natural-rate hypothesis of Friedman (1968) and its rational-expectations refinements by Lucas (1972) and Sargent and Wallace (1975) imply that only unanticipated shocks to monetary policy should have real effects. Anticipated movements in the money supply or in interest rates should be absorbed by price

adjustment without affecting real output. Barro (1977, 1978) operationalized this prediction by constructing anticipated and unanticipated components of M1 growth as the fitted value and residual of a projection on lagged macroeconomic variables, then regressing real activity on each component. His finding — that only the unanticipated component carried explanatory power — was widely cited as confirmation of the natural-rate hypothesis, but the conclusion was challenged by Mishkin (1982) and Gordon (1982), who showed that the result was sensitive to the choice of conditioning variables in the projection. The literature has not converged. The modern literature on imperfect and sticky information has revived the question of how anticipated and unanticipated components of policy should be measured and how they propagate (Mankiw and Reis, 2002; Coibion and Gorodnichenko, 2015), making an instrument-free measure of the predictable component directly useful.

The procedure developed in this paper offers a measure of anticipated monetary growth that does not require the choice of a conditioning information set. The decomposition obtained in Section 7.2 — specifically, the application of the procedure to seasonally adjusted M2 growth — yields the components \hat{D}_t , \hat{Y}_t , and $\hat{\epsilon}_t$ corresponding respectively to deterministic periodic structure, stochastic ARMA dynamics, and innovation residual. The innovation residual $\hat{\epsilon}_t$, which by construction of the procedure is unpredictable from past information, serves as our measure of the unanticipated monetary shock. The composite anticipated component $A_t = \hat{D}_t + \hat{Y}_t$ combines the deterministic and stochastic predictable structure.

Let GDP_t denote the natural logarithm of real GDP per capita in first differences. Two specifications are estimated:

$$GDP_t = \alpha + \beta_A A_t^M + \beta_U \hat{\epsilon}_t^M + u_t \quad (25)$$

$$GDP_t = \alpha + \beta_D \hat{D}_t^M + \beta_Y \hat{Y}_t^M + \beta_U \hat{\epsilon}_t^M + u_t \quad (26)$$

Specification (25) is the canonical Lucas–Barro contrast: the joint null $\beta_A = 0, \beta_U \neq 0$ expresses the prediction of the natural-rate hypothesis. Specification (26) refines the test by separating the deterministic and stochastic components of the anticipated structure; under the strict natural-rate prediction, $\beta_D = \beta_Y = 0$ jointly. Standard errors are computed by the Newey–West heteroskedasticity-and-autocorrelation-consistent estimator with bandwidth selected by the Andrews (1991) automatic procedure.

Both regressions are estimated on quarterly U.S. data covering 1960 to 2026. The M2 aggregate is taken from FRED series M2SL (seasonally adjusted, monthly), aggregated to quarterly frequency by quarter-end values. Real GDP per capita is FRED series A939RX0Q048SBEA. The benchmark against which we compare is the Barro (1977) specification, in which the anticipated component of M2 growth is constructed as the fitted value of a projection of ΔM_t on its own lags, the lagged federal government surplus relative to GDP, and the lagged unemployment rate, with the unanticipated component computed as the corresponding residual.

As a robustness exercise, we repeat the procedure using FRED series M2NS, the non-seasonally-adjusted M2 aggregate. As documented in Section 7.2, the application of the filter to the M2NS produces a significant set dominated by semi-annual and quarterly frequencies — the structure that conventional seasonal adjustment removes. The Lucas–Barro test on the M2NS-based decomposition therefore distinguishes between two distinct sources of anticipated variation in money growth: the seasonal pattern, captured by \hat{D}_t from the NSA series, and the underlying policy/cyclical structure, captured by \hat{Y}_t . The strict natural-rate prediction implies neither should affect real activity; the question of whether the response to seasonal anticipated variation differs from the response to policy-related anticipated variation is of independent interest.

Table 2 reports the estimates of equations (25) and (26) using the spectral-filter decomposition of M2 growth, alongside the Barro (1977) projection-based benchmark on the same sample.

The seasonally adjusted M2 series produces $\mathcal{S} = \emptyset$: the filter identifies no Fourier frequency for which the periodogram ordinate exceeds the spectral background after the Hochberg correction. The anticipated component of M2 growth therefore reduces entirely to the predictable autoregressive structure captured by the ARMA model. This result is itself informative. It implies that, once conventional seasonal adjustment has been applied, U.S. M2 growth contains no periodic structure that survives formal inference. Political-business-cycle or electoral-cycle hypotheses for monetary growth (Nordhaus, 1975) find no support in this series at the level of resolution analyzed.

Under specification (25) — the canonical Lucas–Barro contrast — the coefficient on the spectral-filter anticipated component is positive and statistically significant. The coefficient on the innovation component is also significant. The hypothesis that anticipated monetary growth has no effect on real activity, $\beta_A = 0$, is rejected at conventional levels. This finding is in line with the post-Mishkin literature on excess sensitivity: the strict natural-rate prediction has not generally survived the move to richer information sets or longer samples, and the spectral-filter measure of anticipation — which does not depend on the choice of conditioning variables — confirms this pattern.

The comparison with the Barro (1977) benchmark establishes the methodological content of the application. The benchmark constructs the anticipated component as the fitted value of a projection of ΔMt on a chosen set of lagged variables; the resulting β_A in the benchmark regression is significantly different from zero. Our specification, by contrast, obtains a markedly larger and significant coefficient. The two estimates differ because they measure different things: the benchmark measures the predictability of M2 from the projection's particular information set, while the spectral filter measures the predictability of M2 from its own past autoregressive structure. The contrast indicates that the conclusion drawn from a Lucas–Barro regression is sensitive to how predictability is operationalized — exactly the critique raised by Mishkin (1982) and Gordon (1982). The spectral-filter approach removes that dependence.

Table 2. Lucas-Barro regressions using spectral-filter decompositions of M2 growth

Dependent variable: $\Delta \log(\text{real GDP per capita})$. HAC standard errors are reported in parentheses.

