

Penalties for Hedging and Synchronized Inaction under the Fundamental Review of the Trading Book

Scott Condie*

July 2026 — Preliminary draft

Abstract

Under the Fundamental Review of the Trading Book, most trading banks compute market risk capital using the standardized approach. This sensitivities-based method requires the entire charge to be computed three times, under a high, a medium, and a low prescribed correlation scenario, with the largest of the three binding. Taking the worst of a set of correlation models echoes the maxmin criterion of [Gilboa and Schmeidler \(1989\)](#). This paper studies what that design does to the banks that face it. The first result is that the rule *penalizes hedging*. A hedged book is always charged at the regulator's lowest correlation, so hedging is credited only up to that correlation and charged beyond it, and a fully matched hedge carries strictly more capital than an open position whenever the prescribed correlation is below two thirds. That threshold covers the equity buckets, the base-metals bucket, and most cross-bucket aggregations. The second result is that a bank optimizing against the constraint has an *inaction region*, an interval of news over which its capital holdings against risk do not adjust to information, and, where it is the marginal investor its quotes, do not move. The same freezing occurs when the traded factor's volatility changes and its correlation does not, because the charge is computed from the regulator's risk weights rather than the asset's true current volatility. The region's width is the shadow cost of regulatory capital times the regulator's scenario spread, so it widens in stress, and because every bank faces the same three correlation matrices these regions lie at the same portfolio configurations for all banks. What is synchronized across the system is not how much risk banks perceive but where they stop responding to news. Calibrated at the prescribed Basel correlations, the scenario spread reaches 40 percentage points and the mechanism is shown to be quantitatively meaningful at current policy correlations. Policy suggestions are given to help reduce the negative consequences of this design without undermining its prudential intent.

Keywords: FRTB; standardized approach; correlation scenarios; ambiguity aversion; hedging; capital regulation; financial stability; procyclicality; market liquidity; price discovery.

*Department of Economics, Brigham Young University, Provo, UT 84604. Email: scott_condie@byu.edu.

JEL classification: G21; G28; G18; D81; G14.

1 Introduction

Since the implementation of the Fundamental Review of the Trading Book (FRTB), the Basel Committee’s post-crisis overhaul of market risk capital, most trading banks compute their market risk capital using the standardized approach. The internal-models alternative is costly to obtain and to maintain. Facing the profit-and-loss attribution test, the risk-factor eligibility test, and non-modellable risk factor add-ons, most large dealers have elected not to seek approval at all. Industry surveys report that banks holding roughly 85% of trading-book capital under internal models under Basel 2.5 plan internal-model coverage of only 15–40% under FRTB, and that about ten banks worldwide intend to apply ([International Swaps and Derivatives Association and EY, 2024](#); [KPMG, 2025](#)). The sensitivities-based method (SBM) of the standardized approach is therefore not a backstop for small institutions. It is the marginal capital metric in most of the world’s fixed income, equity, credit, and derivatives markets.

This paper studies one step of that calculation. After a bank computes its delta, vega, and curvature charges using regulator-prescribed risk weights and correlations, paragraph MAR21.6 requires that, “[i]n order to address the risk that correlations may increase or decrease in periods of financial stress,” the entire aggregation be repeated under three prescribed correlation scenarios ([Basel Committee on Banking Supervision, 2019a](#)). The medium scenario leaves the specified correlations unchanged; the high scenario multiplies each correlation by 1.25, capped at 100%; and the low scenario replaces each correlation ρ by $\max(2\rho - 100\%, 75\% \times \rho)$. The bank’s capital requirement is the largest of the three resulting totals. That rationale is empirically grounded. Correlations are unstable and can move either way in stress, hurting diversified books when they rise and hedged books when they fall ([Basel Committee on Banking Supervision, 2019b](#)).

However, this design has a specific analog in decision theory that reveals some of the implications of its use. Taking the worst outcome over a set of correlation models is the maximin expected utility criterion of [Gilboa and Schmeidler \(1989\)](#). That criterion describes an agent who treats the correlation not as a random variable with a known distribution but as a Knightian uncertainty, and who acts on the worst model in an admissible set. The scenario-max installs that criterion in the capital constraint of every standardized-approach bank, whether or not the bank itself holds precise probabilistic beliefs. The question in this paper is what this criterion does to the banks that face it, and the answer is not that they hold more capital against the same book.

The first result concerns hedging. Because the capital charge is decreasing in the correlation for any hedged position, the maximum over scenarios always selects the regulator’s *lowest*

correlation, so a hedged book is evaluated as though its hedges were the least reliable the rule admits. Proposition 1 shows what follows. The charge rewards hedging only up to the low-scenario correlation and rises strictly beyond it, so on the interval between the low-scenario hedge ratio and the variance-minimizing one, each additional unit of hedge lowers true risk and raises required capital. Corollary 1 sharpens this into a statement about the rule as calibrated. A fully matched hedge requires strictly more capital than an open position whenever the prescribed correlation is below two thirds, a threshold that covers the equity buckets, the base-metals bucket, and most cross-bucket aggregations. Any hedge leaning on a prescribed correlation of 80% or higher has its standalone charge inflated by exactly $\sqrt{2}$, however good the hedge. The rule does not merely under-credit hedging. It charges for it.

The second result concerns news. Proposition 2 shows that a dealer optimizing against the constraint has an *inaction region*. The marginal capital cost of a trade jumps discontinuously where the binding scenario switches, at the boundary between hedged and directional books. Over an interval of news realizations the desired trade is not worth that jump, so the position does not move, and neither do the dealer’s quotes where it is the marginal investor. The same inaction appears when the traded factor’s volatility changes with its correlation held fixed, because the charge is computed from the regulator’s risk weights rather than the desk’s own volatilities (Proposition 3). The region’s width is the shadow cost of regulatory capital times the regulator’s scenario spread. Both factors are quantifiable. The spread is a policy constant that reaches 40 percentage points at a prescribed correlation of 0.80, the range of many rates-curve and cross-bucket correlations. In this setup, the shadow cost is a state variable that spikes when balance sheet is scarce (He and Krishnamurthy, 2013; Du et al., 2018). Inaction regions therefore widen precisely when markets most need dealers to respond to news.

These theoretical results are inspired by those in Illeditsch et al. (2021), which studies investors who face ambiguity about the correlation between a public signal and asset fundamentals, and shows that these investors’ optimal portfolios exhibit *information inertia*. Over an interval of signal realizations the position does not respond to the signal at all, and in equilibrium neither does the price. The inertia there is not inattention or a transaction cost. It is the response of a maxmin agent whose worst-case correlation shifts endogenously with the news, so that a signal’s effect on the worst-case mean is exactly offset by its effect on worst-case risk. The SBM methodology of the FRTB imposes a similar structure to bank capital requirements. Intermediaries calculating capital requirements according to the SBM will not see benefits from hedging under certain scenarios because of the imposition of the scenario-max in a way that is similar to ambiguity averse investors in the model. Model-risk and prudent-valuation expectations invite banks to adopt the regulator’s scenario set as the admissible model class for their own risk management. Where that happens, the belief-based mechanism operates on top of the constraint-based one.

At a macro level, since standardized-approach banks in most jurisdictions faces the *same*

three scenario matrices, the implications described above have the potential to become a constant of the financial system. Other work has shown that common regulatory policies can lead to market-wide effects. For example, value-at-risk and risk-based leverage constraints amplify positions in booms and force deleveraging in stress (Adrian and Shin, 2010; Brunnermeier and Pedersen, 2009), and a common regulatory risk model imposes a monoculture that synchronizes behavior and manufactures endogenous risk (Danielsson, 2002). Our departure is that the potential harm here does not come from the *level* of a shared risk estimate nor from leverage targeting, but from a shared *non-convexity*. What is synchronized is not how much risk banks perceive but where they stop responding to news. The effect lives in price discovery and dealer liquidity provision (He and Krishnamurthy, 2013), propagates through the funding constraints that liquidity spirals require (Brunnermeier and Pedersen, 2009), and is invisible in any single institution’s capital adequacy.

