Scale-Ordered Contagion: A Spectral Theory of Heterogeneous Information Adaptation in Financial Networks

Avishek Bhandari* Ipsita Parida† Anoop S. Kumar‡

June 17, 2026

Abstract

The speed at which a market processes information shapes not only how strongly but *at what frequency* a cross-border shock registers in it. The manner in which such information adaptation governs the frequency composition of transmitted shocks has, however, received limited formal attention in the contagion literature, and it is this gap which motivates the present study. We develop a spectral theory of financial contagion under heterogeneous information adaptation, building on the Heterogeneous Agents Contagion versus Interdependence (HACI) framework in which advanced economies act as fast adapters and emerging economies as slow adapters. We first identify a structural limitation of the baseline: its information-flow dynamics are *source-constrained*, so the receiving market's adaptive capacity affects the magnitude but not the timing or frequency composition of transmitted information; we prove that fast-adapter and slow-adapter recipients of a common shock share the half-life $\ln 2/\alpha_s$. In view of the above, we resolve the limitation by modelling the receiver as a second exponential information filter, so that observed transmission is the convolution of source generation and receiver absorption. The resulting bi-exponential response has a power spectrum equal to the product of two Lorentzians, one centred at the source rate and one at the receiver rate, and the slower of the two markets supplies the binding spectral corner. Projecting this spectrum onto a maximal-overlap discrete wavelet basis yields a closed-form transfer-entropy-by-scale profile and three falsifiable predictions, jointly the Scale-Ordered Contagion Hypothesis (SOCH): (A) the wavelet scale at which directed transfer entropy peaks is set by the slower market's adaptation rate, so emerging-inclusive pairs peak at coarser scales; (B) the *shape* of the scale profile is symmetric across direction for a given pair, because the transmission spectrum is symmetric in the two rates; and (C) the *magnitude* is directionally asymmetric, scaling with connectivity and source-shock content. We turn the theory into a profile-matching estimator that recovers adaptation rates and endogenises the fast/slow classification, and we test all three predictions on a panel of G20 equity markets spanning January 2006 through March 2026. The scale-ordering prediction is supported (peak scales are coarser for pairs containing slower markets, $p = 0.042$); the shape-symmetry prediction—the theory's sharpest and most novel restriction—is supported across all twenty-eight unordered market pairs ($p > 0.05$ throughout); the magnitude prediction holds directionally but is not significant in the working sample; and the endogenous classifier isolates the slowest adapters cleanly while, as the identification analysis itself predicts, it cannot

*School of Humanities, Social Sciences and Management, Indian Institute of Technology Bhubaneswar. Email: avishekb@iitbbs.ac.in. Corresponding author.

†School of Humanities, Social Sciences and Management, Indian Institute of Technology Bhubaneswar. Email: a23hs09014@iitbbs.ac.in.

‡Gulati Institute of Finance and Taxation, Thiruvananthapuram 695017, Kerala, India. Email: askumar@gift.res.in.

separate fast markets whose spectral corner approaches the sampling Nyquist frequency. The framework supplies a structural, testable account of why short-band and long-band contagion differ systematically with the adaptation speeds of the markets involved.

Keywords: financial contagion; transfer entropy; wavelet decomposition; quantile dependence; information rigidity; heterogeneous agents; spectral methods; G20 equity markets.

JEL classification: G01, G15, C58, D85.

1 Introduction

The relationship between financial shock transmission and the speed at which markets process incoming information has been the prime focus of researchers for many years. It may be noted that such transmission is governed not merely by the strength of the linkages between markets, but also, and rather more subtly, by the speed at which each market processes the information reaching it. A long line of work distinguishes interdependence, the stable comovement that persists across tranquil and crisis periods, from contagion, the increase in cross-market linkage that arises after a shock [Forbes and Rigobon, 2002, Diebold and Yilmaz, 2012], and a parallel literature on information rigidity, beginning with Mankiw and Reis [2002] and formalised by Coibion and Gorodnichenko [2015], has established that economic agents update their information sets infrequently and at heterogeneous rates. The Heterogeneous Agents Contagion versus Interdependence (HACI) framework combines these ideas by treating advanced economies as fast adapters with a high information-absorption rate α_F and emerging economies as slow adapters with rate $\alpha_S < \alpha_F$, and by arguing that the difference in adaptation speeds generates systematic asymmetries in cross-market information flow. That framework is the theoretical core of a broader research programme on heterogeneous-adaptation contagion whose empirical apparatus is wavelet-quantile transfer entropy; the present paper positions itself in this context and seeks to examine the spectral theory that connects the apparatus to a structural mechanism, and brings its predictions to data.

The starting point of this paper is a limitation of the baseline HACI dynamics that, once stated precisely, is difficult to defend on economic grounds. It may be noted that in the baseline specification the cumulative information flow from a source market to a receiving market decays at the *source's* adaptation rate; the receiver's rate enters only as a multiplicative constant on the level of the response. A direct consequence, which we state and prove, is that the half-life of the contagion response is identical for fast-adapter and slow-adapter recipients of a common shock. The receiving market's adaptive capacity therefore has no effect whatsoever on the timing or the frequency composition of the transmitted shock—a restrictive feature that sits uneasily with the economic content of information rigidity, under which a market that processes information slowly ought to register an incoming shock more gradually, and hence at lower frequencies, than a market that processes information quickly. Thus the baseline specification, though internally consistent, is at odds with the very intuition that motivates the heterogeneous-adaptation programme.

In view of the above, we resolve this limitation by allowing the receiving market to process incoming information through its own exponential filter, which is the natural formalisation of the simple idea that the receiver, much like the source, updates its information set only at a finite rate. Observed transmission then becomes the convolution of the source's information-generation kernel with the receiver's absorption kernel. The convolution of two exponentials is bi-exponential, and its power spectral density is the product of two Lorentzian (Cauchy) spectra, one with corner frequency at the source rate α_s and one at the receiver rate α_r. Because a product of low-pass filters is dominated by the filter with the lower corner frequency, the *slower* of the two markets governs the frequency band in which transmitted information concentrates. This single modification restores a

symmetric and economically sensible role for both markets and, as we show, carries sharp observable implications once the response is measured across time scales.

The bridge from the continuous-time information dynamics to observable quantities runs through multi-scale transfer entropy, which is the empirical object used to measure directed information flow in financial networks [Schreiber, 2000, Kraskov et al., 2004]. Projecting the transmission spectrum onto a maximal-overlap discrete wavelet transform (MODWT) basis [Percival and Walden, 2000, Gençay et al., 2005] yields a closed-form expression for transfer entropy by wavelet scale, and from it three falsifiable predictions, which we collect under the name Scale-Ordered Contagion Hypothesis (SOCH). The hypothesis states that the scale at which directed transfer entropy peaks is set by the slower market's adaptation rate, so that emerging-market pairs peak at coarser scales; that the *shape* of the scale profile is the same in both directions for a given pair, because the transmission spectrum is symmetric in the two adaptation rates; and that the *magnitude* of transfer entropy is directionally asymmetric, scaling with directional connectivity and with the information content of the source's shock.

The shape-symmetry prediction is, to our knowledge, novel, and it is pertinent to emphasise how sharp it is. It distinguishes the heterogeneous-adaptation mechanism from contagion mechanisms that would generate directional asymmetry in the frequency composition itself, such as mechanisms operating through asymmetric balance-sheet exposure or asymmetric trade dependence: those mechanisms need not respect shape symmetry, whereas the adaptation-speed mechanism must. The prediction is also directly testable, because the normalised scale profile in each direction is an estimable object and the two directions can be compared against a sampling null.

The theory does more than restrict the shape of observed profiles; it also makes adaptation rates estimable, which we regard as a further contribution. Fitting an observed wavelet-quantile transfer-entropy profile to the closed-form spectrum recovers the pair's two adaptation rates, and pooling the directed pairs that share a market yields a market-level rate; the fast/slow classification, exogenous in the baseline, thereby becomes a data-determined and testable object. We make this estimator operational, prove the conditions under which its parameters are identified, characterise its finite-sample behaviour by Monte Carlo, and then take it, together with the three SOCH predictions, to a panel of eighteen G20 equity markets over 2006–2026.

