

A Supervisory Diagnostic for Sensitivities-Based Capital Requirements

Scott Condie*

July 2026

Abstract

Under the FRTB’s sensitivities-based method (SBM), the covariance matrix that sets trading-book capital is a regulatory instrument. This paper shows that in such a setting, common indicators of regulatory tightness (shadow prices on capital, hurdle rates, headroom over requirements) can point the wrong way. In a static equilibrium in which banks hold capital against a delta-only SBM charge, the shadow cost of the requirement is hump-shaped in the rule’s stringency while the pricing distortions it creates rise monotonically, so measured tightness is slackest where its effect on prices is largest. This paper proposes an alternative statistic supervisors could watch, the *sensitivity gap*, assembled from the weighted sensitivities banks already report and requiring no new data collection. This statistic rises monotonically with the distortion. Theoretically, the model shows that expected returns follow a two-beta CAPM in which identically risky exposures price differently according to how well the regulatory factor list covers them. Mathematically, the wedge between regulatory and true risk splits into a calibration component the regulator can remove and a coverage component that no calibration of risk weights can remove. A recalibration counterfactual and a welfare calculation measure the two components.

Keywords: FRTB; sensitivities-based method; capital requirements; CAPM; intermediary asset pricing; regulatory risk weights.

JEL codes: G12, G21, G28.

*Department of Economics, Brigham Young University, 2146 West View Building, Provo, UT 84602. Email: scott_condie@byu.edu.

1 Introduction

Under the Basel Committee’s Fundamental Review of the Trading Book (FRTB), the standardized approach was provided to serve as a backstop for banks too small to run internal models. Facing the profit-and-loss attribution test and the non-modelable risk factor framework, most large trading banks have declined to seek internal-model approval. Industry survey evidence indicates that institutions holding roughly 85% of trading-book capital under internal models plan internal-model coverage of only 15–40% under the new regime ([International Swaps and Derivatives Association and EY, 2024](#)). The regime is arriving jurisdiction by jurisdiction. It is already live in Canada, Japan, and several other Basel Committee member countries. Following relief measures adopted in June 2026, it will go live in the European Union from January 2027. The United Kingdom plans to adopt in 2027–28. Recently the discussion has reopened in the United States with the timeline unresolved ([Basel Committee on Banking Supervision, 2026](#); [European Commission, 2026](#)). The landscape of bank regulatory regimes seems to be shifting toward a world where the marginal dealer in most of the world’s fixed income, equity, and derivatives markets uses a sensitivities-based method (SBM) of risk capital assessment. The SBM is a parametric portfolio risk model. Banks supply net sensitivities to a prescribed list of risk factors and the regulator supplies the risk weights, the correlations, and the list itself.

This paper asks two questions about such a rule. The first is what it does to asset prices. That question is interesting because the SBM’s risk model may differ from the true nature of underlying risk in at least two distinct ways. The first is calibration. Its weights and correlations are statutory, so they adjust only slowly to changing market conditions. The second is coverage. The prescribed factor list is finite, and risks that fall outside it carry no charge at all, however large they are. The same is true of risks that the netting inside the sensitivity calculation removes. When the capital standard is binding then, it imposes a portfolio-level pressure on asset returns given the differences between bank-perceived risk and regulator perceived risk. This paper gives a framework for understanding how this affects asset prices. The second question is whether the officials who supervise the requirement can see any of this happening. The natural indicators of regulatory tightness (internal shadow prices on capital, desk-level hurdle rates, reported headroom over requirements) all track the shadow cost of the constraint. The paper’s central result is that the

shadow cost gives an incomplete characterization of the market impacts of the regulation, and it proposes an additional metric, computed from data already being generated in the SBM calculation, to give a more full picture.

The paper answers these questions in a deliberately transparent setting: a static CAPM economy with two classes of investors. Traditional investors solve the usual mean–variance problem. Banks solve the same problem subject to a stylized capital requirement. Their riskless asset holdings must be at least as large as a delta-only SBM charge, computed as the square root of a quadratic form in the bank’s portfolio under a regulatory covariance matrix that the prescribed risk weights and correlations define. That matrix covers only the regulated factors, which are fewer than the factors that actually drive payoffs, so it is rank-deficient relative to the true covariance of returns. The model thus allows for the study of errors in both calibration and coverage. Given the nature of the model, prices can be solved for in equilibrium, and portfolios can be characterized.

