Technology High-Yield Debt and Banking Fragility: Regime-Dependent Contagion, Tail Risk, and Stress-Test Miscalibration

Małgorzata Snarska[a,1,*]

[a] *Cracow University of Economics, Cracow, Poland*

Abstract

This paper studies whether stress in technology-sector high-yield (HY) debt has become a distinct financial-stability vulnerability for U.S. regional banks. The question is deliberately framed as a stress-testing and monitoring problem: we do not ask whether technology HY debt is the ultimate causal source of bank stress, but whether the technology-HY–bank link is sufficiently state-dependent that constant-parameter supervisory calibrations average away the relevant risk. Using 2,869 daily observations for regional-bank equity (KRE), a technology-HY total-return index and semiconductor equity over 2015–2025, we combine a static benchmark, a Markov-switching transmission model, DCC-GARCH correlations, regime-conditioned ΔCoVaR/MES and copula tail dependence. The central result is a sharp amplification in refinancing-stress regimes. The technology-HY transmission coefficient rises from 1.03 in tranquil periods to 2.91 in the stress regime, while residual variance is almost seven times larger. The same nonlinearity appears in the left tail: regime-conditioned ΔCoVaR equals −119 basis points in the stress regime, compared with −37 basis points in tranquil periods. Dynamic correlations, contemporaneous regime correlations and tail-dependence estimates support the same interpretation. The evidence is consistent with a common-creditor/fire-sale mechanism in which technology credit repricing and regional-bank equity losses become tightly coupled when intermediary balance-sheet capacity is constrained. The policy implication is direct: stress tests that use average HY-bank correlations understate the banking-sector loss associated with technology refinancing stress. We propose a regime-adaptive calibration in which the technology-HY block enters financial-stability monitoring as a state-contingent multiplier rather than a constant risk weight.

Keywords: technology high-yield debt, regional banks, regime-dependent contagion, fire sales, ΔCoVaR, Markov-switching, stress testing

JEL: G01, G12, G21, C32, C58

---

[*] Corresponding author.

Email address: malgorzata.snarska@uek.krakow.pl (Małgorzata Snarska)

[1] ORCID: 0000-0003-1743-9848.

1. Introduction

Financial-stability risks increasingly originate outside the traditional bank-balance-sheet perimeter but return to banks through securities holdings, credit lines, common investors and market-based funding. Technology-sector high-yield (HY) debt is a useful case in point. During the low-rate decade, software and technology issuers became more visible in leveraged finance as private-equity sponsorship, recurring-revenue collateralisation and the search for yield expanded the sector's debt capacity. The resulting refinancing calendar is not only a corporate-finance issue. When technology HY returns fall or technology credit spreads widen relative to technology equity, regional banks may be affected through mark-to-market losses on securities portfolios, drawdowns on revolvers extended to leveraged borrowers, and common-creditor sales by funds that hold both credit and financial equities. The relevant stability question is therefore whether technology HY stress behaves like an ordinary component of aggregate credit risk or whether it produces a state-dependent amplification channel for banks.

The mechanisms considered in this paper are inherently nonlinear. When intermediary balance sheets are unconstrained, a technology HY shock can be absorbed by dealers, funds and long-horizon investors with limited price impact. When balance-sheet capacity is scarce, the same shock is transmitted through fire sales, collateral repricing, redemption pressure and funding constraints. In that state, the marginal seller of technology HY exposure may also be reducing financial-equity exposure, so a credit repricing becomes a banking-equity event. Intermediary asset-pricing and fire-sale models therefore imply that the technology-HY coefficient in a bank-return equation should not be treated as a fixed number. It should be moderate in tranquil periods and materially larger in refinancing-stress periods.

We test this prediction using daily data from 2015 to 2025 for KRE, a technology-HY total-return index and semiconductor equity. The empirical strategy is cumulative. An OLS benchmark establishes the average contemporaneous transmission from technology HY returns to regional-bank returns, controlling for lagged bank returns and semiconductor equity. A Markov-switching model then allows both the mean transmission coefficient and residual variance to change across latent regimes. DCC-GARCH correlations, regime-conditioned ΔCoVaR/MES and copula tail dependence test whether the same state dependence appears in co-movement and in the lower tail. This triangulation is important for a financial-stability journal: a result that appears simultaneously in conditional means, dynamic correlations and conditional quantiles is more relevant for stress-test calibration than a single coefficient estimated in isolation.

The evidence identifies a refinancing-stress regime covering 18.5% of the sample. In that regime, the technology-HY coefficient equals 2.91, compared with 1.03 in the tranquil regime, implying a 2.82$\times$ amplification. Residual variance is 6.97 times larger in the stress regime. The tail-risk estimates deliver the same message: ΔCoVaR equals -119 basis points in stress periods and -37 basis points in tranquil periods, an amplification ratio of 3.18. DCC-

GARCH correlations rise during the same episodes and a t-copula estimate implies non-zero lower-tail dependence. These results are consistent with a balance-sheet interpretation in which technology HY stress becomes systemically relevant only when refinancing pressure and intermediary constraints coincide.

The paper contributes to the financial-stability literature in three ways. First, it identifies technology-sector HY debt as a sector-specific credit block whose banking-sector relevance is regime-dependent rather than constant. Second, it shows that the amplification appears not only in conditional means but also in regime-conditioned systemic-risk measures, especially ΔCoVaR. Third, it translates the estimates into a stress-testing implication: a calibration based on full-sample or tranquil correlations understates losses in the very regime stress tests are designed to capture. The article is intentionally limited to regime-dependent transmission and stress-test calibration. It does not claim to identify the ultimate causal ordering of technology HY and bank returns, nor does it decompose the channel into a separate risk-appetite mechanism. Those questions require different empirical designs and are left to companion analyses. Here, the object of interest is narrower and directly macroprudential: whether a constant-parameter representation of the technology-HY–bank link is adequate for financial-stability monitoring.

The remainder of the paper is organised as follows. Section 2 combines the economic mechanism, literature review and hypothesis development. Section 3 describes the data and empirical strategy. Section 4 reports the main evidence. Section 5 presents additional robustness and validation results. Section 6 discusses stress-testing and macroprudential implications. Section 7 concludes.

2. Economic Mechanism, Literature Review and Hypothesis Development

2.1. Technology HY debt, bank balance sheets and direct transmission

Technology issuers historically relied more heavily on equity finance than traditional cyclical borrowers. Since the mid-2010s, however, three forces have increased their relevance for leveraged credit markets. Private-equity-sponsored software transactions have made leveraged loans and HY bonds a standard financing tool for technology assets. Subscription and recurring-revenue business models have made some technology cash flows more legible to credit investors. The reach-for-yield environment documented by Becker and Ivashina (2015) pushed institutional investors toward lower-rated credit instruments. The resulting technology-HY exposure is important for financial stability not because each issuer is systemically large, but because the sector is embedded in the same market-based intermediation chain that connects banks, CLOs, HY funds, insurance accounts and multi-asset investors.

This direct-transmission argument is grounded in three strands of literature. The first is the credit-spread channel, where corporate bond market conditions affect financing, investment and the valuation of bank assets (Gilchrist and Zakrajšek, 2012). The second is the literature on bank liquidity provision, in which corporate credit lines are valuable precisely because borrowers can draw them when market finance becomes unavailable (Holmström and Tirole, 1998). The third is the post-crisis systemic-risk literature, which shows that market-based measures such as ΔCoVaR, MES and SRISK are informative because they use prices to reveal forward-looking interdependence among institutions and sectors (Adrian and Brunnermeier, 2016; Acharya et al., 2017; Brownlees and Engle, 2017). In this paper, the technology-HY index is interpreted as a sectoral credit block whose deterioration can affect regional banks through securities values, credit-line conversion and investor rebalancing.

Regional banks are a particularly relevant equity proxy for this channel. They are less diversified than money-center institutions, more sensitive to funding and deposit-franchise repricing, and more exposed to local or sectoral balance-sheet narratives during stress. KRE therefore captures a part of the banking system for which credit-market repricing can become an equity-market event. Semiconductor equity returns are included because a technology-credit shock must be separated from broad technology-equity co-movement. The first empirical implication is that, after controlling for semiconductor equity and bank-return persistence, technology HY debt returns should still carry information for regional-bank returns.