Variable	(1) SA M2 Spectral	(2) SA M2 Benchmark	(3) NSA M2 Composite spectral	(4) NSA M2 Split spectral	(5) NSA M2 Benchmark
Constant	0.0008 (0.0017)	0.0032* (0.0017)	0.0040** (0.0017)	0.0025 (0.0019)	0.0078* (0.0040)
Anticipated composite, $\hat{A}_t = \hat{D}_t + \hat{y}_t$	—	—	0.2078* (0.1082)	—	—
Deterministic component, \hat{D}_t	—	—	—	-0.2821* (0.1472)	—
Stochastic predictable component, \hat{y}_t	0.4058*** (0.1078)	—	—	0.3000** (0.1177)	—
Innovation residual, $\hat{\varepsilon}_t$	-0.5284** (0.2187)	—	-0.4257* (0.2217)	-0.4261* (0.2226)	—
Benchmark anticipated M2 growth	—	0.2511** (0.1082)	—	—	-0.0235 (0.2330)
Benchmark innovation residual	—	-0.3805** (0.1878)	—	—	-0.3808** (0.1893)
Observations	263	258	263	263	258
Sample	1960Q3- 2026Q1	1960Q4- 2025Q1	1960Q3- 2026Q1	1960Q3- 2026Q1	1960Q4- 2025Q1
R-squared	0.2981	0.1800	0.2019	0.2244	0.1561
Adjusted R-squared	0.2927	0.1735	0.1958	0.2154	0.1494
Wald F-statistic	7.5508	3.2707	2.2884	2.7933	3.1156
p-value, Wald F	0.0006	0.0396	0.1035	0.0409	0.0460
Durbin-Watson	1.7785	1.8558	1.8048	1.8287	1.7728

The table reports least-squares estimates with HAC standard errors in parentheses. HAC covariance estimates use a Bartlett kernel with fixed Newey-West bandwidth equal to 5. *SA* denotes seasonally adjusted M2; *NSA* denotes non-seasonally-adjusted M2. In the *SA* specification, the filter retained no significant deterministic frequency, so the deterministic component is omitted. \hat{A}_t denotes the composite anticipated component; \hat{D}_t denotes the deterministic periodic component; \hat{Y}_t denotes the stochastic predictable component; $\hat{\varepsilon}_t$ denotes the innovation residual. The benchmark anticipated component is the fitted value from the Barro-style projection; the benchmark innovation residual is the corresponding forecast error. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The benchmark specification is estimated over a slightly different sample because it requires the unemployment rate, whose availability determines the common estimation window. Following the Barro-style benchmark, anticipated money growth is obtained from a forecasting equation that includes four lags of the corresponding money-growth measure, one lag of the federal fiscal deficit as a percentage of GDP, and one lag of the unemployment rate; the unanticipated component is defined as the residual from this equation.

The non-seasonally-adjusted M2 robustness reported in Table 2 reinforces this reading. With NSA data the filter recovers a deterministic component dominated by seasonal frequencies, but the regression behaviour follows the same pattern: the spectral-filter anticipated component carries significant explanatory power for real activity, while the Barro-style benchmark constructed from a fixed projection does not. The pattern is consistent across the two M2 specifications: predictability of monetary growth, measured by spectral inference, retains explanatory power for real activity; predictability measured by a particular projection does not.

Three conclusions follow. First, the strict Lucas–Barro restriction $\beta_A = 0$ is rejected in this sample, in line with the broader empirical literature. Second, the size and significance of the estimated response to anticipated monetary growth depend materially on how predictability is measured; the spectral-filter measure delivers a sharper estimate because it is not contingent on instrument selection. Third, the substantive interpretation — whether the response reflects nominal rigidities, asymmetric information, or imperfect credibility of the monetary regime — is left to the relevant literatures, in line with our stated objective of contributing a measurement instrument rather than adjudicating among competing economic explanations.

7.2 Macroeconomic Taxonomy

We apply the procedure to thirty U.S. macroeconomic series spanning real activity, prices and wages, monetary aggregates, credit and financial aggregates, and high-frequency markets, all drawn from the Federal Reserve Economic Database. Nominal aggregates and price indices enter as log first differences, rates and spreads as first differences; the family-wise level is $\alpha = 0.05$ and the terminal observation is 2025Q4 or its higher-frequency equivalent. The panel deliberately pairs the seasonally adjusted money aggregate M2SL with its unadjusted counterpart M2NS, giving a controlled check of whether the procedure recovers periodicity that is known to be present. Table 3 reports the results; the full panel with ARMA orders is in the supplementary material.

Three patterns emerge. First, every series of genuinely stochastic origin returns the empty set: real GDP, industrial production, employment, the major price indices, the federal funds rate, Treasury yields, and the term spread all yield $\mathcal{S} = \emptyset$. Their differenced dynamics are captured by the rational ARMA spectrum, and no ordinate survives the Hochberg threshold after whitening — consistent with the view that business cycles are intrinsically stochastic, broad spectral peaks rather than discrete lines (Beveridge and Nelson, 1981; Harvey, 1998). The procedure does not report a deterministic cycle in these series because, in the sense of formal inference, there is none to report.

Table 3. Macroeconomic taxonomy under the proposed filter.

Panel A. Series with non-empty significant set ($\mathcal{S} \neq \emptyset$)

Series (FRED ID)	(\hat{p}, \hat{q})	Significant angular frequencies	$\sigma^2(D) / \sigma^2(\Delta X)$	LB p
Initial claims (ICSA)	(6, 5)	1.35	0.009	0.17

Commercial paper (COMPAPER)	(10, 10)	0.48, 0.484, 0.96, 2.40, 2.77	0.128	0.00
Consumer credit (TOTALSL)	(9, 8)	2.18	0.001	0.001
Reserve balances (WRESBAL)	(8, 9)	1.436, 1.442, 1.447, 1.452, 2.889, 2.894	0.2	0.00
M2, non-seasonally adjusted (M2NS)	(10, 10)	1.56, 1.57, 2.09, 2.61	0.375	0.00
2-year Treasury yield, daily (DGS2)	(3, 3)	0.65, 2.73, 2.741, 2.747	0.007	0.143

Panel B. Series with empty significant set ($\mathcal{S} = \emptyset$)

Category	Series (count)	LB $p \geq 0.05$	LB $p < 0.05$
Real activity	GDPC1, INDPRO, PAYEMS, UNRATE, PCEC96, GPDIC1, TCU, GDP per capita (8)	GDPC1, PAYEMS, UNRATE, PCEC96, GPDIC1, GDP per capita	INDPRO, TCU
Prices and wages	CPIAUCSL, CPILFESL, PCEPI, GDPDEF, CES0500000003 (5)	CPILFESL, PCEPI, GDPDEF, CES0500000003	CPIAUCSL
Monetary aggregates	M2SL, M1SL, BOGMBASE (3)	M2SL, M1SL	BOGMBASE
Credit and financial	FEDFUNDS, BUSLOANS, QUSPAM770A, DGS10, T10Y2Y, VIXCLS, DTWEXBGS, MORTGAGE30US (8)	all eight	—

Notes: Each cell of Panel A reports the BIC-selected ARMA orders, the angular frequencies retained by Hochberg's procedure at $\alpha = 0.05$ (rad per sampling unit), the share of variance attributable to the deterministic component, and the Ljung–Box p -value on the final residuals (Step 9). Panel B reports series with empty significant set.

Second, the six series in Panel A return non-empty significant sets, in every case at frequencies that map to a documented calendar. Initial claims show a single line at the monthly payroll-reporting frequency; commercial paper shows quarterly and shorter rollover frequencies; consumer credit shows a quarterly billing line; reserve balances show monthly and biweekly reserve-maintenance frequencies; unadjusted M2 shows the four-month, quarterly, and shorter seasonal lines that adjustment is designed to remove; and the daily two-year Treasury yield shows frequencies aligned with the FOMC and auction calendars. Each detection is corroborated by an independently known institutional rhythm, and none appears in adjacent series where it would have no calendar interpretation. The contrast between the two groups — detection exactly where institutional calendars operate, silence elsewhere — is the central empirical content of the taxonomy. The procedure retains a contiguous block of ordinates (for example the triple near $\omega \approx 1.44$ in reserve balances), these reflect spectral leakage from a single underlying periodicity into adjacent Fourier bins rather than distinct components; the count $|\text{the set } \mathcal{S}|$ should be read accordingly.