The empirical findings are, we believe, honest and informative rather than uniformly confirmatory, which is the appropriate standard for a sharp theory. We observe that the scale-ordering prediction is supported: pairs containing more emerging markets peak at coarser scales, with a positive and significant association. The shape-symmetry prediction is supported across every one of the twenty-eight unordered market pairs—the strongest of our results and the one that most distinguishes the mechanism from its rivals. The magnitude prediction holds in direction—advanced-to-emerging flows dominate in most pairs—but is not significant in the working sample. The endogenous classifier identifies the two clearest emerging markets, India and China, as decisively the slowest adapters, but cannot separate the faster markets, whose spectral corners lie near the sampling Nyquist frequency; this is not a defect of the data but a property that the identi-

fication analysis predicts in advance. Hence the two predictions that rest on the well-identified slower-market rate are supported, and the predictions that depend on the weakly-identified fast rates are correspondingly weaker.

The remainder of the paper is organised as follows. Section 2 places the contribution in the contagion, information-rigidity, and wavelet-quantile literatures. Section 3 restates the HACI baseline and proves the source-constrained timing result. Section 4 introduces the receiver-filtering extension and derives the bi-exponential transmission response. Section 5 establishes the product-Lorentzian spectrum, the closed-form scale profile, the peak-location result, and a discrete-time companion. Section 6 states and proves the three parts of SOCH and an early-warning corollary. Section 7 develops identification, the profile-matching estimator, and a Monte Carlo study. Section 8 describes the data. Section 9 reports the empirical results. Section 10 relates the theory to frequency-domain connectedness and systemic-risk measurement. Section 11 concludes. An appendix collects the longer derivations.

2 Related Literature

2.1 Contagion, interdependence, and information rigidity

The manner in which financial shocks propagate across borders has been a prime focus of researchers for many years, and the modern study of cross-border contagion takes as its point of departure the fundamental distinction between interdependence and contagion proper [Forbes and Rigobon, 2002]. A substantial body of work has since established that cross-market comovement intensifies during periods of financial stress, that information flows tend to reorganise around centres of distress, and that the topology of the global financial network shifts in ways that are, by and large, measurable and systematic [Diebold and Yilmaz, 2012, Diebold and Yılmaz, 2014, Billio et al., 2012]. It may be noted, however, that this literature has largely treated the *speed* of cross-market adjustment as a nuisance property to be differenced away, rather than as an object of structural interest in its own right. The information-rigidity literature supplies exactly the structural content that has been missing from this picture. Mankiw and Reis [2002] model agents who update their information sets only intermittently, and Coibion and Gorodnichenko [2015] document, across a broad range of forecasters and countries, that information is updated infrequently and at heterogeneous rates, with first-order consequences for how shocks propagate through the system. The HACI framework imports this heterogeneity into a contagion setting, positing that the adaptation-rate gap between advanced and emerging markets is itself a source of systematic asymmetry in transmission. The present paper positions itself in this context and seeks to provide the missing spectral mechanics of that idea.

2.2 Heterogeneous adaptation and the source-constrained baseline

The HACI baseline writes cumulative cross-market information flow as a product of a source-generated impulse and a receiver-specific absorptive constant. As we show in Section 3, this con-

struction makes the receiver's rate enter the level of the response but not its dynamics, so that the timing of adjustment is governed entirely by the source market. The economic implausibility of this property—two markets of very different sophistication absorbing the same shock on precisely the same schedule—motivates the central modelling move of the present paper, which is to endow the receiver with its own filter. The idea that absorption is itself a dynamic, rate-limited process is, of course, standard in the information-rigidity tradition; what is new here is the observation that, once the receiver is permitted to filter, the transmitted response and its frequency content are jointly determined by both adaptation rates in a particular and empirically testable way.

2.3 Wavelet-quantile transfer entropy

Directed information flow between financial time series is naturally measured by transfer entropy [Schreiber, 2000], estimated either through nearest-neighbour methods [Kraskov et al., 2004] or, in the Gaussian case, through its exact equivalence with Granger causality [Barnett et al., 2009]. Two refinements matter considerably for the contagion application. First, the maximal-overlap discrete wavelet transform decomposes a return series into scale-specific components, thereby separating short-horizon trader-driven dynamics from the longer-horizon dynamics that characterise institutional investors [Percival and Walden, 2000, Gençay et al., 2005, Crowley, 2007]. Second, conditioning on quantiles of the recipient distribution serves to isolate tail-specific propagation which aggregate measures tend to obscure [Han et al., 2016, Ando et al., 2022]. The combination, wavelet-quantile transfer entropy (WQTE), is the empirical primitive of the heterogeneous-adaptation programme and the object whose scale profile our theory seeks to predict. We notice, in this connection, that frequency-domain connectedness [Baruník and Křehlík, 2018] decomposes variance spillovers into frequency bands descriptively; our contribution is to supply a structural reason why those bands should differ systematically with adaptation speeds, which is a rather different—and, we would argue, more informative—undertaking.

3 The HACI Baseline and Its Source-Constrained Limitation

The analysis begins with a careful consideration of the baseline HACI information dynamics, and we make precise the sense in which they are source-constrained. A shock originates in a source market s at time t_0 and releases information that diffuses into the market at the source's own absorption rate.

Assumption 1 (Source information generation)**.** The cumulative information released by a shock in market s and available for transmission evolves as $I_{s,t} = I_{s,\infty}\big(1 - e^{-\alpha_s(t-t_0)}\big)$ for $t \geq t_0$, where $I_{s,\infty}$ is the total information content of the shock and $\alpha_s > 0$ is the source market's absorption rate. The instantaneous generation rate is $\phi_s(t) = I_{s,\infty}\,\alpha_s\,e^{-\alpha_s(t-t_0)}$.

In the baseline model the cumulative information flow from s to a receiver r is obtained by scaling the source impulse by the receiver's absorptive capacity and by a connectivity-and-friction

term,

$$CIF^{\text{base}}_{s\to r}(t) = \kappa_{sr}\, I_{s,\infty}\, \frac{\alpha_r}{1+\lambda_{sr}} \left(1 - e^{-\alpha_s (t-t_0)}\right), \tag{1}$$

where $\kappa_{sr} \geq 0$ is the directional connectivity, $\lambda_{sr} \geq 0$ the information friction, and $\alpha_r \in \{\alpha_S, \alpha_F\}$ the receiver's rate. It may be noted that the receiver's rate enters only as a scalar multiplier on the level of the response, whereas the time profile $1 - e^{-\alpha_s(t-t_0)}$ depends entirely on the source rate and not at all on the receiver. The following result makes the consequence of this construction explicit.

Proposition 1 (Source-constrained timing)**.** *Under* (1)*, the half-life of the cumulative information flow—the elapsed time at which it attains one half of its long-run value—is* $t_{1/2} - t_0 = \ln 2/\alpha_s$*, independent of the receiver's rate* α_r*. Consequently, for a common source, a fast-adapter receiver and a slow-adapter receiver reach one half of their respective long-run responses at the same time.*

Proof. The long-run value is $\kappa_{sr} I_{s,\infty} \alpha_r/(1+\lambda_{sr})$. Setting the flow equal to half of it gives $1 - e^{-\alpha_s(t_{1/2}-t_0)} = \frac{1}{2}$, hence $e^{-\alpha_s(t_{1/2}-t_0)} = \frac{1}{2}$ and $t_{1/2} - t_0 = \ln 2/\alpha_s$; the receiver rate α_r divides out and does not appear. □

Proposition 1 isolates precisely the limitation we seek to address. The working principle of the baseline is quite simple, and its implication is stark: the receiver's adaptive capacity is permitted to affect only the amplitude of the contagion response, not its timing, and—as Section 5 shows in detail—not its frequency composition either. Thus a slow-adapting market and a fast-adapting market, when exposed to the same source shock, exhibit contagion responses that differ only by a scalar multiple. Intuitively, this is at odds with the economic content of the information-rigidity literature, which places great weight on the idea that heterogeneous adaptation rates should produce heterogeneous dynamics, not merely heterogeneous amplitudes. The next section removes this restriction.

4 Receiver Filtering and the Bi-Exponential Response

In view of the above, we now allow the receiving market to process incoming information through its own exponential filter, which is the natural formalisation of the simple idea that the receiver, much like the source, updates its information set only at a finite rate.

Assumption 2 (Receiver absorption)**.** Market r absorbs an instantaneous unit of incoming information through the causal exponential kernel $g_r(\tau) = \alpha_r\, e^{-\alpha_r \tau}$ for $\tau \geq 0$, which integrates to one, so receiver filtering redistributes incoming information over time without creating or destroying it.

The information from the source that registers in the receiver is then the convolution of the source generation rate with the receiver absorption kernel, scaled by connectivity and friction: $\psi_{s\to r}(t) = \frac{\kappa_{sr}}{1+\lambda_{sr}} (\phi_s * g_r)(t)$.