> ***H1 (Direct transmission).*** *Technology HY debt return shocks transmit to regional-bank equity returns with a statistically and economically significant coefficient after controlling for semiconductor equity returns and lagged bank returns. In return notation, negative technology HY shocks should coincide with negative regional-bank return shocks.*

H1 is deliberately modest. It is not a causal identification claim and it does not require that technology HY debt be the only source of regional-bank stress. It is the necessary first condition for the financial-stability argument: if technology HY carries no incremental information for regional-bank returns even on average, there is no reason to build a regime-dependent stress-testing module around it.

2.2. Fire sales, common creditors and regime-dependent amplification

The central claim of the paper is not that the average technology-HY coefficient is positive. It is that the coefficient is a state-contingent object. Fire-sale models show why constrained sellers can depress prices below values justified by fundamentals (Shleifer and Vishny, 1992, 2011). The common-creditor mechanism of Coval and Stafford (2007) explains how selling pressure can move across assets that are not linked by cash flows but are held by the same investors. Work on corporate bond mutual funds and asset fire sales reinforces this point for credit markets, where redemption pressure can turn portfolio liquidity management into price pressure on bonds held by other investors (Choi et al., 2020). Intermediary asset-pricing models provide the state variable: when intermediary capital is abundant, shocks are absorbed; when capital is scarce, shocks are amplified through leverage, margins and risk-bearing constraints (Adrian and Shin, 2010; He and Krishnamurthy, 2013; Brunnermeier and Pedersen, 2009).

This literature implies that a constant-parameter model is not the natural null for financial-stability analysis. Hamilton (1989) introduced Markov-switching models precisely to capture discrete changes in macro-financial dynamics, and subsequent work shows that asset correlations, volatility and risk premia often differ sharply across regimes (Pelletier, 2006; Guidolin and Timmermann, 2007). In banking and sovereign-risk applications, regime-switching frameworks are useful because they allow the same exposure to have stabilising or destabilising implications depending on the volatility state of the financial system. For stress testing, this matters because supervisors care less about the average coefficient and more about the coefficient that applies when refinancing markets are closed, risk-bearing capacity is scarce and the financial sector is already fragile.

Systemic-risk measurement provides the second part of the same hypothesis. ΔCoVaR translates a conditional dependence into the change in the left tail of the system when the conditioning institution or sector is distressed (Adrian and Brunnermeier, 2016). MES asks how much the bank or system loses on average when the conditioning asset is in its own tail (Acharya et al., 2017). Copula-based systemic-risk models emphasise that linear correlation can understate dependence when joint crashes are more frequent than Gaussian assumptions imply (Karimalis and Nomikos, 2018; Reboredo and Ugolini, 2016). If the technology-HY channel is truly a financial-stability channel, regime dependence should therefore appear in four places at once: the conditional mean coefficient, residual volatility, dynamic correlations and tail-risk measures. A result confined to only one of these objects would be a model-specific statistical feature; a result visible across all of them is an economically interpretable amplification mechanism.

> ***H2 (Regime-dependent financial-stability amplification).*** *The technology-HY–bank transmission channel is materially stronger in refinancing-stress regimes than in tranquil periods. This amplification should appear jointly in the conditional mean coefficient, residual volatility, dynamic correlations and tail-risk measures such as ΔCoVaR and MES.*

H2 is the core hypothesis of the article. It connects the balance-sheet mechanism to the empirical design and to the policy question. The same technology-HY shock is expected to have limited systemic content in normal times but substantial systemic content when intermediaries are constrained and common investors sell credit and bank equity together. This is also the point at which Paper 1 is separated from the companion analyses: the present article tests whether the regime multiplier exists and whether it matters for stress-test calibration, not whether the direction of causality is fully identified or whether the risk-appetite channel can be isolated structurally.

3. Data and Empirical Strategy

3.1. Data construction and variable definitions

The empirical design starts from a deliberately compact daily dataset. The baseline sample runs from 2 January 2015 to 31 December 2025 and contains $T = 2{,}869$ aligned trading-day observations after constructing returns and dropping non-overlapping dates. The daily frequency is appropriate for the question addressed in this paper because the hypothesised channel is a market-based financial-stability channel. Technology credit losses, ETF repricing, funding news and bank-equity reactions are incorporated into prices quickly, often within the same trading day. A lower-frequency design would be useful for balance-sheet exposure measurement, but it would obscure the short-lived stress episodes that the Markov-switching model is designed to identify.

The dependent variable is the daily log return on KRE, the SPDR S&P Regional Banking ETF. KRE is not used as a literal balance-sheet measure of regional-bank exposure to technology borrowers. It is used as a traded, forward-looking proxy for the equity value of the regional-bank sector. This distinction is important. The paper studies whether market participants price technology-HY refinancing stress as a regional-bank fragility signal in some states of the world. Regional banks are the appropriate listed-bank proxy because their equity valuations are particularly sensitive to securities-book losses, deposit-franchise repricing and funding narratives. Large money-center banks are more diversified and appear later in the paper only as robustness comparators through broader financial-sector and bank ETF controls.

The main explanatory variable is the daily log return on a technology-sector high-yield bond total-return index, denoted $r_{TD,t}$. Since the index is a total-return credit instrument, negative realizations capture credit-market losses generated by spread widening, downgrading pressure, refinancing stress or default-risk repricing. The scale of this variable is smaller than that of bank or semiconductor equity returns, but its tail is much more pronounced. This feature is economically meaningful rather than a purely statistical inconvenience: technology-HY debt is normally quiet, yet it can reprice abruptly when refinancing windows close. The empirical interpretation of a positive coefficient on $r_{TD,t}$ is therefore straightforward. Positive technology-HY returns coincide with stronger regional-bank equity returns, while negative technology-HY shocks coincide with weaker regional-bank equity returns.

Semiconductor equity returns, proxied by the VanEck Semiconductor ETF (SMH), are included to isolate the credit component of the technology shock. Without this control, the technology-HY coefficient could partly reflect broad technology-equity news, risk-on/risk-off conditions in growth stocks or semiconductor-sector repricing. SMH is not intended to exhaust all technology-sector information. Its role is narrower: it removes a high-frequency, liquid technology-equity factor so that the remaining technology-HY coefficient is more closely related to debt-market stress than to generic technology-sector equity co-movement.

All traded price series are converted into continuously compounded returns, $r_{i,t} = \ln(P_{i,t}/P_{i,t-1})$, after aligning trading dates. The baseline variables are therefore measured on a common daily return scale. The empirical files also contain two technology-stress transformations used only for validation. The first is $Spread\_Tech_t = r_{TD,t} - r_{SMH,t}$, an equity-neutral relative credit-stress proxy. It becomes more negative when technology credit underperforms technology equity and more positive when technology credit is relatively resilient. The second is the sign-adjusted and normalized stress index $SCSI\_norm$, built from the same relative credit-equity component and scaled so that higher values correspond to stronger technology-credit stress. These transformations are not substitutes for the main technology-HY return series. They are included to verify whether the conclusions depend on the sign convention, on using a credit-equity relative measure, or on treating stress as an index rather than a raw return.

A separate macro-financial control dataset is used in robustness checks rather than in the baseline specification. It begins on 16 May 2023 and contains 687 observations because some controls are available only over the later part of the sample. The controls include broad financial-sector equity returns (XLF), community-bank returns (KBWB), changes in VIX and MOVE, changes in aggregate HY option-adjusted spreads, changes in the 2-year and 10-year Treasury yields, the 10Y–2Y yield-curve slope and changes in the St. Louis Fed Financial Stress Index. This dataset is useful because it asks whether the tail-risk result survives after broad equity, credit, volatility and rate conditions are partialled out. It is not used as the main sample because doing so would remove COVID-19, the early post-pandemic reopening period and most of the rate-rise transition from the estimation window. The baseline model is therefore estimated on the full 2015–2025 sample, while the shorter control sample serves as a conservative validation layer.

This organisation of the data follows the logic of the paper. The full sample is needed to identify latent regimes and to observe several distinct stress episodes. The compact baseline specification avoids over-conditioning the regime model on variables that would shorten the sample and absorb part of the stress state itself. The richer control dataset is then used to check whether the main ordering survives when the technology-HY shock competes with broad credit, rate and volatility variables. In this sense, the data architecture is sequential rather than redundant: each block answers a different question.