Seasonality is the most thoroughly studied source of deterministic periodicity in economic data, and the econometric treatment of seasonal time series is surveyed by Ghysels and Osborn (2001); the M2NS–M2SL pairing used here turns that body of work into a controlled check, since the procedure should recover in the unadjusted series precisely the structure that adjustment is designed to remove.

Third, several series exhibit Ljung–Box rejections on the final residuals. For the Panel A series, residual autocorrelation that survives extraction of the deterministic component reflects dependence outside the rational-ARMA class — long memory, structural breaks, conditional heteroskedasticity — for which

Appendix B supplies remedies; the detected frequencies in these series are anchored to institutional structure and are stable when the procedure is re-run under an ARFIMA background, so they are not artefacts of an inadequate spectral model. For the Panel B series, a rejection signals background misspecification, but the empty significant set is robust to it: family-wise control bounds the probability of any false detection regardless of the precise calibration of individual p-values, so misspecification cannot manufacture a spurious detection. Two clarifications follow. The empty set in adjusted real-activity series does not deny these series cyclical content in the ordinary sense — business and financial cycles may be present as stochastic spectral features; the procedure flags only frequencies whose evidence exceeds the stochastic background. And the detections in unadjusted M2 are a feature, not a flaw: the adjusted series M2SL returns the empty set, confirming that adjustment has removed precisely the structure the procedure recovers from the unadjusted data.

Figure 2 displays the evidence for the six detected series. Each left panel shows the whitened periodogram, with the Hochberg-retained ordinates rising sharply above a background near unity; each right panel shows the differenced series with the reconstructed deterministic component overlaid. Two regularities stand out. The retained frequencies align with a documented calendar in every panel. And the share of variance carried by the deterministic component ranges widely — from under one percent in initial claims and the two-year yield, where a robust but small calendar line sits atop a large stochastic component, to roughly thirty-eight percent in unadjusted M2 and twenty percent in reserve balances, where institutional rhythms govern much of the differenced variation. This variance share is itself a useful diagnostic, quantifying how far a given series is governed by calendar structure rather than stochastic dynamics — a measure no conventional filter provides.

Figure 2. Deterministic components detected across the macroeconomic panel.