One policy implication is that the *functional form* of prudence is itself a policy choice. Nothing here argues against conservatism about unstable correlations, which is well founded; the argument is that the worst-case scenario-max imposes costs on doing things that might otherwise be desirable both at an individual and at a market level. Section 6 develops three low-cost responses, from collecting the scenario figures banks already compute, to using the staggered rollout as a natural experiment, to replacing the maximum with a smooth aggregation of the same three scenarios.

Throughout, the scenario-max refers to MAR21.6’s requirement that the SBM capital requirement be the largest of the three correlation-scenario aggregations. The delta, vega, and curvature charges, the bucket structure, and the risk-weight calibration raise their own questions, and the deliberately low inter-name correlations already impose severe haircuts on hedges. But they are smooth features, and we take them as given. The object of interest is the non-smooth layer applied on top of them, because that is the layer with a known behavioral signature. The mapping from a capital constraint to a maxmin preference is functional rather than literal, and it operates under three conditions made explicit in Section 4.4. The charge must bind, the dealer must be marginal in the relevant market, and internal capital allocation must transmit the regulatory shadow price to the desk. We establish the mechanism’s existence, its sign, and its observable footprint. The varied roll-out of the FRTB allows for potential measurement of the effects of this policy change, among others. This is discussed in Section 6.

The remainder of the paper proceeds as follows. Section 2 describes the scenario-max and quantifies its capital footprint for three stylized banks. Section 3 reviews what the theory of maxmin behavior over correlations implies. Section 4 develops the constraint and beliefs channels and calibrates the inaction region. Section 5 draws out the consequences, Section 6 the policy responses, and Section 7 concludes. Appendix A reproduces the prescribed correlation parameters and Appendix B gives steps to reproduce the numerical example. All proofs are in Appendix C.

2 The scenario-max in the standardized approach

The SBM computes capital for seven risk classes (general interest rate risk, three credit spread classes, equity, commodity, and foreign exchange). Within each class, first-order (delta and vega) and second-order (curvature) sensitivities to prescribed risk factors are multiplied by regulator-set risk weights and aggregated twice, first within buckets using prescribed correlations ρ and then across buckets using prescribed correlations γ , via variance-style formulas of the form

$$K_b = \sqrt{\sum_k \text{WS}_k^2 + \sum_k \sum_{l \neq k} \rho_{kl} \text{WS}_k \text{WS}_l}, \quad \text{Capital} = \sqrt{\sum_b K_b^2 + \sum_b \sum_{c \neq b} \gamma_{bc} S_b S_c}, \quad (1)$$

where WS_k is the risk-weighted net sensitivity to risk factor k (Basel Committee on Banking Supervision, 2019a). Every parameter in (1) is set by the regulator; the bank supplies only the sensitivities. The SBM is therefore a parametric portfolio risk model whose covariance matrix is prescribed rather than estimated. Capital is levied as if the covariance of trading-book payoffs were the regulator’s matrix, whatever the truth.

The scenario-max is then applied at the top of this structure. To address the risk that correlations move in stress, the bank must compute the *entire* risk-class aggregation three times (Basel Committee on Banking Supervision, 2019a):

- under the **medium** scenario: the correlations ρ_{kl} and γ_{bc} as specified in the standard;
- under the **high** scenario: every correlation replaced by $\min(1.25 \cdot \rho, 100\%)$;
- under the **low** scenario: every correlation replaced by $\max(2\rho - 100\%, 75\% \cdot \rho)$.

The portfolio-level capital requirement is the largest of the three totals. Writing $\mathcal{C} = \{\text{low, med, high}\}$ for the scenario set and $K(\theta; c)$ for the SBM charge on portfolio θ under scenario c , the rule is

$$K(\theta) = \max_{c \in \mathcal{C}} K(\theta; c). \quad (2)$$

Three institutional observations matter for everything that follows.

First, the max is taken portfolio-wide, not position-by-position. A single scenario is selected, the one worst for the book as a whole. Which scenario binds is therefore an equilibrium property of the entire portfolio’s composition. A book dominated by directional exposures is charged under the high-correlation scenario (correlated longs offer little diversification); a book dominated by hedged, long-short exposures is charged under the low-correlation scenario (hedged are assumed to work less well). As the book’s composition moves, the binding scenario can and does switch.

Second, the scenario set is identical for every bank, in every jurisdiction, at every date. The multipliers 1.25 and the low-scenario formula are hard-coded in the standard and have

Table 1: SBM capital under the three correlation scenarios for three stylized books (\$mm). The binding requirement is the largest of the three scenario totals; the loading is that maximum less the medium-scenario figure.

Portfolio	Scenario total			Binding scenario	SBM (max)	Scenario-max loading	
	Low	Medium	High				
A. Community Bank (rates, directional)	113.0	115.3	117.5	High	117.5	+2.2	(1.9%)
B. Commodity Bank (commodity spreads)	197.4	168.1	124.5	Low	197.4	+29.3	(17.4%)
C. Atlantic Bank (equity/rates/FX)	531.1	487.7	431.3	Low	531.1	+43.4	(8.9%)

been adopted essentially verbatim in the EU (CRR3), the UK, Canada, Japan, and the other jurisdictions that have gone live. Under internal models, banks’ worst-case beliefs differed because their models and estimation windows differed. Under the SBM, the “worst case” is a public constant of the financial system.

Third, the scenario-max is not a remote tail provision; it is the everyday binding margin. For any non-degenerate book, one of the outer scenarios is strictly binding, since the medium scenario is the maximizer only on a knife edge. The marginal capital cost of every trade is therefore evaluated under either the high- or the low-correlation matrix. Every risk-return decision made by a standardized-approach trading desk that is charged for regulatory capital is therefore made each day against a worst-case correlation model.

2.1 A numerical illustration: three banks, one formula

To illustrate the potential consequences of this policy, consider three stylized standardized-approach banks, each with a trading-and-hedging book characteristic of its franchise. *Community Bank* is a local home- and small-business lender whose only material market risk is the interest-rate exposure of its banking-book hedges, a single-currency, net-long-duration position that is essentially directional. *Commodity Bank* is a commodity-finance house whose energy desk holds trading positions in oil, base metals, natural gas, and small rates and FX financing positions. *Atlantic Bank* is an international dealer holding both long and short equity positions, a hedged rates book, a multi-currency FX book, and an investment-grade credit book. Table 1 reports each bank’s SBM capital under the low, medium, and high correlation scenarios, the binding scenario, and the scenario-max loading, which is the extra capital the worst-case operator imposes over the medium (“point-estimate”) figure.¹

The table makes three points. First, *the binding scenario is a function of the business*

¹Positions are net risk-weighted sensitivities ($WS = \text{risk weight} \times \text{sensitivity}$), in \$mm. The correlation parameters are the values prescribed by the Basel standard and reproduced in Appendix A: the GIRR tenor formula (MAR21.46), liquid-combustibles 95% and base-metals 60% within-commodity correlations (MAR21.83), equity within-bucket correlations of 15–25% (MAR21.78), and cross-bucket correlations of 20% (commodity), 15% (equity), 50% (cross-currency GIRR), and 60% (FX). Scenario scaling is that found in MAR21.6. The sensitivities are illustrative and are not calibrated to any actual institution.

model, not of bank size or risk appetite. The directional lender is charged under the *high* scenario, because correlated exposures diversify less when correlations rise, while both hedged intermediaries are charged under the *low* scenario. The “ambiguity aversion” the rule imposes is therefore not uniform. Its impacts are felt differently across business models. Second, *the magnitude of the imposed loading varies by nearly an order of magnitude across business models.* For the community lender it is trivial (1.9%). Its exposures are already highly correlated, so scaling the correlations up, capped at 100%, barely moves the charge. For the commodity house it is large, a 17% increase in market-risk capital, \$29mm, arising *entirely* from the correlation scenario and not from any change in position, price, or risk weight. For the global dealer it is intermediate (8.9%). Third, and least obvious, *the penalty concentrates not on the tightest hedges but on those built from mid-range correlations.* Commodity’s oil positions lean on prescribed correlations of 95–99%, which the low scenario cuts only to 90–98%, so they keep most of their hedge credit; the bite falls instead on its base-metals spread (correlation 60%, cut to 45%) and its interest-rate and FX financing hedges. Atlantic’s loading is likewise driven by its rates book (GIRR charge +34% under the low scenario) and its FX book (+15%), not by its equity relative value (+3%), whose prescribed correlations of 15–25% the SBM already declines to credit as a hedge. The additive cut $\rho_M - \rho_L$ that the low scenario imposes is largest for mid-range correlations, peaking at $\rho_M = 80\%$ (Corollary 1). The scenario spread, the gap between the high and low totals, runs from \$125mm to \$197mm for Commodity Bank. This will be discussed more in Section 4.