Table 2 reports distributional statistics and diagnostic tests. The descriptive evidence already contains the main empirical tension of the paper. KRE and SMH have similar unconditional daily volatility, with standard deviations of 1.94%. Technology-HY returns have a much smaller standard deviation of 0.24%, but their excess kurtosis is 22.921 and skewness is negative. The technology-HY series is therefore not a high-volatility equity-like series. It is a low-volatility credit series with a thick left tail. This is exactly the type of object that can be underweighted by a stress-test calibration based on average correlations or normal approximations. In tranquil periods the series may appear benign; in refinancing-stress periods the same series can contain information about abrupt credit repricing.

The relative stress proxy, $Spread\_Tech$, has a standard deviation close to the equity vari-

Table 1: Data architecture and empirical role of each variable block

Block	Variables	Transformation and interpretation	Empirical role
Regional banks	KRE price, r_{KRE}	Daily log return; forward-looking regional-bank equity proxy	Baseline outcome and systemic-risk target
Technology credit	TechDebt price, r_{TD}	Daily log return on technology-HY total-return index; negative values indicate credit losses	Main stress/transmission variable
Technology equity	SMH price, r_{SMH}	Daily log return on semiconductor equity ETF	Control for technology-equity co-movement
Relative tech-credit stress	$Spread\_Tech$, $SCSI\_norm$	Relative credit-equity component and normalized sign-adjusted stress index	Alternative stress proxies in robustness
Macro-financial controls	XLF, KBWB, dVIX, dMOVE, dHY_OAS, dDGS2, dDGS10, 10Y–2Y, dSTLFSI4	Broad risk, volatility, rates, credit and financial-sector controls; shorter 2023–2025 sample	Conservative validation, not baseline identification

The baseline sample contains 2,869 daily observations. The control sample contains 687 observations and is used only for robustness because imposing it on the baseline would remove earlier stress episodes.

ables because it combines the technology-HY return with semiconductor equity movements. It should therefore be read differently from $r_{TD,t}$. The baseline TechDebt return is the cleaner credit-market variable, while $Spread\_Tech$ is an aggressive stress transformation that captures relative credit-equity deterioration. This distinction explains why the alternative-proxy robustness tests later in the paper are interpreted as stress-envelope checks rather than as the preferred baseline estimates.

The stationarity and diagnostic tests support a return-based modelling strategy but reject simple Gaussian constant-variance inference. ADF and Zivot–Andrews tests reject unit roots in returns, while KPSS does not reject stationarity. Jarque–Bera, ARCH-LM and Ljung–Box tests reject normality, conditional homoskedasticity and serial independence. These results motivate heteroskedasticity-robust standard errors in the benchmark regression, regime-specific residual variances in the Markov-switching model and GARCH-based validation in the dynamic-correlation layer.

Bai–Perron sequential tests select four breakpoints: 27 May 2019, 7 July 2020, 29 November 2021 and 19 May 2023. The dates line up with economically recognisable changes in the market environment: the late-cycle pre-COVID period, the COVID shock and policy response, the shift toward the post-pandemic rate-rise/refinancing cycle and the spring 2023 regional-bank turmoil. The break evidence is not used as a mechanical crisis classification. Its function is more limited and more important for model design: it shows that a single coefficient is

Table 2: Descriptive statistics and diagnostics, daily log returns (2015–2025)

	Ret_KRE	Ret_SMH	Ret_TechDebt	Spread_Tech
Mean	0.0002	0.0009	−0.0000	−0.0009
Std. dev.	0.0194	0.0194	0.0024	0.0187
Skewness	−0.285	−0.264	−1.132	0.090
Excess kurtosis	8.604	5.546	22.921	5.433
Minimum	−0.168	−0.156	−0.026	−0.165
Maximum	0.143	0.158	0.023	0.134
ADF p	<0.01	<0.01	<0.01	<0.01
KPSS p	>0.10	>0.10	>0.10	>0.10
ARCH-LM p	<0.001	<0.001	<0.001	<0.001
JB p	<0.001	<0.001	<0.001	<0.001

n = 2,869. ARCH-LM uses 10 lags. JB denotes the Jarque–Bera normality test.

unlikely to describe the full sample adequately. Figure 1 provides the corresponding model-selection evidence. The minimum BIC occurs at four breaks, which supports a regime-explicit specification before any Markov-switching estimates are examined.

Table 3: Bai–Perron structural-break diagnostics

Number of breaks	0	1	2	3	4	5
BIC	−15340.30	−15403.63	−15435.38	−15426.71	**−15436.72**	−15428.02
Selected break dates	2019-05-27; 2020-07-07; 2021-11-29; 2023-05-19					

Lower BIC indicates a preferred specification. The four-break model is selected.

3.2. *Empirical strategy*

The empirical strategy follows the two hypotheses. To test H1, we estimate the static benchmark

$$r_{KRE,t} = \alpha + \beta_1 r_{KRE,t-1} + \beta_2 r_{TD,t} + \beta_3 r_{SMH,t} + \varepsilon_t, \tag{1}$$

where β_2 is the average technology-HY transmission coefficient. The inclusion of $r_{SMH,t}$ is important because it prevents the estimated technology-credit coefficient from simply reproducing the contemporaneous technology-equity factor. HAC standard errors are used because the diagnostics indicate non-normality, serial dependence and conditional heteroskedasticity.

To test H2, we estimate the Markov-switching extension

$$r_{KRE,t} = \alpha^{(S_t)} + \beta_1^{(S_t)} r_{KRE,t-1} + \beta_2^{(S_t)} r_{TD,t} + \beta_3^{(S_t)} r_{SMH,t} + \varepsilon_t^{(S_t)}, \tag{2}$$

where $S_t \in \{1, \ldots, k\}$ follows a first-order Markov chain with transition probabilities $P_{ij} = \Pr(S_t = j \mid S_{t-1} = i)$. All regression coefficients and the residual variance are allowed to switch across regimes. This unrestricted formulation is chosen because the economic mechanism

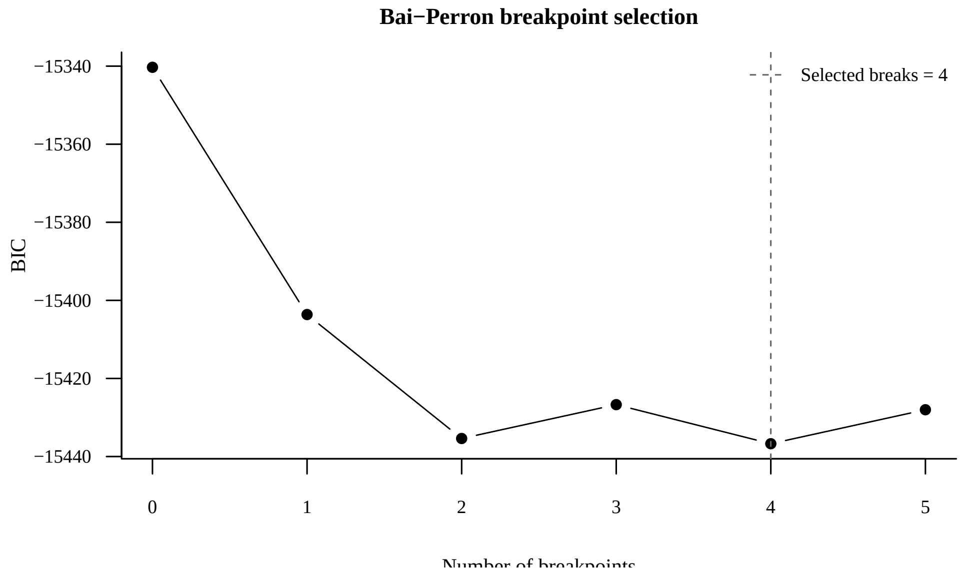

Figure 1: Bai–Perron structural-break model selection. The minimum BIC occurs at four breaks, supporting the use of a regime-explicit specification rather than a single constant coefficient.

predicts both a change in the slope and a change in volatility when balance-sheet constraints bind. The number of regimes is selected by comparing two- and three-state specifications; the decoded regime path is then used for regime-conditioned ΔCoVaR, MES, DCC and robustness tests.

The third layer of the design asks whether the same amplification is visible outside the conditional mean. DCC(1,1)-GJR-GARCH(1,1) is estimated for the trivariate system $(r_{KRE}, r_{TD}, r_{SMH})$ to obtain dynamic conditional correlations. Regime-conditioned ΔCoVaR is estimated by quantile regression following Adrian and Brunnermeier (2016); MES is computed as the mean KRE loss conditional on the technology-HY variable being in its own lower tail following Acharya et al. (2017). Finally, t- and Clayton-copula specifications are fitted to the KRE–TechDebt pair to examine lower-tail dependence. These methods are not presented as competing alternatives. They are a triangulation device: the paper's financial-stability claim is credible only if the regime multiplier is visible in several objects that supervisors actually monitor.