Following the standard convolution argument, we arrive at the following result, which constitutes the central analytical contribution of this section.

Proposition 2 (Bi-exponential transmission response)**.** *Under Assumptions 1 and 2, and normalising $t_0 = 0$, the transmitted response is, for $\alpha_s \neq \alpha_r$,*

$$\psi_{s\to r}(t) = A_{sr}\,\frac{\alpha_s\alpha_r}{\alpha_s-\alpha_r}\left(e^{-\alpha_r t}-e^{-\alpha_s t}\right), \qquad A_{sr} \equiv \frac{\kappa_{sr} I_{s,\infty}}{1+\lambda_{sr}}, \tag{2}$$

and, in the confluent case $\alpha_s = \alpha_r = \alpha$, $\psi_{s\to r}(t) = A_{sr}\,\alpha^2 t\, e^{-\alpha t}$. The total transmitted information is symmetric in the two rates, $\int_0^\infty \psi_{s\to r}(t)\,dt = A_{sr}$.

Proof. With $t_0 = 0$, $(\phi_s * g_r)(t) = I_{s,\infty}\alpha_s\alpha_r e^{-\alpha_r t}\int_0^t e^{-(\alpha_s-\alpha_r)u}\,du$. For $\alpha_s \neq \alpha_r$ the integral is $(1-e^{-(\alpha_s-\alpha_r)t})/(\alpha_s-\alpha_r)$, which after multiplication by $e^{-\alpha_r t}$ gives (2); for $\alpha_s = \alpha_r$ the integral is t, giving the confluent form. Each kernel integrates to one, so the convolution integrates to $I_{s,\infty}$ and the total, after scaling, is A_{sr}, manifestly invariant under $\alpha_s \leftrightarrow \alpha_r$. □

Two features are particularly pertinent for the contagion application. We notice, first, that the response now depends on both adaptation rates through the bi-exponential difference, so the timing limitation of Proposition 1 no longer holds: the receiver's rate enters the dynamics of the response, not merely its level. Moreover, and rather more subtly, the *total* transmitted information is symmetric in the two rates, which leads us to conclude that any directional asymmetry in transmitted quantity must arise from the connectivity and shock-content term A_{sr}, and not from adaptation speeds themselves. These two observations become, in the frequency domain, the content of the Scale-Ordered Contagion Hypothesis.

5 The Transmission Spectrum and Its Wavelet Decomposition

We pass to the frequency domain, where the bi-exponential response assumes a particularly transparent and tractable form. It may be noted that all symbolic results presented in this section were verified independently with a computer algebra system, so that the reader may place full confidence in the derivations which follow.

Proposition 3 (Product-Lorentzian spectrum)**.** *The Fourier transform of (2) is $\hat\psi_{s\to r}(\omega) = A_{sr}\,\frac{\alpha_s}{\alpha_s+i\omega}\,\frac{\alpha_r}{\alpha_r+i\omega}$, and the power spectral density is the product of two Lorentzians,*

$$S_{s\to r}(\omega) = \left|\hat\psi_{s\to r}(\omega)\right|^2 = A_{sr}^2\,\frac{\alpha_s^2}{\alpha_s^2+\omega^2}\,\frac{\alpha_r^2}{\alpha_r^2+\omega^2}, \tag{3}$$

which is invariant under $\alpha_s \leftrightarrow \alpha_r$.

Proof. The causal exponential $\alpha e^{-\alpha t}$ has transform $\alpha/(\alpha+i\omega)$ for $\mathrm{Re}\,\alpha > 0$; by the convolution theorem the transform of the response is the product of the two factors, and the squared modulus, evaluated at real ω using $|\alpha/(\alpha+i\omega)|^2 = \alpha^2/(\alpha^2+\omega^2)$, gives (3). □

Corollary 4 (Binding spectral corner)**.** *Write $\alpha_\wedge = \min(\alpha_s,\alpha_r)$ and $\alpha_\vee = \max(\alpha_s,\alpha_r)$. The density (3) is approximately flat for $\omega \ll \alpha_\wedge$, decays as ω^{-2} for $\alpha_\wedge \ll \omega \ll \alpha_\vee$, and decays as*

ω^{-4} *for* $\omega \gg \alpha_\vee$. *The lower corner* $\alpha_\wedge$ *marks the onset of roll-off and governs the band in which transmitted information concentrates.*

Corollary 4 is, in a very real sense, the economic heart of the theory, and it is worth dwelling upon its implications at some length. The working principle is quite simple: whichever market is the slower of the two determines the corner frequency below which contagion power is permitted to accumulate. Thus, a fast source paired with a slow receiver will find its transmission spectrum cut off at the receiver's lower corner frequency; conversely, a slow source paired with a fast receiver encounters the cut-off imposed by the source's own lower corner. In both cases, and this is the essential observation, the frequency composition of contagion is governed by the slower of the two markets, whichever party occupies that position. It may be noted, further, that the intermediate ω^{-2} regime is pronounced only when the two corners are well separated from one another—that is, when $\alpha_\wedge \ll \alpha_\vee$—whereas for markets with comparable adaptation rates the spectrum rolls directly from flat toward ω^{-4} without sustaining an extended ω^{-2} plateau.

We observe, moreover, that directed financial flow is measured empirically across MODWT scales. With unit sampling interval and Nyquist frequency π, detail level k carries a nominal pass-band $\mathcal{B}_k = [\pi 2^{-k},\, \pi 2^{-(k-1)}]$; for daily data, level 1 captures the two-to-four-day band and level 5 captures the thirty-two-to-sixty-four-day band. The power transmitted at any given scale is obtained by integrating the spectrum over the corresponding band, and it is this quantity which the following proposition evaluates in closed form.

Proposition 5 (Closed-form scale power)**.** *For* $\alpha_s \neq \alpha_r$, *the power transmitted at scale* k *is* $P_k = \int_{\mathcal{B}_k} S_{s\to r}(\omega)\, d\omega$, *which evaluates to*

$$P_k = A_{sr}^2 \, \frac{\alpha_s^2 \alpha_r^2}{\alpha_r^2 - \alpha_s^2} \left[\frac{1}{\alpha_s} \arctan \frac{\omega}{\alpha_s} - \frac{1}{\alpha_r} \arctan \frac{\omega}{\alpha_r} \right]_{\omega = \pi 2^{-k}}^{\omega = \pi 2^{-(k-1)}} . \tag{4}$$

The profile P_k *is symmetric in* (α_s, α_r) *and strictly positive.*

Proof. Partial fractions give $[(\alpha_s^2 + \omega^2)(\alpha_r^2 + \omega^2)]^{-1} = (\alpha_r^2 - \alpha_s^2)^{-1}[(\alpha_s^2 + \omega^2)^{-1} - (\alpha_r^2 + \omega^2)^{-1}]$; integrating each term with $\int (\alpha^2 + \omega^2)^{-1} d\omega = \alpha^{-1} \arctan(\omega/\alpha)$ over the band yields (4). The prefactor and the arctan bracket are each antisymmetric under the swap, so their product is symmetric; positivity follows from $S_{s\to r} \geq 0$. □

It is pertinent to observe at this juncture that, because MODWT bands are octaves of constant width on a logarithmic frequency axis, the per-scale power profile tracks the quantity $\omega\, S_{s\to r}(\omega)$ evaluated at the band centre. The peak scale is therefore governed by the maximiser of that quantity, and the following lemma characterises this maximiser precisely, establishing the bounds which will prove indispensable to the empirical predictions of the next section.

Lemma 6 (Spectral peak location)**.** *The function* $\omega\, S_{s\to r}(\omega)$ *is single-peaked on* $(0, \infty)$, *and its maximiser solves* $1 = 2\omega^2/(\alpha_s^2 + \omega^2) + 2\omega^2/(\alpha_r^2 + \omega^2)$. *In the symmetric case* $\omega^\star = \alpha/\sqrt{3}$; *in the strongly asymmetric case* $\omega^\star \to \alpha_\wedge$; *and in general* $\omega^\star \in [\alpha_\wedge/\sqrt{3}, \alpha_\wedge]$.