Why the SBM does not reward hedging. Commodity’s loading points out a more general property of the capital charge that we show with the following example. Consider a bank that must carry a position of size W_1 in a particular asset. It wishes to hedge this position and will do so with a position of size $h \geq 0$ placed in an offsetting instrument (net sensitivity $W_2 = -h$), with bucket charge $K(h; \rho) = \sqrt{W_1^2 + h^2 - 2\rho W_1 h}$ and scenario-max $\bar{K}(h) = \max_c K(h; \rho_c)$.

Proposition 1 (Hedge recognition under the scenario-max). *Let $\rho_L < \rho_M < \rho_H$. Then:*

- (i) *for every $h > 0$, the scenario-max selects the low correlation: $\bar{K}(h) = K(h; \rho_L)$;*
- (ii) *$\bar{K}$ is minimized at the hedge ratio $h^*/W_1 = \rho_L$, with $\bar{K}(h^*) = W_1 \sqrt{1 - \rho_L^2}$; it is strictly decreasing for $h < \rho_L W_1$ and strictly increasing for $h > \rho_L W_1$.*

Consequently, for any true correlation $\rho_{\text{true}} > \rho_L$, hedging over the interval $h \in (\rho_L W_1, \rho_{\text{true}} W_1]$ strictly raises required capital while strictly lowering the book’s true variance.

The proposition says the framework rewards hedging only up to the ratio ρ_L and charges for it beyond that point (Figure 1). The following corollary shows that this means that in some cases not hedging at all carries a lower capital cost than hedging fully.

Corollary 1 (Full hedges can cost more than none). *Under the low-scenario map $\rho_L = \max(2\rho_M - 1, 0.75\rho_M)$ of [MAR21.6]:*

- (i) *a one-for-one hedge satisfies $\bar{K}(W_1) = W_1\sqrt{2(1-\rho_L)} > \bar{K}(0) = W_1$ if and only if $\rho_L < \frac{1}{2}$, equivalently $\rho_M < \frac{2}{3}$;*
- (ii) *the charge against a matched two-leg hedge is multiplied under the low scenario by $\sqrt{(1-\rho_L)/(1-\rho_M)}$, which equals exactly $\sqrt{2}$ for every $\rho_M \geq 0.80$;*
- (iii) *the difference $\rho_M - \rho_L$ equals $0.25\rho_M$ for $\rho_M \leq 0.8$ and $1 - \rho_M$ for $\rho_M \geq 0.8$, and is maximized at $\rho_M = 0.80$.*

Part (i) covers a wide swath of the SBM. Equity buckets ($\rho_M = 15\text{--}25\%$), base-metals and several agricultural commodity buckets, most cross-bucket aggregations, and the cross-currency FX correlation all have $\rho_M < \frac{2}{3}$, so in each a fully matched hedge is charged *more* capital than the naked inventory. Figure 1 draws the base-metals case ($\rho_M = 60\%$, $\rho_L = 45\%$). Parts (ii)–(iii) demonstrate where this phenomenon is likely to be most strong. Any hedge for an asset with a prescribed correlation of 80% or higher has its standalone charge inflated by 41% in the low scenario, however good the hedge. This effect is widest at $\rho_M = 0.80$. This is why Commodity’s oil spreads ($\rho_M = 95\text{--}99\%$) incur the full $\sqrt{2}$ on their hedged component, though directional length dilutes it at the book level, while the base-metals spread ($\rho_M = 60\%$) sits closest to the worst zone.

All this is not to say that accounting for correlational uncertainty is unwise. Correlation instability in stress is a well-documented fact in financial econometrics. A regulator that does not know which way correlations will move, and that must write one rule for thousands of heterogeneous books, has good reason to refuse to commit to a point estimate. The scenario-max is a defensible expression of genuine Knightian uncertainty on the regulator’s part. This paper isn’t asking whether that uncertainty is real. It is asking what happens when a worst-case operator, rather than a smooth blend of the same three scenarios, becomes the objective against which every marginal trade at every standardized-approach dealer is optimized.

3 Correlation Ambiguity and Information Inertia: What the Theory Says

The economics of maxmin behavior over correlations has been worked out formally. One important result from that literature is that worst-case evaluation can make the optimal market position locally independent of the signal.

In Illeditsch et al. (2021), investors trade a risky asset before and after observing a public signal about its payoff. The signal is known to be informative, in the sense that its correlation with fundamentals is positive, but *how* informative is ambiguous. Investors entertain a family




Figure 1: The scenario-max declines to reward hedging. For a single bucket with fixed inventory $W_1 > 0$ and a hedge of size h in an offsetting instrument, drawn for a base-metals bucket (MAR21.83 within-commodity correlation $\rho_M = 60\%$, so $\rho_L = 45\%$ and $\rho_H = 75\%$), the required charge $\bar{K}/W_1$ is shown against the hedge ratio h/W_1 under each correlation scenario. Because the book is hedged, the scenario-max coincides with the *low*-correlation curve (bold) at every hedge size. It is minimized at $h = \rho_L W_1 = 0.45 W_1$, the only hedge ratio the rule rewards, well short of the economic optimum $h = \rho_{\text{true}} W_1 = 0.60 W_1$ that minimizes true risk (the trough of the dashed medium curve). On the shaded interval $(\rho_L, \rho_{\text{true}}]$ additional, economically useful hedging *raises* required capital; and because $\rho_M < \frac{2}{3}$, a one-for-one hedge ($h = W_1$) is charged *above* the unhedged position. For the tighter oil-spread correlations ($\rho_M = 95\text{--}99\%$) the shape is the same but the band is narrow and full hedges remain credited.

of models indexed by a correlation parameter $\rho \in [\rho_a, \rho_b]$ and, following Gilboa and Schmeidler (1989), evaluate portfolios under the worst model in that interval. This is precisely the structure of MAR21.6, transplanted from a capital rule to a belief. There is an interval of correlation models, and a worst-case criterion over it.

The key insight is that a signal of ambiguous informativeness moves two things at once, in opposite directions. Interpreting bad news under a *high* correlation makes the conditional mean worse (the news really is about fundamentals) but makes the conditional variance smaller (an informative signal resolves more residual uncertainty). Interpreting the same news under a *low* correlation spares the mean but leaves more residual risk. For an ambiguity-averse investor, the worst-case model is whichever interpretation makes the *risk-return trade-off* worst. That trade-off is the conditional Sharpe ratio, and which interpretation minimizes it depends on the

news itself. For good news and for very bad news, one force dominates and the worst case sits at an endpoint of the interval. But there is an interior range of *bad* news in which the two forces exactly offset. As the signal deteriorates, the worst-case correlation slides continuously in a way that leaves the worst-case Sharpe ratio, and therefore the optimal portfolio, *unchanged*. Over this range the investor’s position is constant in the news. It coincides with the position of an investor who believes the signal carries no information at all, though everyone agrees that it does.