Five sets of results follow. The first, and the paper’s main result, concerns the diagnostics supervisors use. The model shows that the shadow cost of the capital requirement is not monotone in the stringency of the rule. This is proven in a single-factor economy and demonstrated numerically in the calibrated multi-asset equilibrium. As risk weights scale up, the shadow cost of the requirement rises, peaks, and then falls, because banks shrink the charged book until a marginal dollar of capital buys little. The pricing distortions rise monotonically throughout, reaching roughly 200 basis points on equity exposure in the calibration even as measured tightness declines. A supervisor watching shadow prices might conclude that the requirement had stopped mattering where its effect on prices is largest. The model supplies a replacement, the *sensitivity gap*: the distance between the banking sector’s factor-level weighted sensitivities and their frictionless benchmark. The gap is proportional to the same equilibrium object that drives every pricing distortion, it rises monotonically through the region where the shadow cost falls, and its data requirement is zero. Its inputs are the weighted sensitivities of the SBM calculation itself, reported at each risk factor by every standardized-approach bank, and Section 5.3 shows how a supervisor assembles the indicator from data the regime already collects. A companion indicator, the family of coverage-basis spreads, is observable daily.

The second set of results concerns the prices behind the sensitivity gap. In equilibrium, every

asset’s capital charge is levied in proportion to its total claim on bank balance sheets, defined as funding plus marginal capital charge, at a uniform rate equal to the shadow cost of the requirement scaled by the banking sector’s share of the economy’s risk-bearing capacity. Expected returns obey a two-beta CAPM in which a regulatory beta, computed under the regulator’s covariance matrix rather than the true one, is priced alongside the market beta, and even zero-beta assets earn a spread. The difference between the true covariance structure and the regulator’s imposed covariance leads to a regulatory wedge that creates cross-sectional distortions. This wedge can be decomposed into a calibration component the regulator could remove and a coverage component that no choice of weights can remove, an irreducibility the paper states as a formal bound (Proposition 1).

The third set of results concerns quantities. The two sectors’ portfolios deviate from their frictionless benchmarks by opposite positions in a single tilt portfolio, whose regulatory part involves combinations of regulated-factor exposure. Two identically risky assets that differ only in the extent to which their risks are captured in the regulatory framework trade at different prices, a coverage basis for which the paper derives an explicit formula. This implies that unmeasured risk gravitates to bank balance sheets, as charge-free funding and hedging at statutory weights and as inventory held outright at punitive ones. In the language of the standard, the low-charge set is concrete, with same-factor bases carrying a charge of zero and cross-curve and bond–CDS bases roughly one twenty-second of an outright position.

The fourth set of results turns the wedge decomposition into a quantified policy experiment. The paper estimates a counterfactual recalibration where the covered factors’ risk weights are set to be proportional to true factor volatilities, their prescribed correlations are set to the true correlations, and the overall scale is chosen to hold the charge on the frictionless bank book fixed. Measured tightness barely moves, but the composition distortions the statutory weights create are repaired. In this counterfactual, the capital charge on a beta-matched hedged-equity book falls from 2.3 times its true risk to approximate parity, and in the stressed equilibria the cross-sectional spread of regulatory premia narrows by roughly a fifth. What remains is the coverage component. The unchanged basis exposure, and the banking sector’s outsized position in it, are untouched by any reweighting, because the factor list, not the weights, produces them.

The fifth result is a welfare accounting. The deadweight loss is second order in the shadow cost

while the price effects are first order, and its burden falls on issuers through the cost of capital rather than on the banks the rule constrains.

Turning to magnitudes, a back-of-the-envelope calculation using the standard’s own delta risk weights puts the cross-sectional pricing hit at roughly 20 basis points per year on ten-year government duration, 30–50 on foreign exchange and equity-index exposure, and 90–165 on single-name and emerging-market equity at a central tightness calibration. The corresponding basis-trade twins earn less than one basis point: the same orders of magnitude as the covered interest parity deviations, negative swap spreads, and funding value adjustments documented since the post-crisis reforms. The general-equilibrium calibration disciplines these numbers downward in normal times, while impairment of the unconstrained sector’s capacity carries them to the top of the range, so the distortions are procyclical.