4. Empirical Results

4.1. Benchmark transmission and structural instability

Table 4 reports the OLS benchmark. The technology-HY coefficient is 1.915 ($t = 13.40$, $p < 0.001$), supporting H1. A one-standard-deviation movement in technology HY returns is associated with an approximately 0.46% contemporaneous movement in KRE, or about 24% of KRE's daily standard deviation. This is an economically meaningful average association, but it is not yet the main financial-stability result. The structural-break evidence in Section 3 shows that the average coefficient is unlikely to be stable across stress and tranquil conditions.

Table 4: OLS benchmark: direct technology-HY transmission to regional banks

Variable	Estimate	Std. Error	t-stat	p-value
Intercept	−0.0001	0.0003	−0.41	0.683
$r_{KRE,t-1}$	−0.0613	0.0165	−3.71	<0.001***
$r_{TD,t}$	1.9153	0.1429	13.40	<0.001***
$r_{SMH,t}$	0.3837	0.0173	22.12	<0.001***
R^2	0.270			
$F(3, 2864)$	352.9			<0.001

HAC standard errors. $^{***}p < 0.001$.

4.2. *Regime-dependent transmission*

The two-regime Markov-switching model is the main test of H2 because it allows the technology-HY slope and the residual variance of regional-bank returns to change with the latent state. Table 5 reports the estimates. The stress-regime coefficient on technology HY returns equals 2.910, compared with 1.033 in the tranquil regime. The implied amplification ratio is 2.82. This number has a direct economic interpretation. Using the unconditional technology-HY standard deviation from Table 2, a one-standard-deviation technology-HY movement corresponds to an approximately 0.70 percentage point contemporaneous KRE movement in the stress regime and an approximately 0.25 percentage point movement in the tranquil regime. Relative to KRE's unconditional daily volatility, the same technology-credit shock therefore has a materially larger banking-sector footprint when the market is in the refinancing-stress state.

The estimated variance switch is equally important. The residual standard deviation rises from 1.09% in Regime 2 to 2.88% in Regime 1, implying a 6.97-fold increase in residual variance. The model is therefore not simply labelling days with a larger slope; it is identifying a state in which the bank-return equation is both more volatile and more tightly linked to technology credit. This joint movement in slope and variance is precisely what the balance-sheet mechanism predicts. When intermediaries are unconstrained, technology-HY repricing can be absorbed with limited spillover to bank equity. When refinancing pressure, redemption pressure and funding concerns coincide, the same technology-credit movement is priced in an environment with lower risk-bearing capacity and higher cross-asset selling pressure.

Regime 1 covers 530 observations, or 18.5% of the sample. Its decoded mean duration is 22.5 trading days. This duration is long enough to be relevant for supervisors and risk managers, but short enough to be interpreted as a stress state rather than a new unconditional normal. The transition probability $P_{11} = 0.894$ indicates persistence after entry into stress, while $P_{22} = 0.972$ shows that tranquil periods are more persistent. The asymmetry is economically plausible. Refinancing-stress episodes arrive intermittently, cluster around market events and then dissipate; tranquil regimes dominate the calendar but are interrupted by episodes during which credit and bank equity become more tightly coupled.

Table 5: Two-regime Markov-switching transmission model

	Regime 1: stress		Regime 2: tranquil	
Variable	Coeff.	p	Coeff.	p
$r_{TD,t}$	**2.910**	<0.001***	**1.033**	<0.001***
$r_{SMH,t}$	0.454	<0.001***	0.353	<0.001***
$r_{KRE,t-1}$	−0.048	0.205	−0.061	0.002**
$\hat{\sigma}$	0.0288		0.0109	
n (obs.)	530 (18.5%)		2338 (81.5%)	
Mean duration	22.5 days		35.6 days	
P_{kk}	0.894		0.972	
Amplification ratio: $\hat{\beta}^{(1)}/\hat{\beta}^{(2)} = 2.82$.				
AIC = −16,199; BIC = −16,088; log$L$ = 8,107.6.				

$^{**}p < 0.01$; $^{***}p < 0.001$. Regime-1 residual variance is 6.97× larger than Regime-2 variance.

Table 6: Economic magnitude of the regime switch

Quantity	Regime 1: stress	Regime 2: tranquil	Interpretation
Technology-HY coefficient	2.910	1.033	Stress-state slope is 2.82 times larger
Residual standard deviation	2.88%	1.09%	Unexplained bank-return volatility rises sharply
Residual variance ratio	6.97		Stress state combines stronger slope and higher volatility
Implied KRE effect of one s.d. TechDebt shock	0.70 pp	0.25 pp	Same credit shock has larger bank-equity footprint in stress
Share of observations	18.5%	81.5%	Stress regime is episodic, not the unconditional norm
Decoded mean duration	22.5 days	35.6 days	Stress episodes are short-lived but persistent enough for monitoring

The implied KRE effect multiplies the regime-specific technology-HY coefficient by the unconditional standard devi
Values are shown in percentage points.

The semiconductor-equity coefficient also rises in the stress regime, from 0.353 to 0.454. This increase is smaller than the technology-HY amplification, but it is useful for interpretation. Stress regimes do not isolate a pure credit effect in a mechanical sense; they are periods when technology equity, technology credit and bank equity are more tightly priced together. The larger increase in the technology-HY slope indicates that the credit component carries incremental information beyond the equity-sector control precisely in the state where refinancing constraints should matter most. The lagged KRE coefficient is significant only in the tranquil regime, which is consistent with mild return reversal or adjustment in normal markets. In stress, contemporaneous repricing dominates the own-lag term.

Figures 2 and 3 translate the coefficient estimates into calendar time. The smoothed probabilities show that Regime 1 clusters around COVID-19, the 2021–2022 rate-rise cycle and the spring 2023 regional-bank crisis. The model therefore does not create a stress state from isolated outliers. It groups together periods in which refinancing conditions, bank-equity volatility and credit-market repricing are jointly elevated. The dominant-regime plot confirms that large bank-equity moves occur disproportionately during Regime 1 periods. This visual evidence is not a substitute for the coefficient estimates, but it helps validate the economic labelling of the latent state as refinancing stress rather than generic noise.

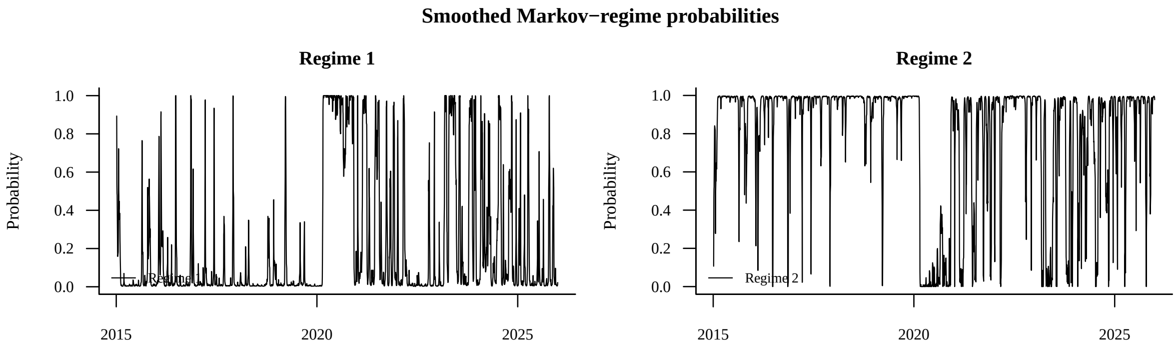

Figure 2: Smoothed Markov-regime probabilities, 2015–2025. Regime 1 clusters around major refinancing and banking-stress episodes.