In equilibrium, the same flatness appears in prices. When the marginal investor is ambiguity averse in this sense, there is an interval of signal realizations, again concentrated on bad news, over which the market price does not depend on the signal. Public information arrives and is simply not impounded. Calibrations in the paper put the probability that a risky portfolio is unresponsive to bad news in the range of 8–20% for plausible parameterizations, rising with the width of the correlation interval $[\rho_a, \rho_b]$ and with the riskiness of the asset. The model’s cross-sectional predictions line up with a familiar family of anomalies. Underreaction and drift are stronger where ambiguity is plausibly larger (low analyst coverage, high forecast dispersion, high volatility), stronger for *systematic* than for idiosyncratic news, and absent for purely idiosyncratic signals. These patterns are consistent with the empirical literature on post-earnings-announcement drift.

Three features of the result deserve emphasis for what follows.

First, *inertia is not caused by pessimism; it is caused by the endogenous switching of the worst case*. A merely pessimistic investor, one who always used ρ_b , say, would respond smoothly to every signal. It is the fact that the identity of the worst-case model *changes with the news* that flattens the response. Any institution whose capital decisions are determined by a max over correlation models inherits this switching property mechanically.

Second, *the inertia region sits on bad news*. The offsetting mean-variance forces operate when news is adverse, while good news is processed normally. An economy of maxmin agents does not underreact symmetrically. It underreacts to market deterioration.

Third, *the mechanism needs no irrationality and no friction*. There are no transaction costs, no inattention, no agency problems in the model. Worst-case evaluation over an interval of correlations is enough. This matters for policy because it means the effect cannot be trained, disclosed, or competed away while the worst-case operator remains in place; it is the optimizing response to the operator itself.

The related literature reinforces the point from several directions. [Dow and Werlang \(1992\)](#) showed that ambiguity generates intervals of prices at which agents neither buy nor sell; [Illeditsch \(2011\)](#) showed that ambiguous information quality produces portfolio inertia and excess volatility; [Epstein and Schneider \(2008\)](#) showed that ambiguity about information quality makes agents treat bad news asymmetrically and demand premia for it; [Caballero and Krishna-](#)

murthy (2008) showed that Knightian uncertainty in intermediaries produces flight-to-quality and market freezes. Relative to this work, what matters here is the object of the ambiguity in Illeditsch et al. (2021). It is the *correlation* between the signal and the payoff, which is exactly the parameter MAR21.6 treats as Knightian.

4 From Capital Rule to Trading Behavior: Two Channels

A capital rule is not a preference. A bank under MAR21.6 may hold precise, probabilistically sophisticated beliefs about correlations and still face the scenario-max only as a constraint. The distinction matters less than it appears. The constraint channel reproduces the inertia mechanism on its own, and the beliefs channel operates on top of it.

4.1 The constraint channel: a stylized model

In this simplified model, a dealer desk with a mean-variance objective chooses net risk-weighted exposures W_1 and W_2 over two risk factors: $W_1 > 0$ is an inventory exposure the desk holds for market-making reasons, and W_2 is the freely chosen position in a correlated instrument, a hedge if $W_2 < 0$ and a directional position if $W_2 > 0$. Regulatory capital is charged at shadow cost $\kappa > 0$ per unit, and the SBM charge under correlation scenario c is the two-asset version of (1):

$$K(W_1, W_2; \rho_c) = \sqrt{W_1^2 + W_2^2 + 2\rho_c W_1 W_2}, \quad \rho_L < \rho_M < \rho_H, \quad (3)$$

with the binding charge $K(W) = \max_c K(W; \rho_c)$ as in (2). The desk observes a public signal s (macro news, an earnings surprise, a ratings action) that shifts the perceived expected return $\mu_2(s)$ on the second factor, and, writing $\eta > 0$ for the desk's coefficient of risk aversion, chooses W_2 to maximize

$$\mu_2(s) W_2 - \frac{\eta}{2} \text{Var}(W_1, W_2) - \kappa K(W_1, W_2). \quad (4)$$

Since $\partial K / \partial \rho = W_1 W_2 / K$, the high-correlation scenario binds whenever the book is *directional* ($W_1 W_2 > 0$) and the low-correlation scenario binds whenever the book is *hedged* ($W_1 W_2 < 0$). The switch occurs exactly at $W_2 = 0$, the boundary between hedging and position-taking. Now compute the marginal capital cost of the position as W_2 crosses that boundary from below (adding the last unit of hedge) and from above (adding the first unit of directional exposure):

$$\left. \frac{\partial K}{\partial W_2} \right|_{W_2 \uparrow 0} = \rho_L \text{sgn}(W_1), \quad \left. \frac{\partial K}{\partial W_2} \right|_{W_2 \downarrow 0} = \rho_H \text{sgn}(W_1). \quad (5)$$

The marginal cost of capital *jumps* by $\kappa(\rho_H - \rho_L)$ at the kink, and the desk's optimal overlay stays at zero on an interval of news.

Proposition 2 (Constraint-channel inaction region). *In problem (4) with fixed inventory $W_1 \neq 0$, true correlation ρ_{true} , and scenario-max charge $\bar{K} = \max_c K(\cdot; \rho_c)$, the optimal position is $W_2^*(s) = 0$, constant in the signal, for every s in the interval*

$$\mu_2(s) \in [\eta \rho_{\text{true}} W_1 + \kappa \rho_L \text{sgn}(W_1), \eta \rho_{\text{true}} W_1 + \kappa \rho_H \text{sgn}(W_1)]. \quad (6)$$

This interval is nondegenerate, and its width in expected value (μ_2) units is $\kappa(\rho_H - \rho_L)$, increasing in the shadow cost of capital κ and in the regulator’s scenario spread $\rho_H - \rho_L$.

The inaction is information inertia generated not by anyone’s beliefs but by the geometry of $\max_c K(\cdot; \rho_c)$. The endogenous switching of the binding scenario plays the same role that the endogenous switching of the worst-case belief plays in Illeditsch et al. (2021).²

Two comparative statics in Proposition 2 are relevant to bank regulators. First, the width $\kappa(\rho_H - \rho_L)$ is the product of a *policy constant* and an *equilibrium outcome*. At $\rho_M = 50\%$ the scenario interval runs from 37.5% to 62.5%, a 25-point band. The shadow cost κ , by contrast, is small when balance sheet capital is abundant, but spikes during times of market stress, when capital is scarce and internal reallocation is slow (He and Krishnamurthy, 2013; Du et al., 2018). The model therefore predicts that *inaction regions widen precisely when markets most need dealers to respond to news*. This is a procyclicality of price discovery, layered on top of the familiar procyclicality of capital itself. Section 4.2 quantifies both factors.

The second comparative static concerns *where* the kinks sit. In the two-asset model the kink is at $W_2 = 0$, but in the full SBM the max is taken portfolio-wide, so the switching locus is the set of book compositions at which two scenario aggregations are equal. That surface is high-dimensional, and every standardized-approach bank shares it, because every bank faces the same correlation matrices. Optimal books cluster *on* kinks (that is what inaction regions mean), so the rule concentrates many institutions’ portfolios near the same switching surface, in the same states of the world. Section 5 returns to what happens when they all come off it at once.

A related phenomenon happens with changes in the relative variance of the two assets. Proposition 2 varied the signal $\mu_2(s)$ and held the risk structure fixed. We now hold the signal fixed and vary the true volatility of the traded factor, keeping the true correlation unchanged. Because the capital charge is built on risk-weighted sensitivities with weights fixed by the regulator, local changes in volatility don’t change these weights, so the same jump at $W_2 = 0$ makes capital holdings constant even with changes in the risk of the assets in the portfolio.

²The microfoundations differ in an instructive way. In the belief-based model the inertia arises from a smooth interior tangency, with the worst-case correlation sliding so that Sharpe-ratio effects offset, while in the constraint version it arises at the kink of the max operator, in the spirit of Dow and Werlang (1992) and Illeditsch (2011). The observable implication is the same. In both there is an interval of news over which positions and quotes do not move.