Three disclaimers are worth stating plainly. First, the model is static and delta-only with a single correlation scenario. The vega and curvature charges, the default risk charge, and the scenario-max layer, which converts the constraint into a maxmin objective, are left for other work. Second, nothing here speaks to the appropriate level of capital, whose benefits accrue against failure states the model does not address. Instead, the analysis concerns the relative price and allocation structure a given level induces, a margin improvable without lowering aggregate capital. Third, no misreporting, ratings management, or evasion is assumed anywhere. Banks report their sensitivities truthfully and optimize honestly, so the distortions are properties of the rule, not of misconduct.

Section 2 describes related results in the literature. Section 3 sets up the economy and the charge. Section 4 derives the equilibrium. Section 5 characterizes the tilt toward unregulated risk, maps the results into FRTB language, constructs the sensitivity gap, and derives the welfare accounting. Section 6 solves the model at FRTB parameters and runs the recalibration counterfactual. Sections 7 and 8 draw out the policy implications and conclude.

2 Related Literature

This paper connects four literatures: equilibrium asset pricing with portfolio constraints, models of risk-measure-based regulation, intermediary asset pricing and the empirics of post-crisis balance-sheet costs, and work on coarse regulatory risk measurement and the behavior it induces. The

discussion is organized by proximity to the present model and results.

Equilibrium with portfolio constraints. The formal machinery descends from [Black \(1972\)](#), whose restricted-borrowing CAPM delivers the positive zero-beta intercept that reappears here as the funding wedge, the uniform rent every asset pays for space on a constrained balance sheet, in [Proposition 3](#). [Cuoco \(1997\)](#) provides the general theory of equilibrium with portfolio constraints in this class of economies, and [Basak and Croitoru \(2000\)](#) is the direct antecedent of the twin-asset result, showing that portfolio constraints allow two *redundant* securities to trade at different prices in equilibrium. The same law-of-one-price break animates the limits-of-arbitrage tradition, in which financially constrained arbitrageurs cannot fully close price gaps between identical or redundant assets ([Shleifer and Vishny, 1997](#); [Gromb and Vayanos, 2002](#)). The coverage basis of [Corollary 3](#) is its regulatory-capital counterpart, with the SBM charge in place of a collateral or equity constraint as the friction that sustains the gap. Closest on the constraint’s functional form are the margin-CAPM of [Gârleanu and Pedersen \(2011\)](#) and the leverage-constrained investors of [Frazzini and Pedersen \(2014\)](#), in which each asset’s premium is proportional to its own margin or leverage cost. The SBM charge differs from a margin schedule in one essential respect. It is a joint quadratic form in the entire portfolio rather than a sum of per-asset terms, so the marginal capital consumed by any one asset depends on the whole book, hedges receive offsets, and the effective “margin” is endogenous to portfolio composition. This is what replaces the per-asset margin premium with a second *beta* ([Proposition 3](#)) and makes the constraint distort relative valuations, not just the level and slope of the security market line.

Risk-measure-based regulation in equilibrium. A second strand studies investors constrained by statistical risk measures: optimal policies and prices under VaR constraints in general equilibrium ([Basak and Shapiro, 2001](#); [Cuoco and Liu, 2006](#)), the endogenous risk, procyclical leverage, and margin-liquidity spirals such constraints induce ([Danielsson et al., 2004](#); [Adrian and Shin, 2014](#); [Brunnermeier and Pedersen, 2009](#)), and, in the central banking literature, their general-equilibrium procyclicality and interaction with the risk-insensitive leverage ratio ([Covas and Fujita, 2010](#); [Gambacorta and Karmakar, 2018](#)). In all of these, the regulatory metric is computed under the *true* return distribution. The departure here is that the metric is computed under the regulator’s

prescribed covariance matrix, which differs from the truth both in calibration and, because the regulated factors span only part of the true factor space, in rank. The results in this paper on the coverage and calibration wedge premia do not seem to have a direct counterpart in the VaR literature.