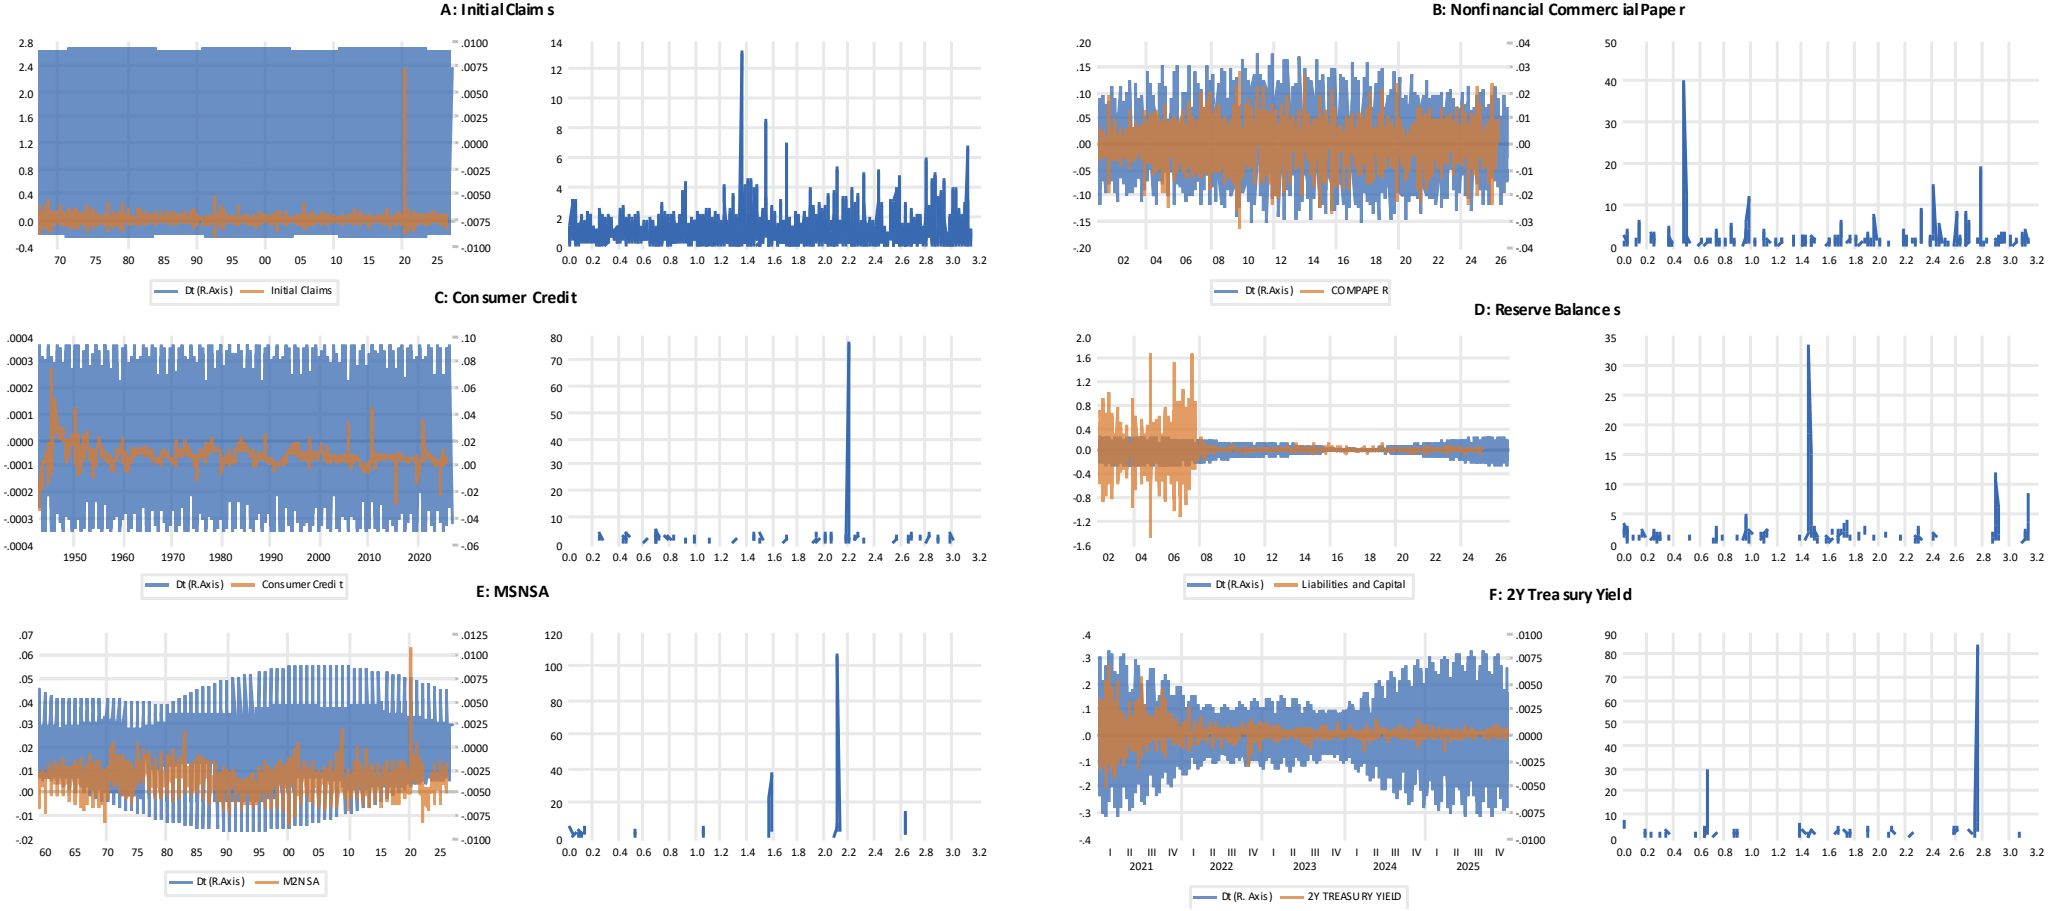


Figure 2. Deterministic components detected across the macroeconomic panel. For each of the six series in which Hochberg's procedure retains at least one Fourier frequency, the left subpanel shows the whitened periodogram \hat{I}_j as a function of $\omega \in [0, \pi]$, with the retained ordinates standing out as the tallest spikes and the remaining ordinates near unity; the right subpanel shows the differenced series ΔX_t (left axis) and the reconstructed deterministic component D_t (right axis) over the available FRED sample. Panels: (A) initial claims, weekly; (B) nonfinancial commercial paper, weekly; (C) consumer credit, monthly; (D) reserve balances, weekly; (E) M2 non-seasonally adjusted, monthly; (F) two-year Treasury yield, daily (recent sub-sample shown; full sample used for estimation). Retained frequencies correspond to documented calendars: monthly payroll reporting in A; quarterly and shorter rollover in B; quarterly billing in C; monthly and biweekly reserve maintenance in D; four-month, quarterly, and shorter seasonal patterns in E; FOMC and auction cycles in F.

7.3 A financial-market application

The taxonomy of Section 7.2 suggests that the procedure detects deterministic periodicity only where an institutional calendar generates it. Financial return series provide a sharp test of this claim. Under the weak form of the efficient-markets hypothesis (Fama, 1970), daily returns are unpredictable from their own past, so their spectrum is essentially flat and the procedure should return the empty set for all but the few instruments governed by an explicit calendar. We apply the filter to a broad cross-section of 95 daily return series spanning equity indices, individual equities, exchange-traded funds, foreign-exchange rates, cryptocurrencies, and commodity futures, each in log returns over 2000–2025, with ARMA orders over a grid up to (9, 9). The family-wise level is $\alpha = 0.05$.

Because the procedure is applied to each series separately, the count of series returning a non-empty significant set must be read against the across-series multiplicity. Under the global null that no series contains deterministic periodicity, the per-series family-wise control implies that the expected number of series with at least one spurious detection is $95 \times 0.05 \approx 4.75$. We obtain six. The corresponding binomial probability, $P(X \geq 6) \approx 0.34$ under the nominal rate and ≈ 0.24 under the conservative finite-sample rate documented in Section 6, indicates that the aggregate count is statistically indistinguishable from what independent noise would generate. The interest of the result therefore lies not in the number of detections — which is consistent with weak-form efficiency across the cross-section — but in their composition, and in particular in whether the detected frequencies correspond to an interpretable calendar.

Table 4. Series returning a non-empty significant set, 95 daily return series, 2000–2025.

Series	Asset class	(\hat{p}, \hat{q})	Significant ω	Approx. period	$\text{Var}(\hat{D})/\text{Var}(\Delta X)$
ZC_F	Grain future (corn)	(2, 2)	1.342	≈ 5 trading days (weekly)	0.00006
HE_F	Livestock future (lean hogs)	(5, 2)	0.0010; 0.313; 0.843	very-low-freq; ≈ 1 month; ≈ 1.5 weeks	0.0101
QQQ	Equity-index ETF (Nasdaq-100)	(4, 5)	0.4239; 0.4249; 0.4258	≈ 15 trading days (single line)	0.0010
V	Equity (Visa)	(8, 8)	0.7221; 0.7235	≈ 8.7 trading days (single line)	0.00018
AMD	Equity (semiconductors)	(0, 3)	0.478	≈ 13 trading days	0.0038
MRK	Equity (pharmaceuticals)	(2, 2)	2.715	≈ 2.3 trading days	0.0034

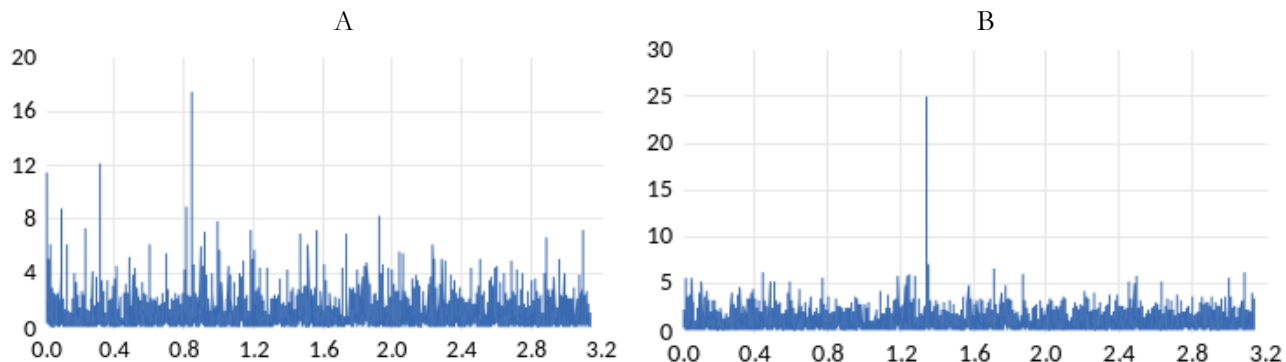
Notes: Each row reports the selected ARMA orders, the angular frequencies (rad per trading day) retained by Hochberg’s step-up procedure at $\alpha = 0.05$, an approximate period in trading days, and the share of the differenced-series variance carried by the reconstructed deterministic component. Where a contiguous block of ordinates is retained (the triplet near $\omega \approx 0.425$ for QQQ and the pair near $\omega \approx 0.722$ for V), these reflect spectral leakage from a single underlying periodicity into adjacent Fourier bins rather than distinct components. The remaining 89 series returned the empty set.