Figure 2: Inaction in volatility. For the stylized desk of Section 4.1 with fixed inventory $W_1 = 1$, edge $\mu_2 = 1.08$, risk aversion $\eta = 1$, and shadow cost $\kappa = 0.8$, the chosen overlay W_2^* is drawn against the true volatility σ_2 of the traded factor, holding the true correlation fixed at $\rho_{\text{true}} = 0.60$. Correlations are the base-metals bucket ($\rho_M = 0.60$, so $\rho_L = 0.45$ and $\rho_H = 0.75$). Under the scenario-max (bold) the desk runs a directional overlay at low volatility and a hedge at high volatility, but over the band $\sigma_2 \in [0.80, 1.20]$ of Proposition 3 it sits at $W_2^* = 0$, so both its position and the capital it consumes are invariant to the underlying's volatility. A desk facing a smooth aggregation of the same three scenarios (dashed) rotates continuously through zero. Generated by `fig_inaction_vol.py`.

Proposition 3 (Inaction in volatility). *Let σ_2 be the true volatility of factor 2 in units of its prescribed risk weight, so that $\text{Var}(W_1, W_2) = W_1^2 + \sigma_2^2 W_2^2 + 2\rho_{\text{true}} \sigma_2 W_1 W_2$ while the regulatory charge $\bar{K}$ is unchanged. Fix the news μ_2 and the true correlation ρ_{true} . For $W_1 > 0$ position in the traded asset is $W_2^* = 0$ for every*

$$\sigma_2 \in \left[\frac{\mu_2 - \kappa \rho_H}{\eta \rho_{\text{true}} W_1}, \frac{\mu_2 - \kappa \rho_L}{\eta \rho_{\text{true}} W_1} \right], \quad (7)$$

an interval of width $\kappa(\rho_H - \rho_L)/(\eta \rho_{\text{true}} W_1)$.

Figure 2 draws the base-metals case. Below the band the desk holds a positive position, and above it a hedge, but across a 50% rise in the underlying's volatility it holds neither, and the regulatory capital it consumes is flat throughout. A desk facing the smooth aggregation of Section 6 rotates continuously from overlay to hedge over the same range. What the rule freezes is not merely the response to news but the response to volatility, with the correlation

the desk is trading against held fixed the whole time. In this case positions, and the capital behind them, decouple from the risk they are meant to track.

4.2 Calibrating the inaction region

Proposition 2 makes the width of the inaction region the product of two objects, $\kappa(\rho_H - \rho_L)$, a regulatory constant and a shadow price. The next example seeks to quantify these.

The scenario spread $\rho_H - \rho_L$ is fixed by [MAR21.6] and is not small. Table 2 evaluates it at representative prescribed correlations. It vanishes for near-perfect hedges (the ρ_H cap at 100% compresses the band) and for the very low equity correlations, but reaches a maximum of *40 percentage points* at $\rho_M = 0.80$, precisely the range of many rates-curve, cross-currency, and cross-bucket correlations on which dealer hedges rely. Formally, $\rho_H - \rho_L = 0.5\rho_M$ for $\rho_M \leq 0.8$ and $2(1 - \rho_M)$ for $\rho_M \geq 0.8$, peaking at $\rho_M = 0.80$.

Table 2: Prescribed correlation-scenario spread $\rho_H - \rho_L$ at representative SBM parameters (MAR21; Appendix A). The spread is the policy constant that scales the inaction region of Proposition 2.

Parameter (bucket)	ρ_M	ρ_L	ρ_H	$\rho_H - \rho_L$
Equity, large-cap advanced	0.25	0.19	0.31	0.13
Credit, different-name (CSR)	0.35	0.26	0.44	0.18
GIRR cross-currency (γ)	0.50	0.38	0.63	0.25
Base metals; FX cross-bucket	0.60	0.45	0.75	0.30
<i>Worst case</i>	<i>0.80</i>	<i>0.60</i>	<i>1.00</i>	<i>0.40</i>
GIRR tenor (2y–10y)	0.89	0.77	1.00	0.23
Energy, liquid combustibles	0.95	0.90	1.00	0.10

The shadow cost κ is an equilibrium object. It is the marginal cost to the desk of the regulatory capital a trade consumes, and the intermediary-asset-pricing literature shows it is both large and sharply countercyclical. Measured through covered-interest-parity bases and other balance-sheet-rent proxies, the shadow price of dealer balance sheet widens from a few tens of basis points in calm markets to well over a hundred in stress (Du et al., 2018; He and Krishnamurthy, 2013). A several-fold increase in κ during stress translates, through Proposition 2, into a several-fold widening of the inaction region.

Combining the two, let $\phi \equiv \kappa(\rho_H - \rho_L)/\sigma_\mu$ be the width of the inaction band relative to the dispersion σ_μ of news-implied edges. In an illustrative cross-section of dealers holding a common systematic inventory of heterogeneous size, the fraction whose optimal position is frozen at the kink rises from about 10% when $\phi \approx 0.25$ (abundant balance sheet) to about one-third when $\phi \approx 1$ (stress, with κ several times higher). That range brackets the 8–20% unresponsiveness frequencies that Illeditsch et al. (2021) obtain from the belief-based model under plausible parameterizations; Appendix B.4 reports the underlying cross-section. Because

the news that moves these desks is systematic and the scenario matrices are common, the freezes are not independent. Dealers enter and leave the inaction region together, which is the synchronization channel Section 5 develops.

4.3 The beliefs channel

The constraint channel operates even on a bank whose risk managers scoff at Knightian uncertainty. But there are institutional reasons to expect the scenario set to leak from the constraint into beliefs. Model-risk management frameworks and prudent-valuation rules push banks to maintain admissible *sets* of parameterizations rather than point estimates, and the regulator’s published scenario band is the natural candidate for that set, and a supervisory-examination-proof one. A desk head who marks correlation risk inside the regulator’s band will rarely be second-guessed; one who insists on a point estimate outside it bears career risk in one direction only. Where that happens, the bank’s *own* risk assessment becomes maxmin over an interval of correlations, and the full mechanism of Illeditsch et al. (2021) applies directly.

4.4 When the mapping fails

The results to this point rely on three assumptions. The first is that the SBM charge is *binding* or near-binding, or that internal capital budgeting charges desks for regulatory capital at a positive shadow price, which is standard practice at large dealers, though with heterogeneous intensity. The second is that the bank is the *marginal* investor in the relevant market, which is plausible for dealer-intermediated OTC markets (credit, rates, FX derivatives, index volatility) and much less so for liquid equities where unconstrained investors set prices. The third is that the desk’s response to news is not dominated by other frictions (position limits, VaR limits, franchise considerations) that would mask the capital kink. Where these fail, the scenario-max is a pricing anomaly rather than a market-functioning issue. However, under periods of market stress, their effect could be pronounced. The empirical footprint proposed in Section 6 is designed to discriminate between the two cases.

5 Likely Consequences

Five families of consequences follow from the mechanism of Section 4, ordered roughly from most to least direct. Each is stated as a prediction about observables.

Implication 1 (Underreaction to systematic bad news in dealer-intermediated markets). *Over intervals of adverse news, constrained dealers’ quotes and positions do not move; where dealers are marginal, prices underreact and subsequently drift. The theory predicts the effect is (a) asymmetric and concentrated on bad news; (b) increasing in the riskiness of the exposure and*

the width of the scenario band; and (c) specific to systematic factors. News about idiosyncratic exposures barely interacts with the correlation matrix and generates no inertia, while news that shifts whole-bucket or cross-bucket exposures generates the most. Post-news drift in credit indices, swap spreads, and index volatility around macro releases, but not in single names on idiosyncratic news, is the signature to look for.