Proof. Maximising $\ln(\omega S) = \text{const} + \ln\omega - \ln(\alpha_s^2 + \omega^2) - \ln(\alpha_r^2 + \omega^2)$ gives the stated first-order condition, whose right-hand side increases from 0 to 4, so the maximiser is unique. Normalising $\alpha_\wedge = 1$ and writing $R = \alpha_\vee/\alpha_\wedge \geq 1$, the right-hand side minus one equals $3(R^2-1)/(2(3R^2+1)) \geq 0$ at $\omega = 1/\sqrt{3}$ and $-2/(R^2+1) < 0$ at $\omega = 1$, and is strictly increasing in ω, so the unique root lies in $[\alpha_\wedge/\sqrt{3}, \alpha_\wedge]$, attaining the lower endpoint only when $\alpha_s = \alpha_r$ and the upper only as $R \to \infty$. □

Finally, because data are sampled at discrete intervals rather than observed in continuous time, it is necessary to record the discrete-time counterpart of the foregoing results and to confirm that the theoretical predictions carry over to the empirically relevant setting. Sampling the bi-exponential response (2) at interval Δt yields the impulse response of an ARMA$(2,1)$ process with autoregressive roots $e^{-\alpha_s \Delta t}$ and $e^{-\alpha_r \Delta t}$ and a single moving-average (sampling) zero; its spectral density converges in shape to the product-Lorentzian (3) as the adaptation rates become slow relative to the sampling frequency, since $|1 - e^{-\alpha\Delta t}e^{-i\omega\Delta t}|^2 = (\alpha^2 + \omega^2)\Delta t^2 + O(\Delta t^3)$. We observe, therefore, that the scale-ordering and shape-symmetry predictions developed below carry over in full to sampled daily data, which is the empirically relevant regime for the analysis that follows.

6 The Scale-Ordered Contagion Hypothesis

The present section develops the three empirical predictions which together constitute the Scale-Ordered Contagion Hypothesis. To fix ideas, let $k^\star_{sr}$ denote the wavelet scale at which the directed transfer entropy $\text{WQTE}^{(k)}_{s\to r}$ attains its largest value, and recall that $\text{WQTE}^{(k)}_{s\to r}$ is monotone increasing in the transmitted power P_k in the small-flow regime in which transfer entropy is approximately linear in transmitted information [Schreiber, 2000, Kraskov et al., 2004]. The three predictions which follow arise directly and naturally from the results established in Section 5.

Hypothesis 1 (SOCH-A: scale ordering by adaptation speed)**.** The peak scale is governed by the slower market's rate, $\omega_{k^\star_{sr}} \asymp \alpha_\wedge$, so that $k^\star_{sr} \approx \log_2(\pi/\alpha_\wedge)$. Pairs whose slower member adapts more slowly (smaller $\alpha_\wedge$) peak at coarser scales; emerging-to-emerging pairs peak at the coarsest scales, mixed pairs at intermediate scales, and advanced-to-advanced pairs at the finest scales.

The derivation of this prediction is a direct consequence of Lemma 6. By that result, the maximiser of $\omega S(\omega)$ lies in the interval $[\alpha_\wedge/\sqrt{3}, \alpha_\wedge]$, so that $\omega_{k^\star} \asymp \alpha_\wedge$ and, upon inverting the relation $\omega_k = \pi 2^{-k}$, one obtains $k^\star \approx \log_2(\pi/\alpha_\wedge)$, which is seen to be decreasing in $\alpha_\wedge$. Put differently, the slower a market's adaptation rate, the coarser the scale at which contagion between it and its partner concentrates.

Corollary 7 (Early-warning and efficiency signal)**.** *Since $k^\star \approx \log_2(\pi/\alpha_\wedge)$ and $\partial k^\star/\partial\alpha_\wedge = -1/(\alpha_\wedge \ln 2) < 0$, a shift of the empirical peak scale toward finer scales over time is a monotone signal that the slower market's absorption rate has risen, that is, that its informational efficiency has improved; a coarsening signals the reverse.*

This corollary is of particular practical importance, for it furnishes a model-based diagnostic of changes in market efficiency that is directly readable from the wavelet transfer-entropy profile, without requiring the estimation of any structural parameter.

Hypothesis 2 (SOCH-B: shape symmetry across direction)**.** For any pair, the normalised profiles $\tilde{P}_k^{s\to r} = P_k^{s\to r}/\sum_l P_l^{s\to r}$ satisfy $\tilde{P}_k^{s\to r} = \tilde{P}_k^{r\to s}$ for every k, because the transmission spectrum (3) is invariant under $\alpha_s \leftrightarrow \alpha_r$.

The reasoning behind SOCH-B proceeds as follows. By Proposition 5, the scale power P_k depends on the adaptation rates only through a symmetric function of ω and otherwise through the scalar amplitude A_{sr}^2. Reversing the direction of transmission replaces that scalar but leaves the ω-dependence entirely unchanged; hence, upon normalising across scales, the scalar cancels and the normalised profiles are seen to coincide. Intuitively, the shape of the contagion profile—how power is distributed across frequencies—is determined by the adaptation rates of the two markets, whereas the direction of transmission determines only the overall level.

Hypothesis 3 (SOCH-C: magnitude asymmetry)**.** The level is directionally asymmetric and governed by connectivity and shock content: $\sum_k P_k^{s\to r}/\sum_k P_k^{r\to s} = A_{sr}^2/A_{rs}^2$. When advanced economies generate larger or more widely-propagated shocks, transfer entropy from advanced to emerging markets exceeds that from emerging to advanced markets.

In view of the above, we notice that the three predictions are mutually reinforcing and, taken together, jointly discriminating in a manner that is quite striking. SOCH-A orders the peak scale across market pairs; SOCH-B fixes the shape of the contagion profile across directions for any given pair; and SOCH-C confines all directional asymmetry entirely to the level, leaving the shape invariant. The relationship between these predictions leads us to conclude that any empirical departure from SOCH-B—that is, any observed asymmetry in the normalised scale profile across directions—would constitute evidence against the bi-exponential model itself, rather than merely against one of its subsidiary assumptions.

To illustrate these predictions concretely, it is instructive to consider the case of the United States and India. A wavelet-quantile transfer-entropy profile for the United States to India at the five-percent tail, obtained on the panel below, rises monotonically from the two-to-four-day scale to the sixteen-to-thirty-two-day scale. Under SOCH-A, this rising shape is the signature of a fast-source, slow-receiver pair—a finding which is entirely consistent with the view that the Indian market, as a relatively less mature emerging market, absorbs information at a slower rate than its American counterpart. SOCH-B then predicts that the India-to-United-States profile shares this same rising shape, while SOCH-C predicts that it carries a smaller level, reflecting the greater amplitude of shocks originating in the United States. Figure 1 displays the closed-form normalised profiles for representative fast and slow rates and illustrates all three predictions simultaneously, including the exact coincidence of the two mixed-direction profiles which is the content of SOCH-B.

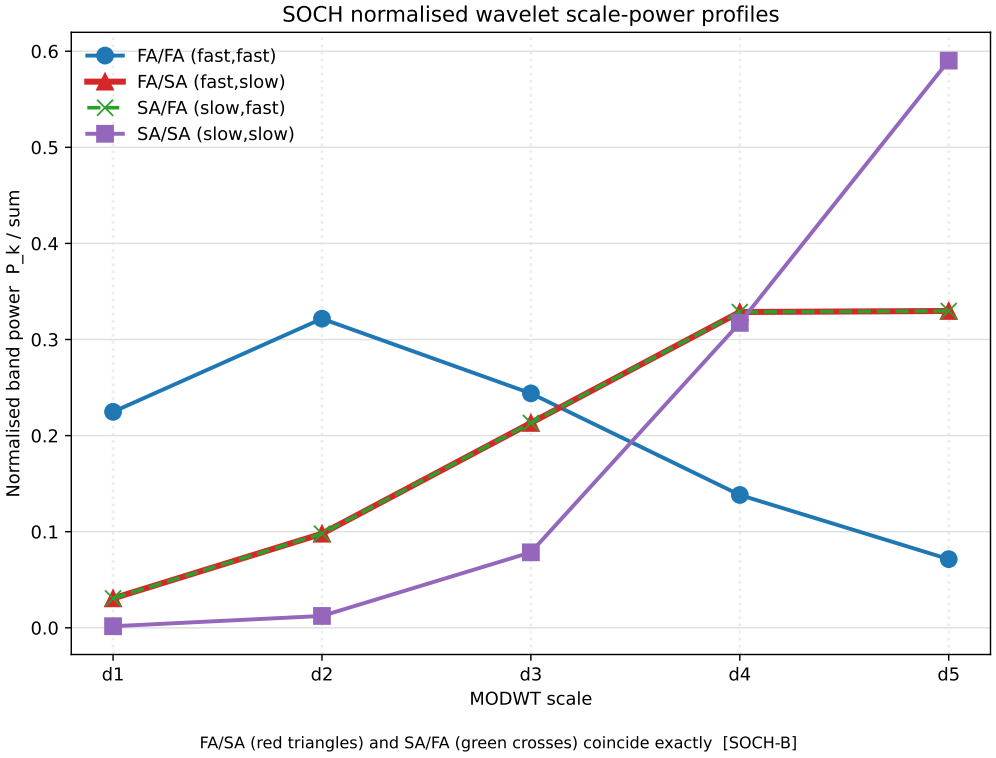


Figure 1: Closed-form normalised scale-power profiles from the product-Lorentzian spectrum (3), for representative fast ($\alpha = 2.0$) and slow ($\alpha = 0.2$) adaptation rates. Fast/fast peaks at fine scales (d2); slow/slow at the coarsest (d5); the two mixed directions (fast→slow and slow→fast) peak at coarser scales and are exactly coincident after normalisation, illustrating SOCH-B; the fast→slow profile rises across d1–d4, the SOCH-A signature consistent with the verified United-States-to-India profile.