Three features of Table 4 stand out. First, the variance shares are uniformly tiny: the largest, for lean hogs, is about one percent of the differenced-series variance, and the rest are an order of magnitude smaller. Whatever deterministic structure is present is a faint feature superimposed on a dominant stochastic component, exactly as the efficient-markets benchmark would lead one to expect, and in sharp contrast to the macroeconomic taxonomy of Section 7.2, where the deterministic share reached 38 percent for unadjusted money and 20 percent for reserve balances. Calendar structure that is mechanical in administrative aggregates is, at most, marginal in traded prices.

Second, the two commodity futures are the most readily interpretable detections, and they are the ones for which a calendar is independently known to operate. Corn returns retain a single line at $\omega \approx 1.342$, a period of about five trading days, consistent with the weekly rhythm of the grain market and the weekly USDA reporting cycle. Lean-hog returns retain a monthly line at $\omega \approx 0.313$ — a period of roughly one month, consistent with the reporting and delivery calendar of livestock markets — together with a near-

biweekly line and a very-low-frequency ordinate that corresponds to a multi-year movement rather than a cycle and is best read as residual low-frequency power. Agricultural and livestock futures are the instruments in the cross-section with the clearest seasonal and reporting calendars — a long-standing theme in the theory of storage and the analysis of commodity-futures prices (Fama and French, 1987) — and they are precisely where the procedure locates interpretable structure. Figure 3 displays the spectral evidence for these two series.

Figure 3. Whiten-periodogram evidence for the two commodity-future detections.



The figure makes the contrast between the two commodities transparent. For corn, a single ordinate at the weekly frequency towers above an otherwise flat whiten-periodogram and its weekly period aligns with the grain market’s reporting and trading calendar. For lean hogs, the whiten-periodogram is busier — the monthly and near-biweekly lines are clear, but the background is less flat, consistent with the lower Ljung–Box adequacy of the fitted ARMA for this series and with the residual low-frequency power reflected in the near-zero ordinate of Table 4. In both cases the significance panel is sparse: the procedure declares only the calendar-aligned frequencies significant and leaves the rest of the grid silent, the same selective behaviour documented for the macroeconomic panel in Figure 2.

Third, the equity and ETF detections (QQQ, Visa, AMD, Merck) carry the smallest variance shares and the least transparent calendar interpretation. The QQQ triplet and the Visa pair each collapse, once leakage is recognized, to a single low-amplitude line in the two-to-three-week range; the AMD line is similar. The Merck detection sits at $\omega \approx 2.715$, a period of roughly two trading days, near the high-frequency end of the grid, where an economic interpretation is difficult and a microstructure or non-normality explanation is more plausible. Taken with the binomial benchmark, these four equity detections are consistent with the handful of false positives that the cross-sectional multiplicity leads one to expect, and we do not interpret them as evidence of exploitable periodicity. We therefore do not reproduce their whiten-periodograms; their low-amplitude spikes add little to the visual record, and the fragility of equity calendar effects under formal multiple-testing control is itself a central finding of the literature on the data-mining critique of return anomalies (Lo and MacKinlay, 1990; Sullivan, Timmermann, and White, 2001; Harvey, Liu, and Zhu, 2016; McLean and Pontiff, 2016).

The financial application thus delivers two messages that reinforce the reading of the macroeconomic taxonomy. The procedure does not manufacture periodicity: across 95 return series it returns the empty set in 89 cases, and the aggregate count of detections is indistinguishable from noise, exactly as weak-form efficiency predicts. And where it does detect structure, the interpretable cases are concentrated in the commodity futures whose seasonal and reporting calendars are independently documented — corn at the weekly frequency, lean hogs at the monthly — mirroring in financial data the pattern established for macroeconomic aggregates: detection where an institutional calendar operates, silence elsewhere.

8. Conclusion

This paper has developed a test of periodicity for economic time series that remains valid when the noise background is serially correlated. The procedure whitens the periodogram with a maximum-likelihood ARMA spectrum, applies a modification of Fisher's g-test, and selects significant frequencies through Hochberg's step-up procedure with family-wise error control. Three results — asymptotic calibration of the global test, local power along Pitman paths, and control of the family-wise error rate — give the test the kind of formal guarantee that informal inspection of a periodogram cannot.

The empirical evidence delivers two messages. The Monte Carlo study shows correct calibration across autoregressive, mixed ARMA, long-memory, and conditionally heteroskedastic backgrounds, and confirms that the procedure reports no periodic structure when applied to a differenced random walk. The taxonomy of thirty series then shows the procedure detecting periodicity exactly where institutional calendars are known to operate — seasonal lines in unadjusted money, reserve-maintenance and rollover rhythms in bank and money-market aggregates, payroll-cycle structure in claims, and FOMC and auction cycles in daily yields — and returning the empty set everywhere else. A financial cross-section of ninety-five daily return series reinforces this specificity from the opposite side of the efficiency divide: across instruments for which weak-form efficiency predicts an essentially flat spectrum, the aggregate count of detections is statistically indistinguishable from noise, and the only interpretable structure surfaces in the commodity futures — corn at the weekly frequency, lean hogs at the monthly — whose seasonal and reporting calendars are independently documented. Detection where an institutional calendar operates, and silence elsewhere, is thus the same regularity in financial as in macroeconomic data. Where the Ljung–Box diagnostic flags residual dependence, the detected frequencies are stable under the richer backgrounds of Appendix B, so the conclusions do not hinge on the ARMA approximation. The natural-rate application in Section 7.1 illustrates the downstream value of the detected component: it furnishes a measure of anticipated money growth free of any conditioning-set choice, and on this measure the strict natural-rate restriction is rejected on post-1959 quarterly data, in line with the post-Mishkin literature.

Several extensions follow naturally. Richer parametric backgrounds, sketched in Appendix B, accommodate long memory and non-rational spectra. A multivariate version built on the cross-spectrum would test for shared periodicity across series or countries. And the detected component opens a set of substantive questions in which an instrument-free measure of predictable variation is useful — the response of consumption to predictable income (Hall, 1978; Flavin, 1981; Campbell and Mankiw, 1989), the decomposition of the term premium, and the real effects of anticipated fiscal shocks. We leave these to future work.

References

- Ahdesmäki, M., H. Lähdesmäki, A. Gracey, I. Shmulevich, and O. Yli-Harja (2007). Robust regression for periodicity detection in non-uniformly sampled time-course gene expression data. *BMC Bioinformatics* 8, article 233.
- Ahdesmäki, M., H. Lähdesmäki, R. Pearson, H. Huttunen, and O. Yli-Harja (2005). Robust detection of periodic time series measured from biological systems. *BMC Bioinformatics* 6, article 117.
- Andrews, D. W. K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59(3), 817–858.
- Barro, R. J. (1977). Unanticipated money growth and unemployment in the United States. *American Economic Review* 67(2), 101–115.
- Barro, R. J. (1978). Unanticipated money, output, and the price level in the United States. *Journal of Political Economy* 86(4), 549–580.