Implication 2 (Cliff-edge hedging). *Inside the inaction region, desks do not adjust hedges as conditions deteriorate; when news finally pushes the desired position through the far edge of the region, the accumulated adjustment executes at once. Hedging flow becomes lumpy, quiet and then discontinuous. For counterparties and for the markets in which the hedges execute, this converts a smooth flow of small trades into occasional large ones. Delayed, lumpy adjustment is more expensive to execute, more visible to predatory strategies, and more likely to move prices sharply than the continuous rebalancing it replaces.*

Implication 3 (Synchronized inaction, synchronized action). *Because the scenario matrices are identical across banks and jurisdictions, the switching surfaces of Section 4.1 are common to (almost) all banks. Books managed against the same kinks cluster near the same configurations; the same macro news moves many institutions across the same boundary in the same direction at the same time. Under internal models, worst cases were heterogeneous, since one bank's stress scenario depended on its position which was almost surely different from another's. In this world, inaction regions, to the extent they existed, were idiosyncratic and offsetting. The SBM homogenizes them. The system-level concern is not that any one bank is slow to react, but that all standardized-approach banks are slow together and then move together, amplifying the one-sided markets and liquidity holes the correlation scenarios were meant to guard against.*

Implication 4 (An uncertainty premium in intermediation prices). *Ex ante, maxmin evaluation over the correlation band commands compensation. In Illeditsch et al. (2021), prices embed an uncertainty premium that increases in risk aversion, in volatility, and in the width of the ambiguity interval, over and above the standard risk premium. The constraint analogue is a capital-cost wedge proportional to $\kappa(\rho_H - \rho_L)$ priced into bid-ask spreads and into the cost of hedging services banks sell to the real economy, to corporates hedging rates and FX and pension funds hedging duration. The premium concentrates around information events, consistent with evidence that a large share of excess returns is earned around announcements. This is a real, ongoing cost of the worst-case functional form, paid whether or not the stress it insures against arrives; against it must be set the (also real) benefit of robustness to correlation breakdown.*

Implication 5 (Procyclical price discovery and migration). *Because the inaction region scales with the shadow cost of capital, price discovery through dealer intermediation degrades endogenously in stress, just when public signals are most valuable and when the volume of adverse news is largest. This has two potential solutions, First, activity migrates toward institutions*

outside the constraint, to hedge funds and principal trading firms. This restores some price responsiveness but relocates risk to more lightly supervised balance sheets and makes intermediation capacity flightier. Second, banks themselves face an incentive to restructure exposures to sit away from switching surfaces (or to relocate them to the shrinking internal-models perimeter), an optimization that consumes real resources and produces books shaped by the geometry of the rule rather than by economic risk.

The correlation scenarios were designed to make individual banks robust to a *estimation* failure, correlations moving away from what had previously been thought to be true. The functional form chosen to deliver that robustness creates a *behavioral* failure mode, synchronized unresponsiveness to bad news followed by synchronized adjustment. That response operates at the level of the market rather than the bank, and it is invisible in any single institution’s capital adequacy. Sitting still is not a capital-adequacy failure, and no supervisory metric records it. The cost shows up instead in the informativeness of prices and the smoothness of hedging flows, which market-risk supervision does not currently measure.

6 What Supervisors and Standard-Setters Could Do

None of the above justifies abandoning conservatism about correlations. The regulator’s ambiguity is genuine, and a worst-case operator is a coherent response to it. The policy question is narrower and separable. It is whether the same robustness can be bought with a functional form that does not manufacture inaction regions, and whether the phenomenon can be measured while the current form is in force. What follows are three policy proposals that could help to ameliorate the problems just discussed. These are ordered in from easiest to hardest to implement.

6.1 Proposal 1: Collect all three scenarios

Every SBM bank already computes all three scenario aggregations at every reporting date. The max operation discards two of them, and the discarded numbers are the diagnostic ones. Supervisors could collect three items per bank and per risk class. The first is which scenario binds. The second is the *scenario spread*, the gap between the largest and second-largest aggregation, which measures distance to the switching surface. The third is the frequency of binding-scenario switches between reporting dates. System-wide, the concentration of banks near the same switching surface at the same time would be useful in estimating the synchronization risk of Section 5, and it is observable at essentially zero incremental reporting cost. A Pillar 3 disclosure of the binding scenario would extend the same visibility to counterparties and researchers.

6.2 Proposal 2: Exploit the staggered rollout

FRTB went live in Canada and Japan in 2024, in the UK in 2025, and in the EU in January 2026, while the US market-risk rule remains unfinalized. Identical products trade in markets intermediated by dealers on different sides of these boundaries. The theory makes sharp difference-in-differences predictions. After roll out, dealer-intermediated prices in adopting jurisdictions should show increased post-announcement drift on adverse systematic news, lumpier hedging flow, and wider announcement-window spreads, relative to the same instruments where the marginal dealer is not yet subject to the scenario-max. The effects should concentrate in periods of expensive balance sheet capital. Supervisory research units and the academic community could run this test now, while the cross-jurisdictional variation still exists.

6.3 Proposal 3: Use smooth aggregation

The inertia mechanism lives in the non-differentiability of $\max_c K(\cdot; c)$, not in its conservatism. Consider a smooth aggregation of the same three scenario figures, for instance a log-sum-exp (“soft-max”),

$$K_\beta(\theta) = \frac{1}{\beta} \log \left(\sum_c w_c e^{\beta K(\theta; c)} \right), \quad (8)$$

which converges to the hard max as $\beta \rightarrow \infty$ and can be calibrated (via β and weights w_c , or an added constant) to deliver the *same average capital* as the current rule. Such an aggregation preserves worst-case sensitivity while making the marginal capital cost of every trade continuous in the portfolio. The inaction interval (6) collapses; the incentive to cluster on switching surfaces disappears; conservatism survives as a smooth penalty rather than a cliff. Decision theory offers principled versions of the same move. The smooth-ambiguity and multiplier-preference families (Klibanoff et al., 2005; Hansen and Sargent, 2001) were developed precisely to retain robustness while restoring differentiability. A targeted revision of MAR21.6’s aggregation using the form above would not need to change the risk weights, buckets, or correlations already in the regulation.

7 Conclusion

The Fundamental Review of the Trading Book made two large choices. It moved the measurement of tail risk from value-at-risk to expected shortfall. And less deliberately, through the incentives embedded in the internal-models tests, it moved most of the world’s trading books onto a standardized, regulator-parameterized risk model. This paper has examined one line of that model, the requirement that capital be the worst of three correlation scenarios. The line encodes a specific and well-understood decision-theoretic framework, maxmin evaluation over a set of correlation models. That theory shows that agents optimizing in that framework stop

responding to information over intervals of bad news, because the identity of the worst case shifts endogenously with the news itself. In [Illeditsch et al. \(2021\)](#) this produces prices that fail to impound public information. Transplanted into a capital constraint, it produces dealers whose positions, hedges, and quotes are flat over ranges of deteriorating conditions, with inaction regions that widen when capital is scarce and that sit at the same portfolio configurations for every bank in the regime.

The scenario-max is a reasonable answer to a real problem, and the magnitude of its behavioral cost is an open empirical question. This analysis establishes a way to think about the costs of this choice, ways that its costs could be measured and potentially alleviated.

A Prescribed SBM correlation parameters

This appendix reproduces the delta correlation parameters of the sensitivities-based method used in the numerical illustration of Section 2.1. All values are those prescribed by the Basel Framework ([Basel Committee on Banking Supervision, 2019a](#)), chapter MAR21; paragraph numbers are given in brackets. These parameters are not a national choice. The European Union (CRR3), the United Kingdom, Canada, and Japan have adopted them essentially verbatim, so the “worst case” embedded in the scenario-max is a common constant across jurisdictions. The only material national discretions lie on the risk-weight side (e.g. dividing the risk weight by $\sqrt{2}$ for Committee-specified currencies and currency pairs [MAR21.44, MAR21.88]), not in the correlations.

Correlation scenarios [MAR21.6]. Every parameter ρ and γ below is the **medium** value. The **high** scenario replaces it by $\min(1.25\rho, 100\%)$; the **low** scenario replaces it by $\max(2\rho - 100\%, 75\% \times \rho)$. The SBM requirement is the largest of the three scenario totals [MAR21.7].