Intermediary asset pricing and balance-sheet costs in prices. That intermediary capital is a priced state variable is the message of [He and Krishnamurthy \(2013\)](#) theoretically and [He et al. \(2017\)](#) empirically. The shadow cost of bank capital, scaled by the banking sector’s share of aggregate risk-bearing capacity, is this paper’s static, regulation-specific counterpart, with Propositions 5 and 6 qualifying its use as a tightness diagnostic. A large post-crisis empirical literature documents funding and regulatory balance-sheet costs in asset prices: covered interest parity deviations, funding value adjustments, balance-sheet “rental” costs, and negative swap spreads ([Du et al., 2018](#); [Andersen et al., 2019](#); [Fleckenstein and Longstaff, 2020](#); [Jermann, 2020](#); [Klingler and Sundaresan, 2019](#); [Duffie, 2018](#); [Baker and Wurgler, 2015](#)). And [Kisin and Manela \(2016\)](#) estimate the shadow cost of capital requirements directly from banks’ use of a costly loophole. These papers typically model the regulatory cost as a per-dollar or per-notional balance-sheet rent, uniform within an asset class. The contribution here is the cross-sectional structure of the rent implied by the SBM’s functional form, together with its equilibrium incidence across regulated and unregulated risks.

Coarse risk measurement and regulatory arbitrage. Closest to the results here is the literature on regulation with coarse or incomplete risk metrics. [Becker and Ivashina \(2015\)](#) document that insurers subject to ratings-based capital requirements reach for yield *within* rating buckets, loading on the risk the metric does not distinguish. [Boermans and van der Kroft \(2024\)](#) find the same tilt in European institutions’ bond holdings, and [Iannotta et al. \(2019\)](#) develop the corresponding theory, and [Opp et al. \(2013\)](#) show how rating-contingent regulation distorts both the ratings themselves and the resulting allocation of capital. [Behn et al. \(2022\)](#) and [Plosser and Santos \(2018\)](#) document miscalibration and strategic slack in model-based capital regulation, [Dubecq et al. \(2015\)](#) show how imperfect observability of bank capital invites risk shifting, [Acharya et al. \(2013\)](#) show pre-crisis securitization moved exposure outside the measured perimeter while retaining the risk, [Ellul et al. \(2011\)](#) trace the price effects of regulatory pressure in the corporate bond mar-

ket, and [Kojen and Yogo \(2015, 2022\)](#) quantify how statutory risk-based capital shapes insurers’ pricing, the closest structural analog to the SBM on the buy side. On the quantity side, tighter requirements also relocate intermediation across the regulated perimeter itself: [Plantin \(2015\)](#) shows theoretically, and [Begenau and Landvoigt \(2022\)](#) in a quantitative general equilibrium, that activity migrates to less-constrained balance sheets when capital is dear. Two points distinguish the present paper. First, no misreporting, ratings management, or perimeter engineering is assumed. The tilt toward unmeasured risk (Proposition 7) is the honest optimum against a rank-deficient metric. Second, the general-equilibrium counterpart matters, since unmeasured risk is not merely relocated to constrained balance sheets but *repriced* for everyone (Corollary 3). Outside capital regulation, benchmarking and mandate frictions reprice covered relative to uncovered assets on the demand side: [Basak and Pavlova \(2013\)](#) show that institutions evaluated against an index bid up its constituents, and [Kashyap et al. \(2021\)](#) that a mandate referencing only a subset of assets subsidizes inclusion, the closest analog of the coverage basis priced here.

The FRTB. Academic work on the FRTB itself is largely computational, for example [Grajales and Medina Hurtado \(2023\)](#) on SBM and expected-shortfall capital for options books and [Avelone et al. \(2025\)](#) on portfolios that minimize the internal-models charge, while industry studies document the migration of large banks onto the standardized approach ([International Swaps and Derivatives Association and EY, 2024](#)), which is what makes the SBM the marginal capital metric in many markets. The rule’s parameters are live policy terrain rather than settled arithmetic. Go-live dates span half a decade across jurisdictions ([Basel Committee on Banking Supervision, 2026](#)), and the European Commission’s temporary relief measures of June 2026, adopted explicitly on competitive-parity grounds, amount to a recalibration debate conducted through a multiplier ([European Commission, 2026](#)). This paper takes the rule’s architecture as given and asks what it does to equilibrium prices. One feature of that architecture is deliberately stripped out: the MAR21.6 requirement that the charge be computed under three correlation scenarios with the maximum binding. A maximum over correlation models is a maxmin objective, under which investors facing correlation ambiguity exhibit portfolio and price inertia ([Illeditsch et al., 2021](#)). The present model, with its single prescribed correlation matrix, is the smooth benchmark on which that kinked, worst-case layer sits.