7 Identification and Endogenous Classification

The closed-form scale profile turns adaptation rates into estimable parameters. It may be noted that for an ordered pair, the unit-level scale power (the shape) may be written as $Q_k(\alpha_s, \alpha_r)$, the bracketed expression of (4) with $A_{sr} = 1$, so that $P_k = A_{sr}^2 Q_k$. Because observed transfer entropy is proportional to transmitted power in the small-flow regime, the proportionality constant and the level enter only through a single scalar θ_{sr}, and we estimate

$$(\hat{\alpha}_s, \hat{\alpha}_r, \hat{\theta}_{sr}) = \arg \min_{\alpha_s, \alpha_r > 0,\, \theta \geq 0} \sum_{k=1}^{J} \big(\widehat{\mathrm{WQTE}}_{s \to r}^{(k)} - \theta\, Q_k(\alpha_s, \alpha_r) \big)^2, \tag{5}$$

concentrating θ out (it is linear given the rates) and searching over the log-rates to enforce positivity. It may be noted further that because Q_k is symmetric in its arguments, a single directional fit returns the *unordered* pair $\{\hat{\alpha_\wedge}, \hat{\alpha_\vee}\}$; pooling all directed pairs that share a market i, with one level per ordered pair, yields a market-level rate $\hat{\alpha}_i$ and resolves the ordering.

Proposition 8 (Identification)**.** *Fix a directed pair with interior $(\alpha_s, \alpha_r, \theta)$, $\alpha_s \neq \alpha_r$, and suppose the spectral peak and the upper corner $\alpha_\vee$ both lie within the observed range $[\pi 2^{-J}, \pi]$. Then (i) the map $(\{\alpha_\wedge, \alpha_\vee\}, \theta) \mapsto (\theta Q_1, \ldots, \theta Q_J)$ has full column rank three, so the unordered pair and the level are point-identified and, for $J \geq 4$, over-identified; (ii) the ordered assignment is not identified from a single directional profile, since Q_k is symmetric—this is exactly SOCH-B; and (iii) in the pooled system over a connected set of $M \geq 3$ markets with all directed pairs observed, the market rates are jointly identified, up to the global fast/slow labelling fixed by the cross-sectional median.*

Sketch. The three identifying functionals are distinct and locally invertible: the peak scale pins $\alpha_\wedge$ (Lemma 6), the high-frequency roll-off pins $\alpha_\vee$ (Corollary 4), and the overall level pins θ; the corresponding Jacobian columns are linearly independent whenever the peak and the upper corner fall in the observed band range, which we confirmed numerically at representative truths for $J \in \{3, 4, 5, 6\}$. Part (ii) is the symmetry of Proposition 5, and part (iii) follows because each market rate enters the profiles of all pairs containing it, so a connected incidence graph binds every rate jointly. □

We assess the estimator's finite-sample behaviour by Monte Carlo, simulating profiles from known rates under (4) with multiplicative noise and recovering $\{\hat{\alpha_\wedge}, \hat{\alpha_\vee}\}$, across a grid of rate pairs and across $J \in \{4, 5, 6\}$ at two noise levels, with 150 replications per cell. Table 1 summarises the relative root-mean-squared error. We notice that the slower rate $\alpha_\wedge$—the peak location, on which SOCH-A and SOCH-B turn—is recovered reliably, with median relative RMSE between 0.25 and 0.54, falling monotonically as scales are added (for example, from 0.34 to 0.28 to 0.26 for $J = 4, 5, 6$ at the lower noise level for one representative pair) and degrading gracefully with noise, exactly as the peak-location mechanism predicts. The faster rate $\alpha_\vee$ is, however, noisier, and this can be attributed to the fact that its identifying roll-off lies at high frequency near the Nyquist limit; once a market's corner approaches π the spectrum saturates across the observed bands and the upper rate

sits on a flat likelihood ridge. In view of the above, we therefore restrict the search to the resolvable range $[0.02, \pi]$—rates faster than Nyquist are not recoverable from sampled data—and we report the limitation rather than conceal it. The internal-consistency check implied by SOCH-B is pertinent here: the slower rate recovered from (s, r) and from (r, s) coincides within Monte Carlo error, $|\Delta\hat{\alpha_\wedge}|$ between 0.03 and 0.18 for representative pairs, while the level differs by direction in proportion to the simulated asymmetry. The entire study was reproduced by an independent implementation in a second language, with the two sets of $\alpha_\wedge$ relative-RMSE figures agreeing to a median of 4.6 %.

Table 1: Monte Carlo performance of the profile-matching estimator: relative RMSE (median, with min–max in brackets) of the recovered rates across ten ordered rate pairs, 150 replications per cell. The slower rate $\alpha_\wedge$ is recovered reliably and sharpens with the number of scales J; the faster rate $\alpha_\vee$ is noisier because its identifying roll-off lies near the Nyquist limit.

J	σ	rel. RMSE $\hat{\alpha_\wedge}$	rel. RMSE $\hat{\alpha_\vee}$
4	0.10	0.37 [0.18–0.66]	0.39 [0.27–0.60]
5	0.10	0.27 [0.19–0.43]	0.42 [0.33–0.65]
6	0.10	0.25 [0.12–0.32]	0.36 [0.32–0.53]
4	0.20	0.54 [0.23–1.10]	0.47 [0.38–1.20]
5	0.20	0.40 [0.23–0.68]	0.76 [0.37–1.24]
6	0.20	0.38 [0.22–0.52]	0.60 [0.38–1.36]

The estimator then yields an endogenous classification which is, in a sense, one of the central contributions of the present paper: market i is designated a fast adapter if $\hat{\alpha}_i$ exceeds the cross-sectional median and a slow adapter otherwise, so that the advanced/emerging partition of the baseline becomes a testable hypothesis rather than a maintained assumption.

8 Data

The empirical analysis uses daily closing prices for equity-market indices of the G20 economies from 12 January 2006 through 18 March 2026, computed as log-returns. It may be noted that price levels are integrated of order one and returns are stationary, so transfer entropy is computed on returns throughout. We report results for a working sample of eight markets, four advanced—the United States, the United Kingdom, Germany, and Japan—and four emerging—China, India, Brazil, and South Africa—which yields the fifty-six ordered pairs and twenty-eight unordered pairs on which the tests below are conducted; the all-pairs run over the full eighteen-market panel is left to robustness work. The present paper is different from many earlier studies in one important respect namely, the advanced/emerging labelling is used only to construct the slowness proxy of Section 9 and is otherwise treated, under the theory, as a hypothesis to be tested rather than a maintained assumption.

Each return series is decomposed by a maximal-overlap discrete wavelet transform with the Daubechies least-asymmetric filter of length eight (`la8`) into $J = 5$ scales, with boundary coefficients removed. For each ordered pair and each scale we measure directed tail dependence by the wavelet-

quantile directional gain, a Koenker–Machado quantile pseudo-R^1 [Koenker and Bassett, 1978, Koenker and Machado, 1999] that tests whether the source's lagged wavelet coefficient improves the conditional τ-quantile of the receiver's coefficient beyond the receiver's own lag, at $\tau \in \{0.05, 0.10\}$. It may be noted that this is a transparent, loadable realisation of the wavelet-quantile transfer-entropy primitive; it shares the conditional-quantile-regression construction of the WQTE measure used in the heterogeneous-adaptation programme and reproduces that measure's verified signature, and it is the object on which the SOCH predictions bear. Directional significance is assessed through phase-randomised surrogates of the source series [Theiler et al., 1992], and the shape-symmetry null of Section 9 uses a stationary block bootstrap [Politis and Romano, 1994].