GIRR (general interest rate risk) [MAR21.45–21.50]. Within a currency: same tenor, different curves = 99.90%; different tenor, same curve as in Table 3; different tenor and different curves, the Table 3 value $\times 99.90\%$; inflation vs. yield curve = 40%; cross-currency basis vs. other curves = 0%. Across currencies (buckets), $\gamma_{bc} = 50\%$. The same-curve tenor correlations in Table 3 are generated by

$$\rho_{kl} = \max\left(e^{-\theta |T_k - T_l| / \min(T_k, T_l)}, 40\%\right), \quad \theta = 3\%.$$

Equity [MAR21.78–21.80]. Within-bucket spot–spot correlation ρ_{kl} : 15% large-cap emerging markets (buckets 1–4), 25% large-cap advanced (buckets 5–8), 7.5% small-cap emerging (bucket 9), 12.5% small-cap advanced (bucket 10), 80% index buckets (12, 13); a spot–repo pair for the same name is 99.90%; the “other” bucket 11 is aggregated by simple sum. Across buckets, $\gamma_{bc} = 15\%$ among buckets 1–10, 0% with bucket 11, and 75% between the two index buckets.

Table 3: GIRR within-bucket tenor correlations ρ_{kl} (same curve), in percent [MAR21.46, Table 2].

	0.25y	0.5y	1y	2y	3y	5y	10y	15y	20y	30y
0.25y	100.0	97.0	91.4	81.1	71.9	56.6	40.0	40.0	40.0	40.0
0.5y	97.0	100.0	97.0	91.4	86.1	76.3	56.6	41.9	40.0	40.0
1y	91.4	97.0	100.0	97.0	94.2	88.7	76.3	65.7	56.6	41.9
2y	81.1	91.4	97.0	100.0	98.5	95.6	88.7	82.3	76.3	65.7
3y	71.9	86.1	94.2	98.5	100.0	98.0	93.2	88.7	84.4	76.3
5y	56.6	76.3	88.7	95.6	98.0	100.0	97.0	94.2	91.4	86.1
10y	40.0	56.6	76.3	88.7	93.2	97.0	100.0	98.5	97.0	94.2
15y	40.0	41.9	65.7	82.3	88.7	94.2	98.5	100.0	99.0	97.0
20y	40.0	40.0	56.6	76.3	84.4	91.4	97.0	99.0	100.0	98.5
30y	40.0	40.0	41.9	65.7	76.3	86.1	94.2	97.0	98.5	100.0

Commodity [MAR21.83–21.85]. Within a bucket, $\rho_{kl} = \rho^{(\text{cty})} \cdot \rho^{(\text{tenor})} \cdot \rho^{(\text{basis})}$, where $\rho^{(\text{cty})} = 100\%$ for the same commodity and otherwise the bucket value in Table 4; $\rho^{(\text{tenor})} = 99\%$ for different tenors; and $\rho^{(\text{basis})} = 99.90\%$ for different delivery locations. Across buckets, $\gamma_{bc} = 20\%$ among buckets 1–10 and 0% with bucket 11.

Table 4: Commodity within-bucket correlations $\rho^{(\text{cty})}$ for distinct commodities [MAR21.83, Table 12].

Bucket	Commodity bucket	$\rho^{(\text{cty})}$
1	Energy – solid combustibles	55%
2	Energy – liquid combustibles (crude, products)	95%
3	Energy – electricity and carbon trading	40%
4	Freight	80%
5	Metals – non-precious (base metals)	60%
6	Gaseous combustibles	65%
7	Precious metals (incl. gold)	55%
8	Grains and oilseed	45%
9	Livestock and dairy	15%
10	Softs and other agriculturals	40%
11	Other commodity	15%

Foreign exchange [MAR21.87–21.89]. A single risk factor per currency pair (no within-bucket correlation); the cross-bucket correlation is $\gamma_{bc} = 60\%$ uniformly.

Credit spread, non-securitization [MAR21.54–21.57]. Within a bucket, $\rho_{kl} = \rho^{(\text{name})} \cdot \rho^{(\text{tenor})} \cdot \rho^{(\text{basis})}$, with $\rho^{(\text{name})} = 35\%$ for different issuers (buckets 1–15; 80% for the covered-bond buckets 17–18), $\rho^{(\text{tenor})} = 65\%$ for different tenors, and $\rho^{(\text{basis})} = 99.90\%$ for different curves (e.g. bond vs. CDS). Across buckets, $\gamma_{bc} = \gamma^{(\text{rating})} \cdot \gamma^{(\text{sector})}$, with $\gamma^{(\text{rating})} = 50\%$ for different rating categories (IG vs. HY/NR) and $\gamma^{(\text{sector})}$ ranging from 0% to 75% by sector pair

[MAR21.57, Table 5].

B Replication of the numerical example

This appendix records the full inputs and intermediate outputs behind Table 1. The correlation parameters are exactly those of Appendix A; the sensitivities are the illustrative net risk-weighted values in Tables 5–7. Two scripts reproduce every numerical result in the paper. `sbm_example.py` produces the three-bank charges of Table 1 and their risk-class breakdown in Table 8, the prescribed scenario spreads of Table 2, the hedge-recognition property of Proposition 1 and Corollary 1 drawn in Figure 1, and the inaction-fraction cross-section of Section 4.2 tabulated in Appendix B.4. `fig_inaction_vol.py` produces the inaction-in-volatility experiment of Figure 2.

B.1 Aggregation procedure

For each correlation scenario $c \in \{\text{low}, \text{med}, \text{high}\}$, that is, the medium parameters of Appendix A rescaled by the rule in [MAR21.6], the capital charge is built in three steps, following (1):

1. *Bucket charge.* For each bucket b ,

$$K_b = \sqrt{\max\left(0, \sum_k \text{WS}_k^2 + \sum_k \sum_{l \neq k} \rho_{kl}^{(c)} \text{WS}_k \text{WS}_l\right)}.$$

2. *Risk-class charge.* Aggregate buckets via

$$\sqrt{\max\left(0, \sum_b K_b^2 + \sum_b \sum_{d \neq b} \gamma_{bd}^{(c)} S_b S_d\right)}, \quad S_b = \sum_k \text{WS}_k,$$

using the [MAR21.4] fallback $S_b^* = \max(\min(S_b, K_b), -K_b)$ in place of S_b where the radicand is negative.

3. *Scenario total.* Sum the risk-class charges across risk classes (no cross-class correlation).

The SBM requirement is the largest of the three scenario totals [MAR21.7], as in (2). All sensitivities WS are net risk-weighted, in \$mm.

B.2 Positions

Table 5: Bank A (Community Bank). A single GIRR bucket (USD), net long duration, a directional book; within-bucket correlations follow Table 3.

Risk class	Risk factor (tenor)	WS
GIRR (USD)	0.5y	−8
	2y	+35
	5y	+50
	10y	+28
	30y	+12

Table 6: Bank B (Commodity Bank). Oil calendar/crack spreads (bucket 2), a natural-gas calendar spread (bucket 6), a base-metals spread (bucket 5), a rates financing hedge, and an FX position.

Risk class (γ)	Bucket (ρ within)	Risk factor	WS
Commodity ($\gamma=20\%$)	Energy–liquid comb. ($\rho^{(\text{cty})}=95\%$)	WTI crude, prompt	+240
		WTI crude, deferred	−200
		Gasoil, near	+150
		Gasoil, far	−120
	Gaseous / nat. gas (same cmdty, 99%)	Nat. gas, near	+120
		Nat. gas, far	−110
	Base metals ($\rho^{(\text{cty})}=60\%$)	Copper	+80
		Aluminium	−70
GIRR (USD)	single bucket	2y	+60
		10y	−55
FX ($\gamma=60\%$)	EUR/USD	—	+40
	USD/JPY	—	−30

Table 7: Bank C (Atlantic Bank). Long-short equity relative value across three buckets, a rates curve/basis book, a four-pair FX book, and an investment-grade credit spread pair.