- Baxter, M. and R. G. King (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81(4), 575–593.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B* 57(1), 289–300.
- Beran, J. (1994). *Statistics for Long-Memory Processes*. New York: Chapman and Hall.
- Beveridge, S. and C. R. Nelson (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle. *Journal of Monetary Economics* 7(2), 151–174.
- Brillinger, D. R. (1981). *Time Series: Data Analysis and Theory*. Expanded edition. San Francisco: Holden-Day.
- Brock, W. A. and D. A. Hsieh (1991). *Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence*. Cambridge: MIT Press.
- Brockwell, P. J. and R. A. Davis (1991). *Time Series: Theory and Methods*. Second edition. New York: Springer.
- Burns, A. F. and W. C. Mitchell (1946). *Measuring Business Cycles*. New York: National Bureau of Economic Research.
- Campbell, J. Y. and N. G. Mankiw (1989). Consumption, income and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual* 4, 185–216.
- Canova, F. (1998). Detrending and business cycle facts. *Journal of Monetary Economics* 41(3), 475–512.
- Christiano, L. J. and T. J. Fitzgerald (2003). The band pass filter. *International Economic Review* 44(2), 435–465.
- Chiu, S.-T. (1989). Detecting periodic components in a white Gaussian time series. *Journal of the Royal Statistical Society, Series B* 51(2), 249–259.
- Cogley, T. and J. M. Nason (1995). Effects of the Hodrick–Prescott filter on trend and difference stationary time series: Implications for business cycle research. *Journal of Economic Dynamics and Control* 19(1–2), 253–278.
- Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105(8), 2644–2678.
- Durbin, J. and S. J. Koopman (2012). *Time Series Analysis by State Space Methods*. Second edition. Oxford: Oxford University Press.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25(2), 383–417.
- Fama, E. F. and K. R. French (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business* 60(1), 55–73.
- Fisher, R. A. (1929). Tests of significance in harmonic analysis. *Proceedings of the Royal Society of London, Series A* 125(796), 54–59.
- Flavin, M. A. (1981). The adjustment of consumption to changing expectations about future income. *Journal of Political Economy* 89(5), 974–1009.
- Friedman, M. (1968). The role of monetary policy. *American Economic Review* 58(1), 1–17.

- Geweke, J. and S. Porter-Hudak (1983). The estimation and application of long-memory time series models. *Journal of Time Series Analysis* 4(4), 221–238.
- Ghysels, E. and D. R. Osborn (2001). *The Econometric Analysis of Seasonal Time Series*. Cambridge: Cambridge University Press.
- Gordon, R. J. (1982). Price inertia and policy ineffectiveness in the United States, 1890–1980. *Journal of Political Economy* 90(6), 1087–1117.
- Hall, R. E. (1978). Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86(6), 971–987.
- Hamilton, J. D. (2018). Why you should never use the Hodrick–Prescott filter. *Review of Economics and Statistics* 100(5), 831–843.
- Hannan, E. J. (1980). The estimation of the order of an ARMA process. *Annals of Statistics* 8(5), 1071–1081.
- Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Harvey, A. C. (1998). Long memory in stochastic volatility. In J. Knight and S. Satchell (Eds.), *Forecasting Volatility in the Financial Markets*, pp. 307–320. Oxford: Butterworth-Heinemann.
- Harvey, A. C. and A. Jaeger (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8(3), 231–247.
- Harvey, A. C. and T. M. Trimbur (2003). General model-based filters for extracting cycles and trends in economic time series. *Review of Economics and Statistics* 85(2), 244–255.
- Harvey, C. R., Y. Liu, and H. Zhu (2016). ... and the cross-section of expected returns. *Review of Financial Studies* 29(1), 5–68.
- Hassani, H., R. Mahmoudvand, and M. Yarmohammadi (2011). Filter-based Fisher g-test approach for periodicity detection in time series analysis. *African Journal of Mathematics and Computer Science Research* 4(7), 220–226.
- Hinich, M. J., J. Foster, and P. Wild (2009). Discrete Fourier transform filters: Cycle extraction and Gibbs effect considerations. *Macroeconomic Dynamics* 13(4), 523–534.
- Hochberg, Y. (1988). A sharper Bonferroni procedure for multiple tests of significance. *Biometrika* 75(4), 800–802.
- Hodrick, R. J. and E. C. Prescott (1980). Postwar U.S. business cycles: An empirical investigation. Working paper, Carnegie-Mellon University.
- Hodrick, R. J. and E. C. Prescott (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking* 29(1), 1–16.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6(2), 65–70.
- Hommel, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika* 75(2), 383–386.
- Hsieh, D. A. (1993). Implications of nonlinear dynamics for financial risk management. *Journal of Financial and Quantitative Analysis* 28(1), 41–64.

- Iacobucci, A. and A. Noullez (2005). A frequency selective filter for short-length time series. *Computational Economics* 25(1), 75–102.
- King, R. G. and S. T. Rebelo (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17(1–2), 207–231.
- Krafty, R. T. and W. O. Collinge (2013). Penalized multivariate Whittle likelihood for power spectrum estimation. *Biometrika* 100(2), 447–458.
- LeBaron, B. (1994). Chaos and nonlinear forecastability in economics and finance. *Philosophical Transactions of the Royal Society A* 348(1688), 397–404.
- Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory* 4(2), 103–124.
- Lo, A. W. and A. C. MacKinlay (1990). Data-snooping biases in tests of financial asset pricing models. *Review of Financial Studies* 3(3), 431–467.
- Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4), 1295–1328.
- McLean, R. D. and J. Pontiff (2016). Does academic research destroy stock return predictability? *Journal of Finance* 71(1), 5–32.
- Mishkin, F. S. (1982). Does anticipated monetary policy matter? An econometric investigation. *Journal of Political Economy* 90(1), 22–51.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica* 29(3), 315–335.
- Nelson, C. R. and H. Kang (1981). Spurious periodicity in inappropriately detrended time series. *Econometrica* 49(3), 741–751.
- Nordhaus, W. D. (1975). The political business cycle. *Review of Economic Studies* 42(2), 169–190.
- Pawitan, Y. and F. O’Sullivan (1994). Nonparametric spectral density estimation using penalized Whittle likelihood. *Journal of the American Statistical Association* 89(426), 600–610.
- Pollock, D. S. G. (2007). Wiener–Kolmogorov filtering, frequency-selective filtering and polynomial regression. *Econometric Theory* 23(1), 71–88.
- Pollock, D. S. G. (2014). Econometric methods of signal extraction. *Computational Statistics and Data Analysis* 78, 161–177.
- Pollock, D. S. G. (2018). Filters, waves and spectra. *Econometrics* 6(3), article 35.
- Priestley, M. B. (1981). *Spectral Analysis and Time Series*. London: Academic Press.
- Quinn, B. G. (2021). Recent advances in periodicity testing. *Journal of Statistical Planning and Inference* 211, 27–37.
- Reegen, P. (2007). SigSpec — I. Frequency- and phase-resolved significance in Fourier space. *Astronomy and Astrophysics* 467(3), 1353–1371.
- Robinson, P. M. (1995). Gaussian semiparametric estimation of long-range dependence. *Annals of Statistics* 23(5), 1630–1661.
- Romano, J. P., A. M. Shaikh, and M. Wolf (2008). Formalized data snooping based on generalized error rates. *Econometric Theory* 24(2), 404–447.