3 Model

The economy is static, with two dates $t \in \{0, 1\}$. Trade occurs at date 0, and payoffs are realized and consumed at date 1. There is a riskless asset and N risky assets whose payoffs are driven by L underlying risk factors, of which the regulator's capital framework recognizes only $K < L$. Two classes of investors trade: *traditional investors*, who solve a standard mean–variance problem, and *constrained investors* (“banks”), who face a stylized sensitivities-based capital requirement modeled on the delta component of the FRTB standardized approach.

3.1 Assets and risk factors

The riskless asset has an exogenous gross return $R_f \geq 1$ and is in perfectly elastic supply. The N risky assets have date-1 payoff vector $X \in \mathbb{R}^N$ generated by an L -factor structure,

$$X = \mu + Mf + \varepsilon, \quad (1)$$

where $\mu \in \mathbb{R}^N$ is the vector of expected payoffs, $f \in \mathbb{R}^L$ is a vector of risk factors with $f \sim \mathcal{N}(0, \Sigma_f)$, $\varepsilon \sim \mathcal{N}(0, \Sigma_\varepsilon)$ is idiosyncratic noise independent of f , and $M \in \mathbb{R}^{N \times L}$ is the matrix of factor exposures. Row i of M collects the *sensitivities* (deltas) of asset i 's payoff to the L risk factors: $M_{i\ell} = \partial X_i / \partial f_\ell$. The payoff covariance matrix is

$$\Sigma := \text{Var}(X) = M\Sigma_f M' + \Sigma_\varepsilon, \quad (2)$$

and the following assumption is maintained throughout:

Assumption 1. $\Sigma_f \succ 0$, $\Sigma_\varepsilon \succeq 0$, and $\Sigma \succ 0$.

The regulatory framework recognizes only $K < L$ of these factors. Order the factors so that the regulated ones come first, and partition

$$f = \begin{pmatrix} f_r \\ f_u \end{pmatrix}, \quad M = [M_r \ M_u], \quad \Sigma_f = \begin{pmatrix} \Sigma_{rr} & \Sigma_{ru} \\ \Sigma_{ur} & \Sigma_{uu} \end{pmatrix}, \quad (3)$$

where $f_r \in \mathbb{R}^K$ collects the *regulated* factors, those for which the framework prescribes risk weights and correlations, and $f_u \in \mathbb{R}^{L-K}$ the *unregulated* factors: risks that are real, systematic, and priced

by investors, but invisible to the capital calculation. $M_r \in \mathbb{R}^{N \times K}$ and $M_u \in \mathbb{R}^{N \times (L-K)}$ are the corresponding blocks of exposures.

Risky-asset prices $p \in \mathbb{R}^N$ are determined in equilibrium. A portfolio is a vector $\theta \in \mathbb{R}^N$ of share holdings together with a dollar position $b \in \mathbb{R}$ in the riskless asset. Short sales of risky assets and borrowing are permitted, subject only to the capital requirement introduced below. The *net regulated sensitivity* of portfolio θ is the vector

$$s(\theta) := M_r' \theta \in \mathbb{R}^K, \quad (4)$$

the model counterpart of the net sensitivities a bank reports to its supervisor for the risk factors the framework prescribes.

3.2 The regulatory capital charge

The regulator prescribes a risk weight $w_k > 0$ for each regulated factor $k = 1, \dots, K$ and a correlation matrix $\Gamma = (\gamma_{k\ell}) \in \mathbb{R}^{K \times K}$, symmetric and positive semidefinite with unit diagonal. Given a portfolio's net sensitivities $s = s(\theta)$, the *weighted sensitivities* are $WS_k = w_k s_k$, and the capital charge is the single-layer, delta-only analog of the SBM aggregation formula:¹

$$\mathcal{K}(\theta) = \sqrt{\sum_{k=1}^K \sum_{\ell=1}^K \gamma_{k\ell} WS_k WS_\ell} = \sqrt{s(\theta)' \Omega s(\theta)}, \quad \Omega := D_w \Gamma D_w, \quad (5)$$