Risk class (γ)	Bucket (ρ within)	Risk factor	WS
Equity ($\gamma=15\%$)	Large-cap advanced ($\rho=25\%$)	single name (long)	+220
		sector index (short)	-190
	Large-cap emerging ($\rho=15\%$)	long	+100
		short	-80
	Small-cap advanced ($\rho=12.5\%$)	long	+45
		short	-30
GIRR (USD)	single bucket	2y	+160
		3y	-150
		5y	+90
		10y	-80
FX ($\gamma=60\%$)	USD/EUR	—	+90
	USD/GBP	—	-70
	USD/JPY	—	+50
	USD/CAD	—	-40
CSR non-sec	IG bucket ($\rho^{(\text{name})}=35\%$)	issuer A (long)	+70
		issuer B (short)	-55

B.3 Intermediate results

Table 8 reports the risk-class charges and scenario totals under each correlation scenario. The binding total per bank (in bold) is the SBM requirement reported in Table 1: for Bank A the high scenario, for Banks B and C the low scenario.

Table 8: Risk-class charges and scenario totals (\$mm). The SBM requirement is the largest scenario total (bold).

Bank	Risk class	Low	Medium	High
A	GIRR (USD)	113.0	115.3	117.5
	Total	113.0	115.3	117.5
B	Commodity	120.8	107.8	93.0
	GIRR (USD)	39.0	27.8	5.0
	FX	37.7	32.6	26.5
	Total	197.4	168.1	124.5
C	Equity	293.9	283.9	273.5
	GIRR (USD)	61.3	45.6	20.0
	FX	99.1	85.9	70.4
	CSR (credit)	76.8	72.3	67.5
	Total	531.1	487.7	431.3

B.4 Calibration of the inaction region

The width of the inaction region in Proposition 2 is $\kappa(\rho_H - \rho_L)$, the product of the shadow cost of capital and the prescribed scenario spread. Table 2 tabulates the spread across representative buckets. To translate the width into a frequency of inaction, consider a cross-section of dealers who share a common long systematic inventory of heterogeneous size, $W_1 \sim \text{Uniform}[0.5, 1.5]$, and who receive a common news shock $\mu \sim \mathcal{N}(\mu_0, \sigma_\mu^2)$ centered at $\mu_0 = \rho_{\text{true}} \mathbb{E}[W_1]$. Normalizing $\kappa = 1$ and $\eta = 1$, a dealer is frozen ($W_2^* = 0$) when its news lands in the inventory-dependent band $\mu \in [\rho_{\text{true}} W_1 + \rho_L, \rho_{\text{true}} W_1 + \rho_H]$ of Proposition 2. The frozen fraction depends on the single statistic $\phi \equiv \kappa(\rho_H - \rho_L)/\sigma_\mu$, the band width relative to the dispersion of news. Table 9 reports it at the base-metals correlation $\rho_M = 0.60$ (so $\rho_L = 0.45$, $\rho_H = 0.75$), from a Monte Carlo of 2×10^5 dealers.

Table 9: Illustrative fraction of dealers frozen at the kink, by inaction-band width $\phi = \kappa(\rho_H - \rho_L)/\sigma_\mu$ (base-metals correlations, $\rho_M = 0.60$; Monte Carlo of 2×10^5 dealers, $W_1 \sim \text{Uniform}[0.5, 1.5]$). Produced by `sbm_example.py`.

$\phi = \kappa(\rho_H - \rho_L)/\sigma_\mu$	0.25	0.50	1.00
Balance sheet	abundant	intermediate	scarce
Dealers frozen ($W_2^* = 0$)	9.9%	19.0%	33.2%

As balance sheet moves from abundant ($\phi \approx 0.25$) to scarce ($\phi \approx 1$), the frozen fraction rises from about one in ten to about one in three, bracketing the 8–20% unresponsiveness frequencies that [Illeditsch et al. \(2021\)](#) obtain from the belief-based model. The same script verifies the hedge-monotonicity of Proposition 1 numerically. The inaction-in-volatility experiment of Figure 2, which holds the news and correlation fixed and varies the traded factor’s true volatility, is produced by `fig_inaction_vol.py`.

C Proofs

Proof of Proposition 1. Since $\partial K/\partial \rho = -W_1 h/K < 0$ for $h > 0$, the maximum of $K(h; \cdot)$ over $\{\rho_L, \rho_M, \rho_H\}$ is attained at ρ_L , giving (i). Hence $\partial \bar{K}/\partial h = (h - \rho_L W_1)/\bar{K}$, which is negative for $h < \rho_L W_1$ and positive for $h > \rho_L W_1$; the first-order condition yields $h^* = \rho_L W_1$ and $\bar{K}(h^*) = W_1 \sqrt{1 - \rho_L^2}$. Under a single correlation ρ_{true} the variance-minimizing hedge ratio is ρ_{true} , so on $(\rho_L W_1, \rho_{\text{true}} W_1]$ true variance is still falling while $\bar{K}$ is already rising. $\square$

Proof of Proposition 2. The marginal capital cost is $\partial K/\partial W_2 = (W_2 + \rho W_1)/K$, which at $W_2 = 0$ equals $\rho \operatorname{sgn}(W_1)$ since $K = |W_1|$ there. The binding correlation switches from ρ_L on the hedged side ($W_1 W_2 < 0$) to ρ_H on the directional side ($W_1 W_2 > 0$), so $\kappa \partial \bar{K}/\partial W_2$ has a jump discontinuity at $W_2 = 0$ with one-sided values $\kappa \rho_L \operatorname{sgn}(W_1)$ and $\kappa \rho_H \operatorname{sgn}(W_1)$. The rest of (4) is smooth and strictly concave, with marginal value $\mu_2(s) - \eta \rho_{\text{true}} W_1$ at $W_2 = 0$. Hence $W_2^* = 0$ is optimal exactly when that smooth marginal value lies between the two one-sided capital costs, which is (6); the width is $\kappa(\rho_H - \rho_L)$. $\square$

Proof of Proposition 3. The charge and its one-sided derivatives $\kappa \rho_L$ and $\kappa \rho_H$ at $W_2 = 0$ are built from fixed risk weights and do not depend on σ_2 . The smooth marginal value at $W_2 = 0$ is now $\mu_2 - \eta \rho_{\text{true}} \sigma_2 W_1$. By the argument of Proposition 2, $W_2^* = 0$ is optimal exactly when that value lies in $[\kappa \rho_L, \kappa \rho_H]$, which rearranges to (7). $\square$

References

- Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, 2010.
- Basel Committee on Banking Supervision. Minimum capital requirements for market risk. Technical report, Bank for International Settlements, January 2019a. Rev. Consolidated in the Basel Framework as MAR20–MAR23.
- Basel Committee on Banking Supervision. Explanatory note on the minimum capital requirements for market risk. Technical report, Bank for International Settlements, January 2019b.

- Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- Ricardo J. Caballero and Arvind Krishnamurthy. Collective risk management in a flight to quality episode. *Journal of Finance*, 63(5):2195–2230, 2008.
- Jón Daníelsson. The emperor has no clothes: Limits to risk modelling. *Journal of Banking & Finance*, 26(7):1273–1296, 2002.
- James Dow and Sérgio Ribeiro da Costa Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1):197–204, 1992.
- Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957, 2018.
- Larry G. Epstein and Martin Schneider. Ambiguity, information quality, and asset pricing. *Journal of Finance*, 63(1):197–228, 2008.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.
- Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–770, 2013.
- Philipp K. Illeditsch. Ambiguous information, portfolio inertia, and excess volatility. *Journal of Finance*, 66(6):2213–2247, 2011.
- Philipp K. Illeditsch, Jayant V. Ganguli, and Scott Condie. Information inertia. *Journal of Finance*, 76(1):443–479, 2021.
- International Swaps and Derivatives Association and EY. Fundamental review of the trading book: Internal models approach adoption. Technical report, ISDA, July 2024.
- Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.
- KPMG. Fundamental review of the trading book: An overview. Technical report, KPMG, July 2025. URL <https://kpmg.com/in/en/insights/2025/07/fundamental-review-of-the-trading-book-an-overview.html>. Notes that banks globally are preferring the standardised approach over the internal model approach for market risk capital, owing to its lower complexity and operational burden.