- Sargent, T. J. and N. Wallace (1975). Rational expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of Political Economy* 83(2), 241–254.
- Sarkar, S. K. (1998). Some probability inequalities for ordered MTP₂ random variables: a proof of the Simes conjecture. *Annals of Statistics* 26(2), 494–504.
- Schuler, Y. S. (2021). On the cyclical properties of Hamilton’s regression filter. Deutsche Bundesbank Discussion Paper, No. 03/2021.
- Simes, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* 73(3), 751–754.
- Siegel, A. F. (1980). Testing for periodicity in a time series. *Journal of the American Statistical Association* 75(370), 345–348.
- Sowell, F. (1992). Maximum likelihood estimation of stationary univariate fractionally integrated time series models. *Journal of Econometrics* 53(1–3), 165–188.
- Sullivan, R., A. Timmermann, and H. White (2001). Dangers of data mining: The case of calendar effects in stock returns. *Journal of Econometrics* 105(1), 249–286.
- Sykulski, A. M., S. C. Olhede, A. P. Guillaumin, J. M. Lilly, and J. J. Early (2019). The debiased Whittle likelihood. *Biometrika* 106(2), 251–266.
- Vaughan, S. (2010). A Bayesian test for periodic signals in red noise. *Monthly Notices of the Royal Astronomical Society* 402(1), 307–320.
- Walker, A. M. (1965). Some asymptotic results for the periodogram of a stationary time series. *Journal of the Australian Mathematical Society* 5(1), 107–128.
- Whittle, P. (1952). The simultaneous estimation of a time series harmonic components and covariance structure. *Trabajos de Estadística* 3(1–2), 43–57.
- Wichert, S., K. Fokianos, and K. Strimmer (2004). Identifying periodically expressed transcripts in microarray time series data. *Bioinformatics* 20(1), 5–20.

Appendix A. Proofs

A.1 Uniform consistency of the spectral estimator (equation 4)

We prove the uniform rate

$$\sup_{\omega \in [-\pi, \pi]} [\hat{f}(\omega) - f(\omega; \theta_0)] = O_p(N^{-1/2}).$$

The rational spectrum $f(\omega; \theta)$ is jointly continuous in (ω, θ) on the compact set $[-\pi, \pi] \times \Theta$, and continuously differentiable in θ . Under (A1), $\inf_{\omega, \theta \in V} f(\omega; \theta) > 0$ for some open neighbourhood V of θ_0 , and the gradient $\nabla_{\theta} f(\omega; \theta)$ is bounded uniformly on $[-\pi, \pi] \times V$. A first-order Taylor expansion gives, uniformly in ω ,

$$\hat{f}(\omega) - f(\omega; \theta_0) = \nabla_{\theta} f(\omega; \theta_0)^{\top} (\hat{\theta} - \theta_0) + O_p(N^{-1}),$$

where the remainder is uniform in ω by boundedness of the Hessian on $[-\pi, \pi] \times V$. The \sqrt{N} -consistency of $\hat{\theta}$ (Brockwell and Davis, 1991, Theorem 8.8.1) yields $\hat{\theta} - \theta_0 = O_p(N^{-1/2})$, and combining with the bounded gradient gives the result. The second equation in (4) follows by continuity of the reciprocal map on a strictly positive function class.

A.2 Proof of Proposition 1

The argument proceeds in three steps.

1. **Distribution of the raw periodogram.** Under (A1)–(A3) and the existence of fourth moments implied by Gaussianity of the innovations, Brillinger (1981, Theorem 5.2.6) gives, for any finite collection $j_1 < \dots < j_r$ of distinct Fourier frequencies,

$$\left(\frac{2 I_N(\omega_{j_1})}{f(\omega_{j_1})}, \dots, \frac{2 I_N(\omega_{j_r})}{f(\omega_{j_r})} \right) \rightarrow^d (2 E_1, \dots, 2 E_r),$$

with $E_1, \dots, E_r \sim \text{iid Exp}(1)$.

2. **Uniform substitution error.** By (4), proven in Appendix A.1, $\sup_j \left[\frac{f(\omega_j; \theta_0)}{\hat{f}(\omega_j)} - 1 \right] = O_p(N^{-1/2})$. Define $E_j = 2 I_N(\omega_j) / f(\omega_j; \theta_0)$ and $\tilde{E}_j = 2 \tilde{I}_j$, giving $\tilde{E}_j = E_j(1 + r_j)$ with $\max_j |r_j| = O_p(N^{-1/2})$.
3. **Stability of the Fisher statistic.** Standard extreme-value theory for exponentials gives $\max_j E_j = \log m + O_p(1)$, and the law of large numbers yields $\sum_{j=1}^m E_j = m + O_p(\sqrt{m})$. A direct algebraic manipulation of the ratio that defines \tilde{g} yields

$$\tilde{g} - g_0 = O_p\left(\frac{\log m}{\sqrt{N}}\right) = o_p(1),$$

where g_0 is Fisher's statistic computed from the unperturbed ordinates E_j . By Fisher (1929), g_0 follows exactly the distribution F_{G_m} . Continuity of F_{G_m} combined with the uniform $o_p(1)$ approximation yields, by Slutsky's theorem,

$$\sup_{x \in (0,1)} [P(\tilde{g} \leq x) - F_{G_m}(x)] \rightarrow 0,$$

completing the proof.

A.3 Proof of Proposition 2

1. **Decomposition of the DFT.** Set $S_t = A_N \cos(\omega_{j_0} t + \varphi)$. By linearity of the discrete Fourier transform, $d_X(\omega_j) = d_S(\omega_j) + d_Y(\omega_j)$, with the DFT normalised so that $d_X(\omega) = N^{-1/2} \sum_t x_t e^{-i\omega t}$.
2. **Signal contribution.** Direct calculation at the target frequency yields $d_S(\omega_{j_0}) = (c/2) e^{i\varphi} + o(1)$. For $j \neq j_0$, orthogonality of sinusoids at Fourier frequencies implies $d_S(\omega_j) = O(N^{-1/2})$, and the contribution to $I_N(\omega_j)$ is asymptotically negligible.
3. **Limiting distribution at the target frequency.** Under the ARMA assumption on $\{y_t\}$, $d_Y(\omega_{j_0}) \rightarrow^d \mathcal{CN}(0, 2\pi f(\omega_{j_0}))$. Combining with Step 2 and applying the whitening transformation,

$$\tilde{I}_{j_0} \rightarrow^d \frac{1}{4\pi f(\omega_{j_0})} [c e^{i\varphi} + Z]^2, \quad Z \sim \mathcal{CN}(0, 4\pi f(\omega_{j_0})),$$

which is a non-central χ_2^2 variate with non-centrality parameter $\lambda = c^2 / [4\pi f(\omega_{j_0})]$. For $j \neq j_0$, $\tilde{I}_j \rightarrow^d \text{Exp}(1)$, asymptotically independent.