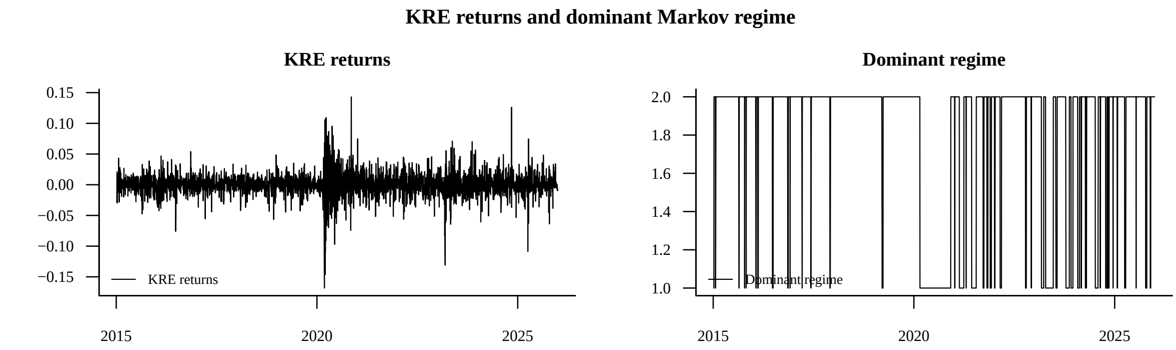

Figure 3: KRE daily returns with dominant Markov regime overlay. Large clustered bank-equity moves coincide with Regime 1 periods.

The interpretation should remain disciplined. The Markov-switching model identifies a state in which the technology-HY–bank association is much stronger; it does not by itself establish that technology-HY debt is the unique causal origin of each bank-equity movement. The result is nevertheless directly relevant for financial stability because stress tests do not require a single-cause narrative. They require a calibration that applies when vulnerabilities are active. Table 5 and Table 6 show that a constant coefficient would average together two materially different pricing environments and would therefore understate the bank-equity sensitivity that applies in the stress state.

4.3. Dynamic correlations, tail risk and co-crash dependence

Figure 4 reports DCC-GARCH dynamic conditional correlations. The KRE–TechDebt correlation rises during stress periods and shows a persistent increase after 2020. The average pairwise correlation is 0.344 in 2023–2024 and 0.390 in 2025, with the difference statistically significant at conventional levels. This supports the financial-stability interpretation: the relationship is not simply a one-off COVID or SVB event, but a channel whose strength appears to grow as technology leveraged finance becomes more embedded in broader financial intermediation.

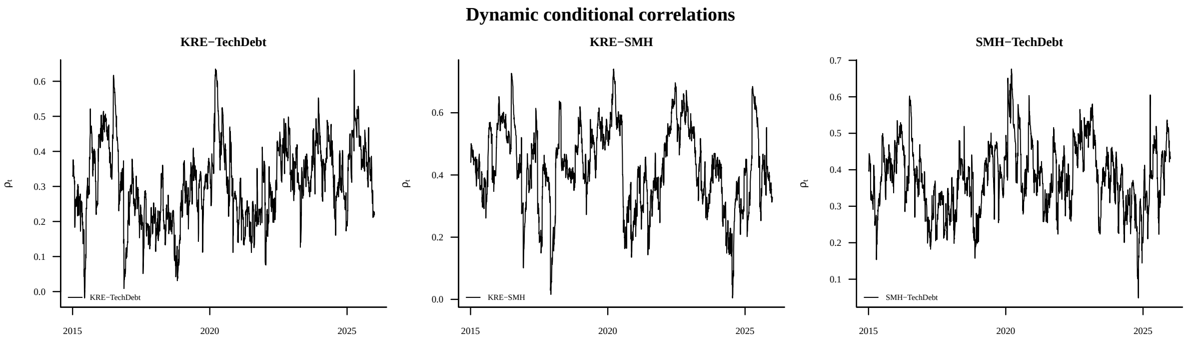

Figure 4: DCC-GARCH dynamic conditional correlations for KRE–TechDebt, KRE–SMH and SMH–TechDebt.

Table 7 reports the regime-conditioned systemic-risk estimates. The baseline full-sample ΔCoVaR is −82 basis points. Conditioning on the Markov regimes reveals that the average conceals a much larger stress-state contribution: −119 basis points in Regime 1 against −37 basis points in Regime 2. The amplification ratio of 3.18 is close to the 2.82 amplification in the Markov-switching transmission coefficient. MES tells the same qualitative story. The average KRE loss conditional on technology-HY stress days is materially more negative than on normal days, and the difference is much larger in the stress regime than in the tranquil regime.

Table 7: ΔCoVaR and MES by Markov regime ($\tau = 0.05$)

Metric	Full sample	Regime 1 (stress)	Regime 2 (tranquil)
VaR_{KRE}	−2.89%	−5.62%	−2.15%
CoVaR (stress)	−3.30%	−5.82%	−2.17%
CoVaR (normal)	−2.49%	−4.64%	−1.80%
ΔCoVaR	−82 bp	−119 bp	−37 bp
Amplification ratio	—	3.18×	
MES difference, TechDebt stress	−192 bp	−377 bp	−111 bp

The copula estimates provide an additional check on co-crash risk. The t-copula is preferred to the Clayton copula by AIC and yields a lower-tail dependence coefficient of 0.069. The magnitude is not large in absolute terms, but it is economically meaningful because Gaussian dependence has zero asymptotic tail dependence. Linear correlation therefore understates the probability of joint lower-tail realisations precisely in the states that matter for financial stability.

Table 8: Copula-based lower-tail dependence between KRE and TechDebt

Model	Parameter 1	Parameter 2	Lower-tail dependence	AIC
t-copula	0.2918	7.0100	0.0694	−311.989
Clayton copula	0.4504	—	0.2146	−235.427

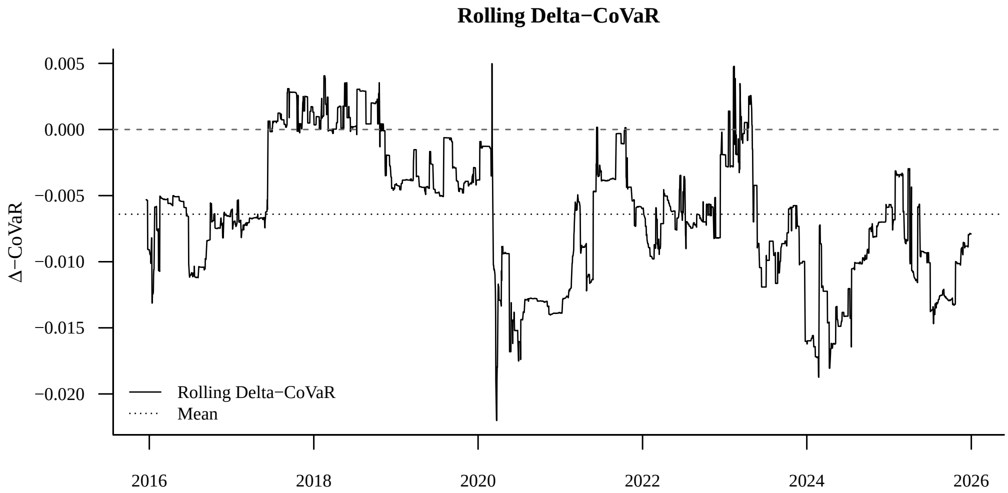

Figure 5: Rolling 252-day ΔCoVaR. The most negative values coincide with COVID-19, the rate-rise period and the spring 2023 regional-bank stress episode.

5. Robustness and Additional Validation

The robustness evidence is organised around four possible objections: the regime result may be a correlation artefact; the tail-risk result may depend on the stress definition; the relationship may be absorbed by broad macro-financial controls; and the dynamic channel may be too weak to matter beyond contemporaneous co-movement. The tests below do not transform the paper into a causal-identification exercise. They are intended to show that the financial-stability interpretation is not driven by one model, one window length or one sign convention.

5.1. Contemporaneous regime correlations and DCC validation

The first validation compares contemporaneous KRE–TechDebt correlations across regimes. Table 9 shows that the correlation is 0.4118 in Regime 1 and 0.3166 in Regime 2. A Fisher z-test rejects equality at the 5% level, with a one-sided p-value of 0.0113 in the direction predicted by H2. This result is useful because it does not rely on the Markov-switching slope estimates. It simply asks whether the same observations classified as stress observations display stronger same-day co-movement.

Rolling ΔCoVaR is also robust to window length. Figure 6 shows that 126-, 252- and 504-day windows all produce the same qualitative pattern: the measure is persistently negative

Table 9: Contemporaneous correlations by regime and Fisher test

Regime	n	Corr(TechDebt,KRE)	95% CI low	95% CI high	Corr(SMH,KRE)
Regime 1: stress	530	0.4118	0.3385	0.4801	0.4584
Regime 2: tranquil	2338	0.3166	0.2796	0.3526	0.5225
Fisher difference		0.0952	$z = 2.2791$		$p = 0.0227$

One-sided p-value for higher KRE–TechDebt correlation in Regime 1 is 0.0113.

and becomes more negative during stress episodes. The shorter window is noisier, as expected, but the regime-aligned timing is unchanged.