where $D_w = \text{diag}(w_1, \dots, w_K)$. Substituting $s(\theta) = M_r' \theta$ and defining the *regulatory covariance matrix*

$$\Sigma_{\mathcal{K}} := M_r \Omega M_r' \in \mathbb{R}^{N \times N}, \quad (6)$$

a positive semidefinite matrix of rank at most K , the charge takes the compact form

$$\mathcal{K}(\theta) = \sqrt{\theta' \Sigma_{\mathcal{K}} \theta}. \quad (7)$$

The function $\mathcal{K}$ is a seminorm on $\mathbb{R}^N$: it is nonnegative, positively homogeneous of degree one

¹Relative to the full standardized approach (Basel Framework MAR21; [Basel Committee on Banking Supervision, 2019](#)), the model makes three simplifications. First, it treats all K regulated factors as members of a single bucket, removing the two-layer structure in which within-bucket aggregation with correlations ρ is followed by cross-bucket aggregation with correlations γ . Second, it considers delta risk only, leaving out vega and curvature risk. Third, it works with a single prescribed Γ rather than the three correlation scenarios of MAR21.6. The scenario-max layer raises separate questions and is left for other work.

($\mathcal{K}(t\theta) = |t| \mathcal{K}(\theta)$), and convex. It is differentiable everywhere except on the cone $\{\theta : \Sigma_{\mathcal{K}}\theta = 0\}$ of portfolios that carry no regulatory risk, with gradient

$$\nabla \mathcal{K}(\theta) = \frac{\Sigma_{\mathcal{K}} \theta}{\mathcal{K}(\theta)} \quad \text{wherever } \mathcal{K}(\theta) > 0. \quad (8)$$

Two features of (7) will be important for what follows.

Remark 1 (The regulator's covariance matrix). Equation (7) says that the SBM charges capital as if portfolio risk were $\sqrt{\theta' \Sigma_{\mathcal{K}} \theta}$, that is, as if the covariance matrix of asset payoffs were $\Sigma_{\mathcal{K}}$ rather than the true Σ . Every parameter of $\Sigma_{\mathcal{K}}$ is set by the regulator. The bank supplies only its sensitivities. Furthermore, since the coverage gap caps the rank of $\Sigma_{\mathcal{K}}$ at $K < L$, whatever the risk weights, the regulator's matrix cannot reproduce true systematic risk unless the unregulated factors are irrelevant to payoffs. To keep track of both failures at once, fix a benchmark scale $\phi > 0$ and define the *regulatory wedge*

$$\Delta := \Sigma_{\mathcal{K}} - \phi^2 (\Sigma - \Sigma_{\varepsilon}). \quad (9)$$

This is the gap between the factor covariance structure the regulator imposes and the true factor covariance structure, scaled to a common benchmark. This wedge drives the pricing distortions in the model. There are two potential source of error. The first, which will be called the *calibration wedge*, is the difference between the regulator's prescribed weights and correlations and the true ones. The second, which will be called the *coverage wedge*, is the difference between the true factor covariance and its projection onto the span of the regulated factors. The next result makes this decomposition precise.

Using $\Sigma - \Sigma_{\varepsilon} = M \Sigma_f M'$ and the partition (3), the wedge decomposes as

$$\Delta = \underbrace{M_r (\Omega - \phi^2 \Sigma_{rr}) M_r'}_{\text{calibration wedge}} - \underbrace{\phi^2 (M_r \Sigma_{ru} M_u' + M_u \Sigma_{ur} M_r' + M_u \Sigma_{uu} M_u')}_{\text{coverage wedge}}. \quad (10)$$

The calibration wedge vanishes under a *neutral calibration* of the covered factors. This happens when risk weights are proportional to the true volatilities of the regulated factors, $w_k = \phi \sigma_{r,k}$, and prescribed correlations equal to the true ones, $\Gamma = \text{corr}(f_r)$, so that $\Omega = \phi^2 \Sigma_{rr}$. The coverage wedge cannot be calibrated away, because it is the risk the framework does not see, and it is irreducible given $K < L$ whenever $M_u \neq 0$. In FRTB terms, the calibration wedge happens when true risk

weights and correlations differ from statutory ones. The coverage wedge captures the risk impact of everything outside the prescribed risk-factor taxonomy: non-modelable risk factors, residual basis risks, and risks the list simply omits.