4. **Limiting power.** By concentration of the denominator, $\sum_j \tilde{I}_j / m \rightarrow^p 1$. The rejection event reduces, in the limit, to the comparison between the non-central variate at the target frequency and the maximum of $m - 1$ iid $\text{Exp}(1)$ variates. Hence

$$\beta = P(W_\lambda > m x_\alpha) + P(W_\lambda \leq m x_\alpha) \cdot P(\max_i E_i > m x_\alpha),$$

where W_λ is the limiting non-central χ^2_2 variate of Step 3 and the E_i are iid $\text{Exp}(1)$.

A.4 Proof of Proposition 3

1. Sufficient condition. Sarkar (1998, Theorem 3.1) establishes that Hochberg's step-up procedure controls FWER at level α under positive regression dependence on the subset of true nulls (PRDS). Asymptotic independence — the strongest form of PRDS — therefore suffices.
2. Asymptotic independence of the periodogram. By Brillinger (1981, Theorem 4.4.1), $\text{Cov}[I_N(\omega_j), I_N(\omega_k)] = O(N^{-1})$ for $j \neq k$, with finite-dimensional distributions converging to products of independent exponentials. The whitened ordinates inherit this property by Slutsky's theorem and the uniform consistency of the spectral estimator (Appendix A.1).
3. Strong-mixing extension. The residual $O(N^{-1})$ pairwise correlations between whitened ordinates satisfy a strong-mixing condition; under this condition the conclusions of Sarkar (1998) extend asymptotically, yielding $\limsup_{N \rightarrow \infty} \text{FWER}_N \leq \alpha$.
4. Conservative fallback. Holm's (1979) step-down procedure, which uses the same critical values $c(i) = \alpha / (m - i + 1)$ but applied in step-down order, delivers $\text{FWER} \leq \alpha$ under arbitrary dependence among the p-values and may be substituted without further conditions.

Appendix B. Extensions beyond the ARMA class

This appendix records the three parametric extensions referenced in Section 3.8. Formally, let $\mathcal{F} = f(\cdot; \theta) : \theta \in \Theta$ denote a parametric class of spectral densities, with $\Theta \subset \mathbb{R}^k$ compact and $f(\omega; \theta)$ continuous in (ω, θ) and twice continuously differentiable in θ . Let $\hat{\theta}$ be a \sqrt{N} -consistent estimator of θ_0 and $\hat{f}(\omega) = f(\omega; \hat{\theta})$ the plug-in spectral estimator. Replace assumptions (A1)–(A4) of Section 3.1 with the following.

(A1*) The true spectral density of x_t satisfies $f = f(\cdot; \theta_0) \in \mathcal{F}$, with $\inf_{\omega} f(\omega) > 0$.

(A2*) The estimator $\hat{\theta}$ satisfies $\sqrt{N}(\hat{\theta} - \theta_0) = O_p(1)$, and the gradient $\nabla \theta f(\cdot; \theta)$ is uniformly bounded on $[-\pi, \pi] \times V$ for some open neighbourhood V of θ_0 .

Under (A1*) and (A2*), the uniform consistency of the spectral estimator at rate \sqrt{N} continues to hold; its proof in Appendix A.1 is unchanged. Propositions 1 to 3 follow without modification. The ARMA case treated in Section 3 is the leading instance of this more general framework, distinguished operationally by the availability of residual-based diagnostics rather than by any feature of the asymptotic theory.

B.1 Long-memory backgrounds: the ARFIMA family

The ARFIMA(p, d, q) family extends ARMA by introducing a fractional differencing parameter $d \in (-1/2, 1/2)$:

$$\Phi(L)(1 - L)^d X_t = \Theta(L)\varepsilon_t.$$

The corresponding spectral density

$$f(\omega; \theta) = \frac{\sigma^2}{2\pi} \frac{|\Theta(e^{-i\omega})|^2}{|\Phi(e^{-i\omega})|^2} |1 - e^{-i\omega}|^{-2d}$$

captures the hyperbolically decaying autocorrelations characteristic of long memory, prevalent in inflation series, realized volatility, and macroeconomic spreads. The exact maximum-likelihood estimator developed by Sowell (1992) is \sqrt{N} -consistent for $|d| < 1/4$; consistent estimators with the same rate are available throughout the stationary range (Beran, 1994; Robinson, 1995). The diagnostic battery of Section 3.7 applies, with the Geweke–Porter–Hudak (1983) test added to verify that long memory has been correctly captured. When the true generating process exhibits long memory and one (mis)specifies an ARMA background, the calibration of Proposition 1 deteriorates; using the ARFIMA background restores nominal size. We document this trade-off empirically in Section 6.

B.2 Non-rational backgrounds: penalized splines

When the true spectral density admits no finite-dimensional rational representation, a flexible alternative is to model $\log f(\omega)$ as a penalised spline of ω , fitted to the log-periodogram. The Whittle penalized likelihood approach of Pawitan and O’Sullivan (1994) and the multivariate extension of Krafty and Collinge (2013) yield estimators with uniform convergence rates of order $N^{-r/(2r+1)}$, where r is the smoothness of the true log-spectrum. The rate is slower than \sqrt{N} , reflecting the non-parametric nature of the estimator; however, the resulting test maintains correct asymptotic size, and the analyst gains protection against arbitrary smooth departures from the rational class. The trade-off is the loss of interpretable diagnostic residuals — there is no analogue of the Ljung–Box statistic in this setting — and a hyperparameter (the penalty weight) that must be selected, typically by generalized cross-validation.

B.3 Continuous-parameter alternatives: the Matérn family

For applications in which the autocorrelation function admits a parsimonious continuous-parameter description, the Matérn family

$$f(\omega; \theta) = \frac{\sigma^2}{(\alpha^2 + \omega^2)^{\nu + 1/2}}$$

offers a three-parameter representation that nests AR(1) ($\nu = 1/2$), continuous-time analogues of higher-order ARMA processes, and a wide range of intermediate-memory behaviours. The debiased Whittle methodology of Sykulski et al. (2019) provides \sqrt{N} -consistent estimation of (σ^2, α, ν) at finite-sample bias of strictly smaller order than classical Whittle estimation. The test calibrates correctly under (A1*)–(A2*); the operational disadvantage relative to ARMA is the loss of the Box–Jenkins diagnostic apparatus and a smaller installed base of econometric software.

B.4 Practical recommendation

In macroeconomic applications, our recommendation is to use ARMA as the default and to escalate to ARFIMA when long memory is detected by the Geweke–Porter–Hudak test or by visual inspection of the autocorrelation function. The penalized-spline and Matérn alternatives are warranted when both ARMA and ARFIMA are rejected by the residual diagnostics, signalling a non-rational spectral background. The Monte Carlo evidence in Section 6 documents the size and power consequences of these choices.