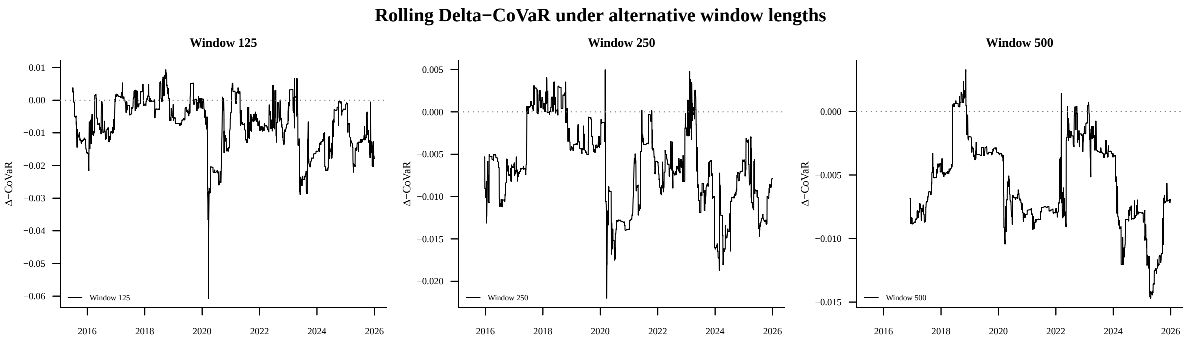

Figure 6: Rolling ΔCoVaR under alternative window lengths. The conclusion does not depend on the benchmark 252-day window.

5.2. *Alternative technology-stress definitions*

Table 10 replaces the baseline technology-HY return stress variable with relative technology-credit stress proxies. The baseline ΔCoVaR is −81.64 bp. The relative spread/stress variants produce a much larger ΔCoVaR of approximately −651 bp. This should not be read as the preferred loss estimate. It is better interpreted as an upper-bound stress proxy because the relative credit-equity measure captures both credit repricing and technology-equity underperformance. The relevant robustness result is that the sign convention and stress transformation do not overturn the qualitative conclusion that technology-credit stress worsens the banking-sector left tail.

Table 10: Alternative technology-stress definitions and ΔCoVaR

Stress variable	VaR	CoVaR stress	CoVaR normal	ΔCoVaR (bp)
Ret_TechDebt	−0.0289	−0.0330	−0.0249	−81.64
Spread_Tech	−0.0289	−0.0904	−0.0253	−651.22
SCSI_norm	−0.0289	−0.0904	−0.0253	−651.22
SCSI_stress	−0.0289	−0.0904	−0.0253	−651.22

5.3. *Macro-financial controls and bank-proxy checks*

The most conservative robustness check uses the shorter 2023–2025 control sample and includes broad risk, credit, volatility and rate controls. Table 11 shows that the regime-conditioned ΔCoVaR result remains economically large after adding dVIX, aggregate HY OAS changes, 10-year Treasury changes and the 10Y–2Y slope. The controlled ΔCoVaR is −120.98 bp in Regime 1 and −46.44 bp in Regime 2. The ratio is smaller than in the full sample, but the ordering is unchanged: the stress-regime left-tail contribution remains materially larger.

Table 11: Macro-controlled regime ΔCoVaR in the shorter control sample

Regime	n	CoVaR stress	CoVaR normal	ΔCoVaR	ΔCoVaR (bp)
Regime 1: stress	147	−0.0418	−0.0297	−0.0121	−120.98
Regime 2: tranquil	538	−0.0210	−0.0163	−0.0046	−46.44

Controls are dVIX, dHY_OAS, dDGS10 and the 10Y–2Y yield-curve slope. The smaller sample reflects the availability of the macro-financial control block.

Table 12 repeats controlled transmission regressions across alternative financial-sector proxies. In the control sample, the direct controlled slope is not statistically significant. This is an important qualification rather than a problem to hide. It means the baseline result should not be sold as a fully identified causal effect after conditioning on all broad risk factors over the short post-2023 window. At the same time, the ordering of coefficients is consistent with the financial-stability narrative: the magnitude is largest for KRE, smaller for KBWB and XLF, and remains negative for the regional-bank-specific residual. The controlled proxy exercise therefore supports the choice of KRE as the main banking-fragility proxy while preserving the scope of Paper 1 as a regime and stress-calibration article, not a causal identification article.

Table 12: Controlled transmission across banking-sector proxies

Proxy	Coverage	$\hat{\beta}_{TD}$	SE	p-value	Adj. R^2
KRE	Regional banks	−0.8274	0.6472	0.2016	0.3693
KBWB	Community banks	−0.1989	0.5620	0.7234	0.4749
XLF	Broad financials	−0.2629	0.3994	0.5105	0.5372
KRE residualized on XLF	Regional-specific component	−0.4039	0.3591	0.2611	0.0664

All specifications use the 2023–2025 control sample and include macro-financial controls. The lack of statistical significance cautions against a strong causal interpretation in Paper 1.

5.4. *Dynamic spillovers, impulse responses and early-warning evidence*

The dynamic layer is included to answer a narrower question than the main Markov-switching and ΔCoVaR results. The baseline evidence is contemporaneous: it shows that the KRE–TechDebt link is much stronger when the market is already in the refinancing-stress state. The dynamic checks ask whether this same state dependence is visible when the system is treated as a short-run multivariate process. This distinction matters for interpretation. Daily credit and equity prices incorporate information quickly, so the absence of a long and statistically precise impulse response should not be read as evidence against the contagion channel. It means that the channel operates mainly through same-day repricing and short-lived balance-sheet pressure rather than through a persistent, multi-week propagation pattern.

Table 13 summarizes the dynamic diagnostics. Forecast-error variance decomposition assigns a larger role to TechDebt innovations in the stress regime than in the tranquil regime, while the Diebold–Yilmaz estimates show a more negative TechDebt net-spillover position in Regime 1. This is consistent with the interpretation that technology credit stress becomes a stronger transmitter or amplifier of banking-sector shocks when investors are forced to reduce risk. The MES and controlled ΔCoVaR rows show that the same ordering remains visible in conditional losses and in the shorter sample with macro-financial controls. The early-warning row is deliberately stated conservatively: the regime signal contains statistically significant linear predictive content at horizons of 3–10 trading days, but the nonlinear logit specification is not strong enough to support a stand-alone early-warning claim.

Table 13: Additional dynamic validation evidence

Block	Object	Regime 1	Regime 2	Interpretation
Diebold–Yilmaz	TechDebt net spillover	−7.7359	−3.5072	More negative net spillover in stress.
MES	ΔMES, TechDebt stress	−377.05 bp	−111.27 bp	Larger conditional bank loss in stress.
Controlled ΔCoVaR	Ret_TechDebt	−120.98 bp	−46.44 bp	Tail-risk ordering survives controls.
Early warning	Significant linear leads	3–10 days	—	Some short-horizon signal, not a standalone early-warning model.

This table summarises additional diagnostics. It is not used as the primary identification strategy, but it checks whether the regime ordering survives outside the main contemporaneous Markov-switching specification.

The impulse-response functions in Figure 7 are estimated separately within the two Markov regimes using regime-specific VAR lag lengths selected by information criteria. The stress-regime VAR uses three lags, while the tranquil-regime VAR uses one lag. Confidence bands are obtained from bootstrap draws and the responses are reported in return units, so their economic interpretation is best made in relative terms across regimes rather than as a structural shock decomposition. In Regime 1, the first-day KRE response to a TechDebt innovation is visibly larger in magnitude and negative, and the cumulative response over the displayed horizon is approximately −42 basis points. In Regime 2, the response is much smaller, with a cumulative value close to zero. The confidence intervals include zero in both regimes, which is why the IRF evidence is treated as validation rather than as the paper's main identification result.

Table 14: Impulse-response diagnostics by regime

Regime	n	VAR lag	Cum. IRF	Cum. 95% CI	Peak IRF	Peak 95% CI
Regime 1: stress	530	3	−0.004249	[−0.006341, 0.006125]	−0.003444	[−0.001632, 0.004207]
Regime 2: tranquil	2338	1	0.000360	[−0.000468, 0.000638]	0.000264	[−0.000447, 0.000620]

IRFs are in decimal return units. Peak responses occur at horizon 1. All 95% confidence intervals include zero.