The calibration wedge is removable, since the neutral calibration above sets it to zero. But the coverage wedge is not, and this is a statement about the *rank* of the regulator's matrix that no choice of risk weights or correlations can escape. The next proposition shows the minimum value of this wedge is bounded below by the risk carried by the unregulated factors, and that bound is strictly positive whenever the unregulated factors are nonzero.

Proposition 1 (Coverage irreducibility). *Let $\Delta = \Sigma_K - \phi^2(\Sigma - \Sigma_\epsilon)$ be the wedge between the regulator's covariance matrix and the (benchmark-scaled) true systematic covariance of payoffs, as in Remark 1, and let $M_u^\perp$ denote the part of the unregulated factors' payoff loadings that lies outside the span of the regulated factors, the exposure the metric structurally cannot see. Then, however the regulator sets the risk weights and prescribed correlations,*

$$\|\Delta\| \geq \phi^2 \|M_u^\perp \Sigma_{uu} M_u^{\perp'}\| \geq 0, \quad (11)$$

where Σ_{uu} is the covariance of the unregulated factors and $\|\cdot\|$ is any measure of matrix size (formally, any unitarily invariant norm). The lower bound depends only on the risk carried by the uncovered exposure, not on the risk weights or correlations. Recalibration relocates and rescales the regulatory risk model but cannot push the wedge below it. The bound is strictly positive whenever the uncovered exposure is nonzero, so the regulatory and true risk models coincide (and the wedge can be driven to zero) only in the special case in which the regulated factors already span every priced source of risk.

Proposition 1 (proved in the supplementary appendix) is the formal counterpart of the paper's policy claim that the factor *list*, not the weights, is the binding design choice. Reweighting can drive the calibration wedge to zero, but the coverage wedge is bounded away from zero by an object the regulator does not control.

Remark 2 (Unmeasured risk is free). Because the charge depends on θ only through the regulated sensitivities $M_r'\theta$, any portfolio in the null space of M_r' carries a capital charge of zero. That null space has dimension at least $N - K$, and the coverage gap makes it economically substantive,

containing not only pure idiosyncratic positions but every position whose risk runs through the unregulated factors f_u . Such a portfolio can carry arbitrarily large true, systematic risk while consuming no regulatory capital. The model thus makes plain the incentive of a sensitivity-based regime to push constrained institutions toward the risks the regime does not measure. Section 5.1 makes this precise in equilibrium.

3.3 Investors

There is a continuum of investors of total mass one, divided into two classes: *traditional investors* (superscript U , for unconstrained) of mass $m_U \in (0, 1)$, and *constrained investors* (superscript C) of mass $m_C = 1 - m_U$. Investors within a class are identical, and all are price takers.

An investor of class $j \in \{U, C\}$ enters date 0 with wealth W_0^j , chooses risky holdings $\theta \in \mathbb{R}^N$ and a riskless position $b \in \mathbb{R}$ subject to the budget constraint $b + p'\theta = W_0^j$, and consumes date-1 wealth

$$W_1 = R_f b + \theta' X = R_f W_0^j + \theta' (X - R_f p). \quad (12)$$

Preferences over terminal wealth are CARA with risk aversion $\eta_j > 0$. Given the Gaussian structure (1), expected utility maximization is equivalent to the mean–variance criterion

$$V^j(\theta) := \theta'(\mu - R_f p) - \frac{\eta_j}{2} \theta' \Sigma \theta. \quad (13)$$

Traditional investors. Traditional investors solve the unconstrained problem

$$\max_{\theta \in \mathbb{R}^N} V^U(\theta), \quad (14)$$

with the familiar solution

$$\theta^U = \frac{1}{\eta_U} \Sigma^{-1} (\mu - R_f p). \quad (15)$$

Constrained investors. Constrained investors face, in addition, a stylized capital requirement: *the dollar position in the riskless asset must be at least as large as the capital charge on the risky book,*

$$b \geq \mathcal{K}(\theta). \quad (16)$$

Using the budget constraint $b = W_0^C - p'\theta$, the constrained problem is

$$\max_{\theta \in \mathbb{R}^N} V^C(\theta) \quad \text{subject to} \quad W_0^C - p'\theta \geq \sqrt{\theta' \Sigma_{\mathcal{K}} \theta}. \quad (17)$$