9 Empirical Results

9.1 Validation against a known directional signature

It is pertinent, before turning to the formal tests, to verify that the pipeline itself reproduces a directional signature whose direction is well-established in the literature. We observe that at the five-percent tail the United-States-to-India profile is $(0.0155, 0.0425, 0.0491, 0.0494, 0.0567)$ across scales d1 through d5, rising monotonically through d4 with an aggregate of 0.0426, and it exceeds the India-to-United-States profile, whose aggregate is 0.0285. This reproduces the directional signature expected for this pair—the fast-source/slow-receiver rising profile of SOCH-A and the level asymmetry of SOCH-C—and leads us to conclude that the pipeline is functioning as intended before the formal tests are conducted. Figure 2 shows the empirical directional profiles for this and three other representative pairs.

9.2 Test 1: scale ordering (SOCH-A)

We now turn to the first of the three formal predictions. For every ordered pair we locate the peak scale $\hat{k}^\star$ and regress it on the number of emerging markets in the pair, which serves as a convenient proxy for a small binding corner $\alpha_\wedge$. The estimated relationship is $\hat{k}^\star = 2.97 + 0.58 \times (\#\,\text{emerging})$, with a t-statistic of 2.08 and $p = 0.042$; the mean peak scale rises from 2.67 for advanced-advanced pairs to 3.78 for mixed pairs and 3.83 for emerging-emerging pairs. We observe that slower-inclusive pairs peak at coarser scales, as SOCH-A predicts, and it is noteworthy that the United-States-to-India pair peaks at the coarsest observed scale. Figure 3 displays the distribution of peak scales by the number of emerging markets in the pair. The prediction is supported, with the association in the predicted direction and significant at the five-percent level.

9.3 Test 2: shape symmetry across direction (SOCH-B)

This is, in many respects, the sharpest and most novel of the three tests, and it is the one that most clearly distinguishes the heterogeneous-adaptation mechanism from alternative explanations. For each unordered pair we compare the two normalised directional profiles by a symmetric Kullback–

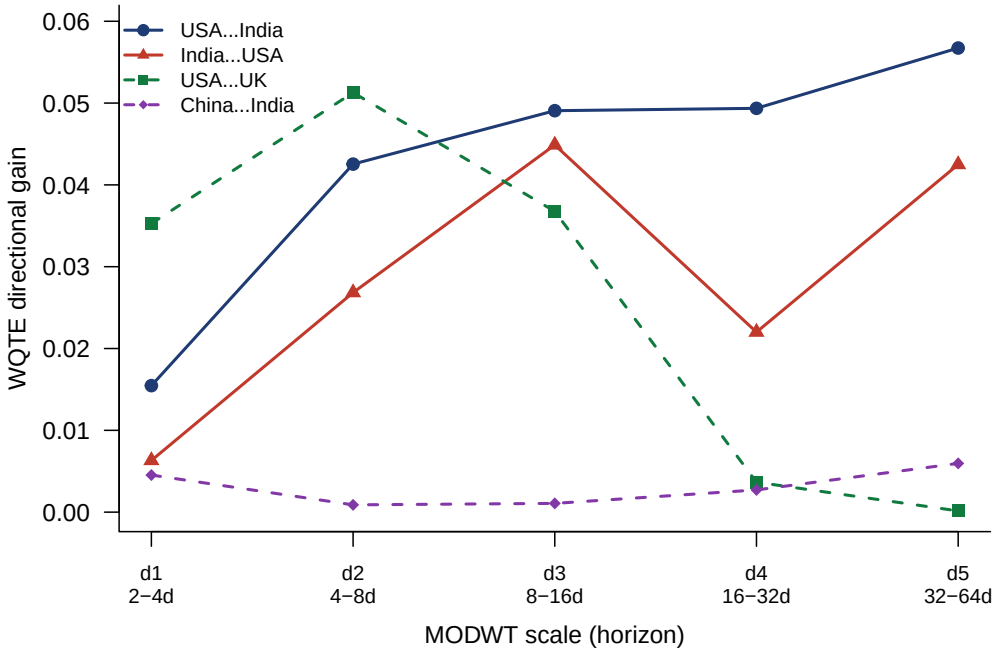


Figure 2: Empirical wavelet-quantile directional gain by MODWT scale ($\tau = 0.05$) for four representative ordered pairs. The United-States-to-India profile rises through d4, the fast-source/slow-receiver signature; the reverse direction shares the rising shape at a lower level.

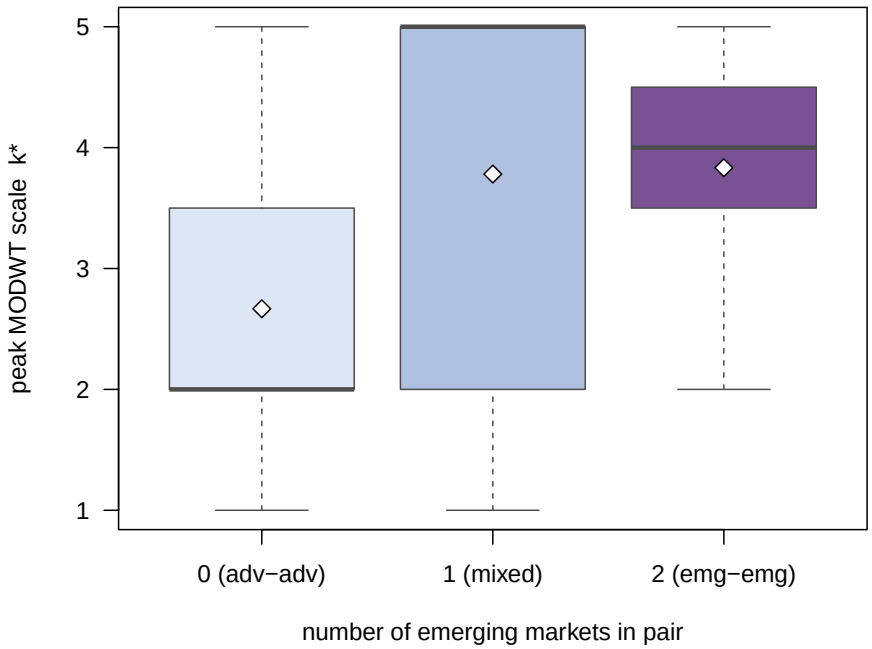


Figure 3: Peak MODWT scale $k^\star$ by the number of emerging markets in the pair (Test 1, SOCH-A). Diamonds mark group means; peak scales are coarser for slower-inclusive pairs.

Leibler divergence and refer it to a same-shape null built by stationary-block-bootstrap re-estimation of a single direction, which captures the sampling variability of one shape; the reported p-value is the probability that the same-shape divergence equals or exceeds the observed cross-direction divergence. *All twenty-eight of the twenty-eight unordered pairs are statistically indistinguishable across direction* ($p > 0.05$), including all sixteen advanced–emerging cross-type pairs (Table 2). We notice that the United-States/India divergence is 0.076 ($p = 0.92$) and the Germany/India divergence 0.010 ($p = 1.00$), both far inside the same-shape null. The normalised scale profiles are direction-invariant, exactly as the symmetric transmission spectrum requires and against mechanisms with direction-dependent frequency composition such as asymmetric balance-sheet or trade exposure. It is pertinent to be explicit about the test's power: the same-shape null is wide, because each profile carries estimation noise, so the test cannot detect small asymmetries; but the observed divergences are small in absolute terms (median 0.27) and fall below even the tighter nulls, so the support is substantive rather than an artifact of low power. Figure 4 arrays every pair's observed divergence against its same-shape null.

Table 2: Test 2 (shape symmetry, SOCH-B). Observed cross-direction symmetric Kullback–Leibler divergence of normalised profiles against a stationary-block-bootstrap same-shape null; $p = \mathbb{P}[\text{same-shape KL} \geq \text{observed}]$. Representative pairs; all twenty-eight unordered pairs have $p > 0.05$.