The table makes the regime contrast visible without overstating the result. The stress-regime cumulative response is negative and roughly an order of magnitude larger than the tranquil-regime cumulative response. The peak response has the same ordering. However, because the bootstrap bands are wide in the smaller stress sample, the IRF evidence supports the interpretation of fast stress-state repricing but does not by itself identify a persistent dynamic causal chain.

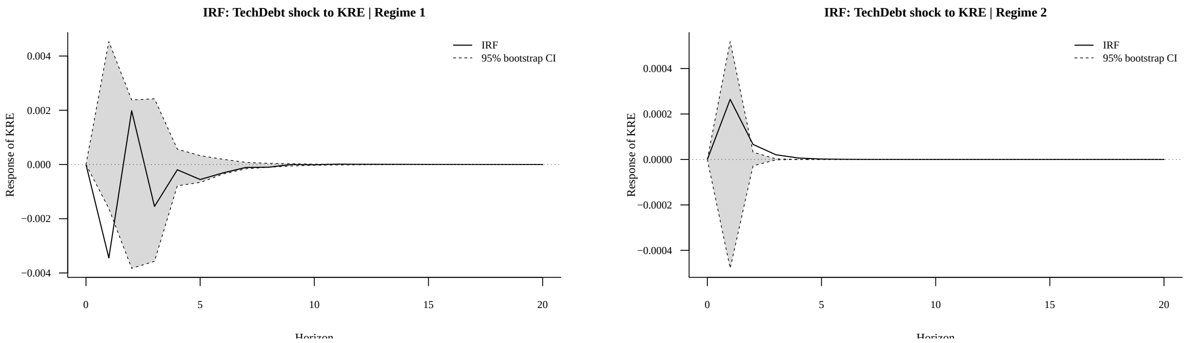

(a) Stress regime

(b) Tranquil regime

Figure 7: Bootstrap impulse responses of KRE to a TechDebt innovation by regime. The stress-regime response is larger and initially negative, but the confidence intervals show that the dynamic layer is less precise than the contemporaneous coefficient and ΔCoVaR evidence.

Figure 8 provides a complementary way to read the same dynamic information. The left panel plots the linear lead coefficients linking the current refinancing-stress signal to future ΔCoVaR. The coefficients become increasingly negative across short horizons and are statistically significant for a subset of leads, implying that the regime probability is not merely a coincident label. The right panel shows that the corresponding logit odds ratios rise with the lead, but the nonlinear classification evidence is weaker. The correct conclusion is therefore cautious: the regime probability is useful in a monitoring dashboard, especially as a warning that tail risk may be building, but it should be combined with credit spreads, bank funding indicators and market liquidity measures rather than used mechanically as a binary alarm.

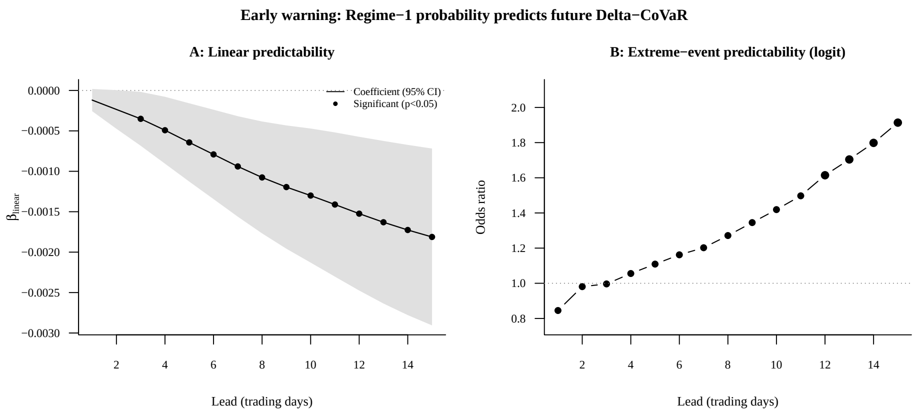

Figure 8: Early-warning validation. The technology-stress signal contains some linear predictive information at short leads, but the evidence is weaker than the contemporaneous regime and tail-risk results.

The dynamic checks therefore refine the main result without converting the paper into a VAR-based causal exercise. The strongest evidence is not a long-lived propagation effect that unfolds gradually over many trading days. It is the rapid regime-dependent repricing of banking-sector risk and the accompanying amplification of conditional left-tail losses. This timing is important for stress testing. A loss multiplier that activates quickly in a stress state can be missed by full-sample dynamic averages even when it is economically relevant for supervisory monitoring.

6. Discussion: Economic and Policy Implications

6.1. The amplification factor as a balance-sheet object

The central numerical object in the paper is the 2.82× amplification of the technology-HY transmission coefficient between tranquil and stress regimes. This number should not be read as a purely statistical preference for a Markov-switching specification. It is an empirical representation of the balance-sheet mechanism developed in Section 2. In tranquil markets, a decline in technology-credit prices can be absorbed by diversified investors, dealers and funds without forcing an immediate revaluation of regional-bank equity. Losses on securities books may remain an accounting or valuation issue rather than a market-wide solvency signal. In a refinancing-stress regime, the same credit-price movement occurs when debt rollover risk is salient, fund redemptions are more likely, collateral values are under pressure and bank investors are already reassessing the boundary between liquidity problems and capital impairment.

This distinction matters because banking fragility is not linear in the size of the initial credit shock. A small technology-HY movement in a liquid market can be absorbed by the usual risk-bearing chain. A similar movement in a constrained market can force investors to reduce exposures across several assets simultaneously. The common-creditor mechanism of Coval and Stafford (2007) and the fire-sale logic of Shleifer and Vishny (2011) imply that prices can become linked through the identity of sellers rather than through direct cash-flow exposure. Corporate bond mutual funds are particularly relevant in this setting because redemption pressure can turn portfolio adjustment into price pressure on the bonds that remain in the system (Choi et al., 2020). Intermediary asset-pricing models add the state variable: the price impact of a shock is small when balance sheets are elastic, but large when dealer capital, funding liquidity or risk-bearing capacity is scarce (Adrian and Shin, 2010; Brunnermeier and Pedersen, 2009; He and Krishnamurthy, 2013).

The Markov-switching estimates map directly onto this economic structure. The stress-regime coefficient of 2.91 means that, conditional on being in the refinancing-stress state, a technology-HY return innovation is associated with a much larger contemporaneous movement in KRE than in tranquil periods. The increase in residual variance is equally important. It shows that the stress state is not merely a high-beta version of the same market environment; it is a state in which unexplained bank-equity risk is also materially larger. The combination of a higher slope and a higher residual variance is exactly what one would expect when investors begin to price technology-credit stress, bank balance-sheet exposure and market liquidity as a single fragile configuration.

The close correspondence between the coefficient amplification and the ΔCoVaR amplification is economically informative. The conditional-mean model and the conditional-quantile model are estimated differently and answer different questions. The Markov-switching coefficient asks how the average KRE response to technology-HY movements changes across states. The regime-conditioned ΔCoVaR asks how the left tail of KRE changes when technology-HY distress is present. The fact that these two quantities produce similar multipliers indicates that the regime mechanism affects both the center and the lower tail of the bank-return distribution. For financial-stability purposes, this is stronger than a result that appears only in the mean or only in the tail. It implies that the stress state changes the entire pricing environment in which banks and technology credit are jointly valued.

A useful way to interpret the amplification factor is as a state-contingent shadow risk weight. In tranquil periods, technology HY is priced mainly as a sectoral credit exposure. In the stress regime, it carries an additional financial-stability premium because the market treats it as connected to refinancing pressure, fund liquidity and regional-bank fragility. A constant risk weight averages these two states and therefore assigns too little weight to the state that matters for stress testing. The paper does not claim that every technology-HY loss mechanically becomes a bank loss. It claims that the market price of regional-bank equity becomes much more sensitive to technology-HY stress when the financial system is already in the regime in which balance-sheet constraints and fire-sale externalities are likely to bind.

6.2. Stress-testing reform: from a constant risk weight to a regime schedule

Current supervisory stress-testing practice generally models HY stress through broad credit factors. A sector-specific technology-HY sensitivity is therefore implicitly absorbed into average HY-bank correlations. The results here suggest that this is not adequate when technology refinancing stress coincides with regional-bank fragility. A stress test calibrated to full-sample or tranquil-period correlations will understate the loss that a genuine refinancing-stress episode would generate, because the calibration sample under-weights exactly the observations stress tests are intended to anticipate.

Table 15: Regime-adaptive stress-test calibration implied by the estimates

Component	Constant-parameter treatment	Regime-adaptive treatment
Credit shock	Aggregate HY shock	Technology-HY refinancing-stress block
Bank sensitivity	Full-sample bank–HY beta/correlation	Stress-regime transmission coefficient $\hat{\beta}_{TD}^{(1)} = 2.91$
Tail-loss calibration	Full-sample or rolling ΔCoVaR	Stress-regime ΔCoVaR of −119 bp
Monitoring variable	Spread level or unconditional volatility	Smoothed probability of Regime 1
Capital implication	Static exposure weight	State-contingent multiplier activated in refinancing stress

The table does not prescribe a new capital formula. It translates the paper's estimates into the minimum components required for a supervisory technology-HY refinancing-stress module.

6.3. *Macroprudential monitoring and limits of interpretation*

The smoothed Regime 1 probability is useful as a monitoring variable because the stress regime is persistent enough to carry supervisory information. A one-day spike in volatility may be difficult to interpret, but a rising probability of the refinancing-stress regime provides a compact signal that several pieces of market information are being jointly repriced. The HY spread level tells a supervisor how wide credit spreads are. The regime probability tells the supervisor whether technology-credit stress is becoming statistically coupled with regional-bank equity risk. These objects are related, but they are not substitutes. A spread can widen because of sector-specific credit news without becoming a banking-system issue. A high regime probability indicates that the market is pricing the technology-credit block and the regional-bank block as part of the same stress environment.

In a macroprudential dashboard, the regime probability would be most useful when combined with observable indicators that are closer to the balance-sheet mechanism. Examples include technology-HY spread changes, fund-flow pressure in corporate-bond funds, bank funding spreads, unrealised-loss indicators, measures of Treasury and credit-market liquidity, and equity-implied volatility for regional banks. The regime estimate should not replace these variables. Its role is to summarize when their interaction has moved into the state in which the paper estimates a materially larger bank-sensitivity parameter. This is why the policy proposal in Table 15 is formulated as a regime-adaptive module rather than as a mechanical trigger. The module would make supervisory stress scenarios conditional on the probability of being in the stress state, while leaving room for judgement about the source, persistence and institutional distribution of the shock.

The monitoring interpretation also clarifies why the dynamic evidence is useful even though the impulse responses are not statistically decisive. The IRFs show that the stress-state response is larger and initially negative, but the confidence bands are too wide to support a claim of persistent dynamic causality. This is not a weakness for the paper's central stress-testing argument. Daily market prices often incorporate credit and liquidity news immediately. A contemporaneous amplification in the coefficient and in ΔCoVaR can therefore be more relevant for supervisory loss calibration than a multi-day response path. The early-warning results should be read in the same restrained manner. They suggest that the regime signal contains short-horizon information about tail-risk build-up, but they do not justify using the Markov state as an automatic alarm.

The robustness checks define the boundary of the contribution. Macro-controlled OLS slopes in the shorter 2023–2025 sample are not statistically significant, and the bank-proxy regressions weaken once broad financial conditions are controlled. Those results prevent an excessive causal interpretation. They indicate that Paper 1 should not be presented as evidence that technology-HY shocks unconditionally cause bank stress across all samples, all banks and all macro-financial environments. The more defensible statement is narrower and stronger for a financial-stability audience: the technology-HY–bank linkage has very different systemic content across regimes, and the stress-state calibration is materially larger than the tranquil-state calibration.

This boundary is also what separates the present paper from the companion articles. The second paper can focus on causal identification, including instruments, event timing or alternative identification restrictions. The third paper can decompose whether the channel works through risk appetite, intermediary constraints or portfolio rebalancing. The present article keeps the object deliberately closer to supervisory practice. It shows that a constant-parameter stress-test calibration averages over states that should not be averaged over. Preserving this nonlinearity is the key macroprudential implication: a monitor that smooths away regime dependence will look most precise in tranquil markets and least informative when the refinancing-stress state is reached.

7. Conclusion

This paper asks whether technology-sector high-yield debt has become a regime-dependent financial-stability vulnerability for regional banks. The answer is yes in the specific sense relevant for stress testing: the technology-HY–bank relationship is not adequately represented by a single constant coefficient. The OLS benchmark shows a significant average association between technology HY returns and KRE returns. The Markov-switching model then shows that this average conceals a materially larger stress-regime transmission coefficient, with a $2.82\times$ amplification relative to tranquil periods and a 6.97-fold increase in residual variance. Regime-conditioned ΔCoVaR provides an independent tail-risk confirmation: the banking-sector left-tail contribution of technology HY stress is -119 basis points in the stress regime and -37 basis points in tranquil periods.

The economic interpretation is a common-creditor and intermediary-constraint mechanism. Technology HY stress is not systemically important in every state of the world. It becomes important when refinancing pressure coincides with limited risk-bearing capacity, so that credit repricing, bank-equity selling and funding concerns reinforce one another. This is why the paper's main object is the regime multiplier rather than the full-sample coefficient.

The policy implication is equally direct. Stress tests calibrated to average HY-bank correlations can understate losses in refinancing-stress states, because the calibration sample gives too much weight to tranquil observations. A regime-adaptive technology-HY module would instead condition the bank sensitivity and the tail-loss estimate on the probability of the stress regime. Such a module does not require a wholesale redesign of supervisory stress testing. It requires adding a sector-specific, state-contingent calibration block for a credit segment whose importance has increased with the growth of technology leveraged finance.

References

Acharya, V., Pedersen, L., Philippon, T., Richardson, M., 2017. Measuring systemic risk. *Review of Financial Studies* 30, 2–47.

Adrian, T., Brunnermeier, M.K., 2016. CoVaR. *American Economic Review* 106, 1705–1741.

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. *Journal of Financial Intermediation* 19, 418–437.

Bai, J., Perron, P., 1998. Estimating and testing linear models with multiple structural changes. *Econometrica* 66, 47–78.

Becker, B., Ivashina, V., 2015. Reaching for yield in the bond market. *Journal of Finance* 70, 1863–1902.

Billio, M., Getmansky, M., Lo, A.W., Pelizzon, L., 2012. Econometric measures of connectedness and systemic risk. *Journal of Financial Economics* 104, 535–559.

Brownlees, C., Engle, R.F., 2017. SRISK: A conditional capital shortfall measure of systemic risk. *Review of Financial Studies* 30, 48–79.

Brunnermeier, M.K., Oehmke, M., 2014. Predatory short selling. *Review of Finance* 18, 2153–2195.

Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22, 2201–2238.

Choi, J., Hoseinzade, S., Shin, S.S., Tehranian, H., 2020. Corporate bond mutual funds and asset fire sales. *Journal of Financial Economics* 138, 432–457.

Coval, J., Stafford, E., 2007. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics* 86, 479–512.

Diebold, F.X., Yilmaz, K., 2012. Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting* 28, 57–66.

Engle, R., 2002. Dynamic conditional correlation. *Journal of Business & Economic Statistics* 20, 339–350.

Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. *American Economic Review* 102, 1692–1720.

Girardi, G., Ergün, A.T., 2013. Systemic risk measurement: Multivariate GARCH estimation of CoVaR. *Journal of Banking & Finance* 37, 3169–3180.

Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.

Guidolin, M., Timmermann, A., 2007. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control* 31, 3503–3544.

Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series. *Econometrica* 57, 357–384.

He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. *American Economic Review* 103, 732–770.

Holmström, B., Tirole, J., 1998. Private and public supply of liquidity. *Journal of Political Economy* 106, 1–40.

Karimalis, E.N., Nomikos, N.K., 2018. Measuring systemic risk in the European banking sector: A copula CoVaR approach. *European Journal of Finance* 24, 944–975.

Pelletier, D., 2006. Regime switching for dynamic correlations. *Journal of Econometrics* 131, 445–473.

Reboredo, J.C., Ugolini, A., 2016. Systemic risk in Spanish listed banks: A vine copula CoVaR approach. *Spanish Journal of Finance and Accounting* 45, 1–31.

Shleifer, A., Vishny, R.W., 1992. Liquidation values and debt capacity: A market equilibrium approach. *Journal of Finance* 47, 1343–1366.

Shleifer, A., Vishny, R.W., 2011. Fire sales in finance and macroeconomics. *Journal of Economic Perspectives* 25, 29–48.