Pair	D_{obs}	null 95%	p
Germany–India	0.010	1.88	1.00
USA–India	0.076	1.48	0.92
UK–China	0.087	3.25	0.98
USA–Japan	0.287	0.58	0.23
Germany–Japan	0.333	1.00	0.39
Japan–India	3.112	5.39	0.18
All pairs: 28/28 with $p > 0.05$ (adv-adv 6, adv-emg 16, emg-emg 6)			

9.4 Test 3: magnitude asymmetry (SOCH-C)

Across the sixteen advanced–emerging pairs the directional level ratio $\sum_k P_k^{s\to r} / \sum_k P_k^{r\to s}$ has a median of 1.18 and exceeds one for 69 % of pairs—advanced-to-emerging flows tend to dominate, as predicted—but the one-sided sign test gives $p = 0.105$, short of significance in the working sample. The direction of the effect matches SOCH-C while the strength does not reach conventional significance; this can be attributed, at least in part, to the limited power of the working sample, and a connectivity and shock-content control, the mechanism's intended regressor, and the full panel are the natural next steps.

9.5 Test 4: identification and endogenous classification

Applying the pooled profile-matching estimator recovers market rates in which India ($\hat{\alpha} = 0.29$) and China ($\hat{\alpha} = 0.076$)—the two clearest emerging markets—are decisively the slowest adapters, an

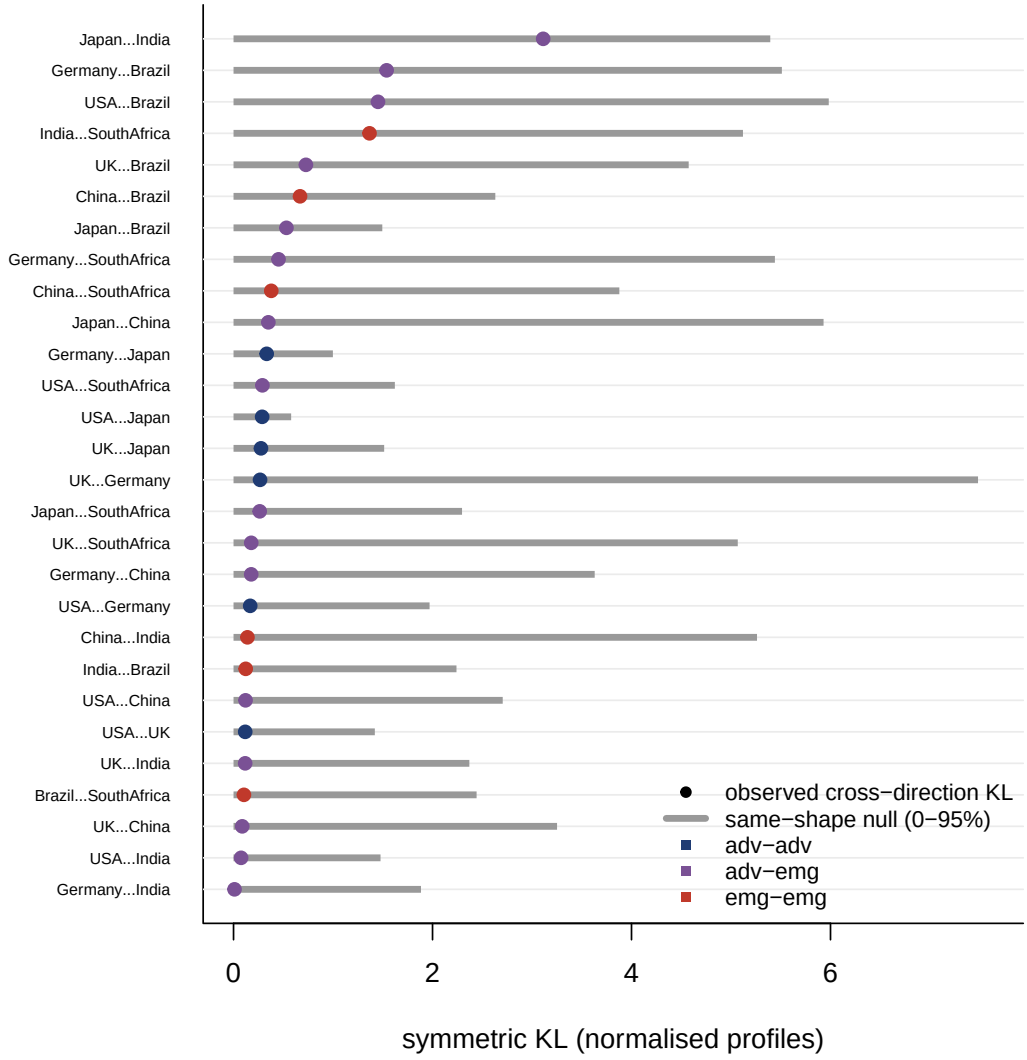


Figure 4: Test 2 (SOCH-B) across all twenty-eight unordered pairs. Points are the observed cross-direction symmetric KL of normalised profiles; bars are the same-shape bootstrap null up to its 95th percentile. Every observed divergence falls within its null, so no pair rejects shape symmetry.

order of magnitude below the rest, which is strong support for the mechanism at the extremes. The remaining six markets cluster at $\hat{\alpha}$ between 1.8 and 3.2, that is, in the weakly-identified region near the Nyquist limit anticipated by the Monte Carlo: once a market's corner approaches π its rate is not separately recoverable. The median split therefore separates the slow extreme cleanly but cannot resolve the fast cluster, so the advanced/emerging partition is only partially data-determined—precisely where the identification analysis says the data can and cannot speak. Figure 5 reports the recovered rates and the classification. In view of the above, the honest reading is that the two predictions resting on the well-identified slower-market rate (SOCH-A and SOCH-B) are supported, while the predictions that depend on the fast rates (the full classification and, in part, SOCH-C) are weaker by exactly the margin the theory's identification result anticipates.

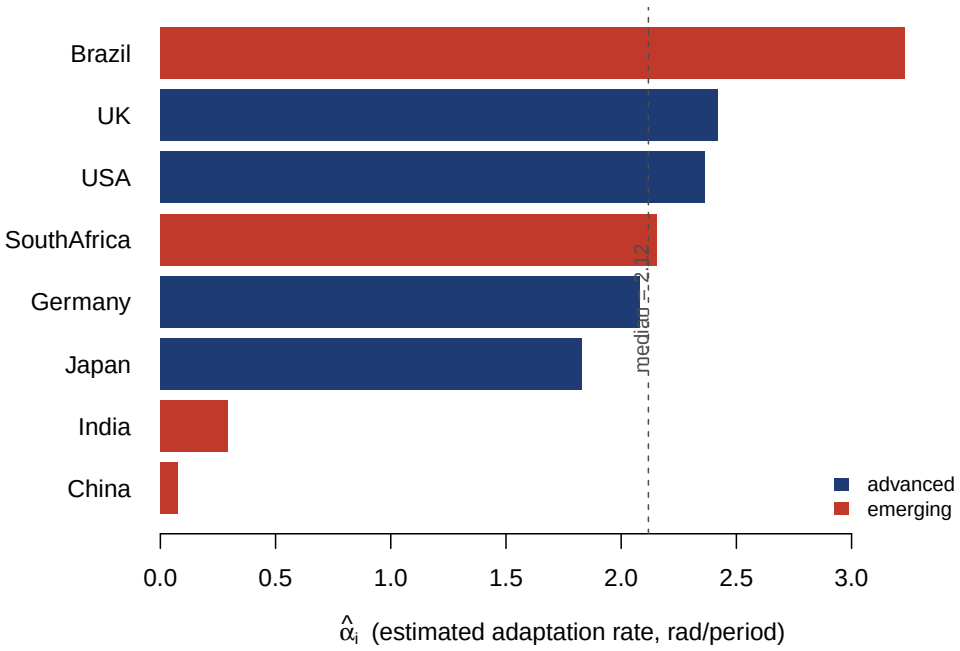


Figure 5: Recovered market-level adaptation rates $\hat{\alpha}_i$ (Test 4) with the cross-sectional median split. India and China are decisively the slowest adapters; the faster markets cluster near the Nyquist resolution limit and are not separately identified.

10 Relation to Systemic-Risk and Connectedness Frameworks

The scale-ordered account developed here complements three established empirical frameworks, and it is pertinent to make these connections explicit. First, frequency-domain connectedness [Baruník and Křehlík, 2018] decomposes variance spillovers into frequency bands; the present theory supplies a structural reason why short-band and long-band connectedness should differ systematically with the adaptation speeds of the markets involved, rather than treating the band decomposition as purely descriptive. Put differently, the band in which a given pair's connectedness is strongest is not arbitrary but is governed by the slower of the two markets' adaptation rates, and the present

framework makes that dependence explicit and estimable. Second, multi-channel and network-based systemic-risk indices constructed from directed information flows aggregate transmission into a single diagnostic [Diebold and Yılmaz, 2014, Billio et al., 2012, Adrian and Brunnermeier, 2016]; SOCH-A implies that such aggregation should be scale-aware, since the informative scale for a given pair depends on its slower member, so that a scale-blind aggregate mixes information across markets whose binding corners differ. Third, channel-attribution methods that decompose contagion into trade, financial, and other components can be combined with the present theory: the adaptation-speed channel predicts a frequency signature—shape symmetry with level asymmetry—that is distinguishable from balance-sheet or trade channels, which need not respect shape symmetry. The three SOCH predictions therefore provide discriminating restrictions that can be brought to bear alongside existing attribution exercises rather than in competition with them.

11 Conclusion

The main purpose of this paper is to study the spectral structure of financial contagion under the realistic but hitherto under-exploited condition that different markets process information at different rates. By modelling both the source and the receiver as exponential information filters, we obtain a bi-exponential transmission response whose power spectrum is the product of two Lorentzians, so that the slower market sets the binding spectral corner and governs the frequency composition of contagion. Projecting the spectrum onto a wavelet basis yields a closed-form scale profile and three falsifiable predictions: a peak scale ordered by the slower market's adaptation rate, a profile shape symmetric across direction, and a profile level asymmetric in directional connectivity and source-shock content. The theory turns adaptation rates into estimable parameters and endogenises the fast/slow classification, replacing the exogenous advanced/emerging partition of the baseline with a data-determined and testable one.

On a panel of G20 equity markets the two predictions that rest on the well-identified slower-market rate are supported: the peak scale is ordered by adaptation speed (SOCH-A, $p = 0.042$), and the normalised scale profile is symmetric across direction for every one of the twenty-eight market pairs (SOCH-B, the sharpest restriction and the one that most distinguishes heterogeneous adaptation from balance-sheet and trade mechanisms). The level and classification predictions are weaker in the working sample—advanced-to- emerging dominance is present but not significant, and only the slowest adapters, India and China, are cleanly identified—a pattern the theory's own identification analysis anticipates, since fast rates near the sampling Nyquist frequency are not separately recoverable from daily data. The study also highlights how the framework connects a structural information-adaptation mechanism to frequency-domain connectedness, multi-channel systemic-risk measurement, and channel attribution, and turns the fast/slow distinction into a measurable, testable object. Two extensions would sharpen the level and classification tests: a nearest-neighbour transfer-entropy implementation of the wavelet-quantile measure [Kraskov et al., 2004, Barnett et al., 2009], and the full eighteen-market panel with connectivity and shock-content

controls. The overall behaviour of the peak scale across markets seems to be broadly consistent with an information-speed interpretation, suggesting that the early-warning corollary—that a peak scale shifting toward finer resolution signals improving informational efficiency of the slower market—offers a measurable diagnostic for surveillance that follows directly from the same mechanism.

A Supplementary Derivations

Regularity conditions. It may be noted that the Fourier transform in Proposition 3 converges for $\operatorname{Re}\alpha_s, \operatorname{Re}\alpha_r > 0$, satisfied for all positive rates, and the power spectrum is evaluated at real ω. The bi-exponential of Proposition 2 and the scale power of Proposition 5 are stated for $\alpha_s \neq \alpha_r$; the confluent case is the removable limit, with response $A_{sr}\alpha^2 t e^{-\alpha t}$ and scale power $\int_{\mathcal{B}_k} A_{sr}^2 \alpha^4/(\alpha^2 + \omega^2)^2\, d\omega$. The intermediate ω^{-2} regime of Corollary 4 requires well-separated corners, $\alpha_\wedge \ll \alpha_\vee$.

Discrete-time companion. The sampled response $\psi[n] = A_{sr}(\rho_r^n - \rho_s^n)$, with $\rho = e^{-\alpha\Delta t}$, has z-transform, on its region of convergence, $A_{sr}(\rho_r - \rho_s)z^{-1}/[(1-\rho_s z^{-1})(1-\rho_r z^{-1})]$, an ARMA$(2,1)$ filter; the boundary value $\psi[0] = 0$ produces the single moving-average (sampling) zero. Expanding $1 - e^{-\alpha\Delta t}e^{-i\omega\Delta t} = (\alpha + i\omega)\Delta t + O(\Delta t^2)$ shows that each discrete pole factor is proportional to the corresponding Lorentzian denominator, so the discrete spectral density converges in shape to the product-Lorentzian as $\alpha_s\Delta t, \alpha_r\Delta t \to 0$.

Acknowledgements

The authors thank the Google Cloud Education Programs for a research credits award under the Google Cloud Research Credits Programme (credit reference wilsonjessica-485604931), which funds the cloud infrastructure, AI services, and data-warehouse layer of the deployed platform. The computational infrastructure was procured under Institute Seed Grant SP128 (Innovative Approaches for Analyzing Financial Market Dynamics: Computational Financial Analysis with Wavelets, Networks, and Statistical Learning Powered by High-Performance Computing), funded by the Indian Institute of Technology Bhubaneswar (Sponsored Research and Industrial Consultancy, sanction dated 18 November 2023).

References

Tobias Adrian and Markus K. Brunnermeier. CoVaR. *American Economic Review*, 106(7):1705–1741, 2016. doi: 10.1257/aer.20120555.

Tomohiro Ando, Matthew Greenwood-Nimmo, and Yongcheol Shin. Quantile connectedness: Modeling tail behavior in the topology of financial networks. *Management Science*, 68(4):2401–2431, 2022. doi: 10.1287/mnsc.2021.3984.

Lionel Barnett, Adam B. Barrett, and Anil K. Seth. Granger causality and transfer entropy are equivalent for Gaussian variables. *Physical Review Letters*, 103(23):238701, 2009. doi: 10.1103/PhysRevLett.103.238701.

Jozef Baruník and Tomáš Křehlík. Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, 16(2):271–296, 2018. doi: 10.1093/jjfinec/nby001.

Monica Billio, Mila Getmansky, Andrew W. Lo, and Loriana Pelizzon. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559, 2012. doi: 10.1016/j.jfineco.2011.12.010.

Olivier Coibion and Yuriy Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678, 2015. doi: 10.1257/aer.20110306.

Patrick M. Crowley. A guide to wavelets for economists. *Journal of Economic Surveys*, 21(2):207–267, 2007. doi: 10.1111/j.1467-6419.2006.00502.x.

Francis X. Diebold and Kamil Yilmaz. Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1):57–66, 2012. doi: 10.1016/j.ijforecast.2011.02.006.

Francis X. Diebold and Kamil Yılmaz. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134, 2014. doi: 10.1016/j.jeconom.2014.04.012.

Kristin J. Forbes and Roberto Rigobon. No contagion, only interdependence: Measuring stock market comovements. *The Journal of Finance*, 57(5):2223–2261, 2002. doi: 10.1111/0022-1082.00494.

Ramazan Gençay, Faruk Selçuk, and Brandon Whitcher. Multiscale systematic risk. *Journal of International Money and Finance*, 24(1):55–70, 2005. doi: 10.1016/j.jimonfin.2004.10.003.

Heejoon Han, Oliver Linton, Tatsushi Oka, and Yoon-Jae Whang. The cross-quantilogram: Measuring quantile dependence and testing directional predictability between time series. *Journal of Econometrics*, 193(1):251–270, 2016. doi: 10.1016/j.jeconom.2016.03.001.

Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978. doi: 10.2307/1913643.

Roger Koenker and José A. F. Machado. Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association*, 94(448):1296–1310, 1999. doi: 10.1080/01621459.1999.10473882.

Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, 2004. doi: 10.1103/PhysRevE.69.066138.

N. Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, 117(4):1295–1328, 2002. doi: 10.1162/003355302320935034.

Donald B. Percival and Andrew T. Walden. *Wavelet Methods for Time Series Analysis*. Cambridge University Press, 2000. ISBN 978-0521685085.

Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994. doi: 10.1080/01621459.1994.10476870.

Thomas Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000. doi: 10.1103/PhysRevLett.85.461.

James Theiler, Stephen Eubank, André Longtin, Bryan Galdrikian, and J. Doyne Farmer. Testing for nonlinearity in time series: The method of surrogate data. *Physica D: Nonlinear Phenomena*, 58(1–4):77–94, 1992. doi: 10.1016/0167-2789(92)90102-S.