Remark 3 (Interpreting the constraint). Requirement (16) is a deliberately stylized, static stand-in for the regulatory requirement that eligible capital be at least as large as the market risk charge. It says that the capital backing the requirement must be held in riskless, liquid form rather than reinvested in risky positions. Every dollar of risky exposure therefore both consumes funding ($p'\theta$) and, through $\mathcal{K}(\theta)$, ties down additional riskless holdings. Note also that since $\mathcal{K} \geq 0$, requirement (16) implies $b \geq 0$: a constrained investor can never finance risky positions by borrowing at the riskless rate, so the constraint embeds a no-borrowing restriction on top of its risk-sensitive component. The constraint set is closed and convex and satisfies Slater's condition whenever $W_0^C > 0$, so (17) has a unique solution characterized by the Karush–Kuhn–Tucker conditions. The particular form of (16) is not essential for the results. Alternative formulations (a charge levied against equity, a leverage bound on borrowing, a requirement against risk-weighted assets (Gambacorta and Karmakar, 2018)) change only the funding component. In any formulation in which the shadow cost multiplies the marginal charge $\partial \mathcal{K} / \partial \theta_i$, the two-beta structure, the wedge premium, and the tilt toward unmeasured risk survive, because those results flow from the shape of $\mathcal{K}$ rather than from the financing assumption. The form used here makes the funding and regulatory channels separately visible with a single multiplier.

3.4 First-order conditions and the distorted demand of banks

Let $\lambda \geq 0$ be the multiplier on the capital requirement in (17). At any solution with $\mathcal{K}(\theta) > 0$, the first-order condition is

$$\mu - R_f p = \eta_C \Sigma \theta^C + \lambda p + \lambda \frac{\Sigma_{\mathcal{K}} \theta^C}{\mathcal{K}(\theta^C)}, \quad \lambda [W_0^C - p'\theta^C - \mathcal{K}(\theta^C)] = 0. \quad (18)$$

Rearranging (18) isolates the two channels through which the requirement distorts demand:

$$\mu - (R_f + \lambda) p = \left[\eta_C \Sigma + \frac{\lambda}{\mathcal{K}(\theta^C)} \Sigma_{\mathcal{K}} \right] \theta^C. \quad (19)$$

A binding capital requirement makes the constrained investor behave as if (i) her funding rate were $R_f + \lambda$ rather than R_f , a uniform funding spread levied on every dollar of risky investment, and (ii) her risk aversion were augmented by an additional term $\lambda/\mathcal{K}(\theta^C)$ applied to the *regulator's* covariance matrix $\Sigma_{\mathcal{K}}$ rather than the true covariance Σ . Channel (i) is a level effect familiar from margin- and leverage-constrained models. Channel (ii) is the distinctive feature of a sensitivities-based regime, re-pricing risk *asset by asset* according to the regulator's risk weights and correlations, and it is the source of every cross-sectional (composition) result in the paper. Under the neutral calibration of Remark 1, $\Sigma_{\mathcal{K}} \propto \Sigma - \Sigma_{\varepsilon}$ and channel (ii) collapses to a nearly uniform increase in effective risk aversion. Away from it, assets whose regulatory risk exceeds their true risk are penalized in the constrained investor's portfolio and must offer higher equilibrium returns to be held.

Lemma 1 in the next section states the resulting demand in closed form, and Sections 4.5–4.6 establish existence of equilibrium and the comparative statics of λ .

3.5 Equilibrium

The risky assets are in fixed per-capita supply $\bar{\theta} \in \mathbb{R}_+^N$. The riskless asset is in perfectly elastic supply at the exogenous rate R_f .

Definition 1 (Equilibrium). An *equilibrium* is a price vector $p \in \mathbb{R}^N$ and portfolios (θ^U, θ^C) such that (i) θ^U solves (14) and θ^C solves (17) given p , and (ii) the market for every risky asset clears:

$$m_U \theta^U + m_C \theta^C = \bar{\theta}. \quad (20)$$

When the capital requirement is slack ($\lambda = 0$), (15) and (18) coincide up to risk aversion and the model reduces to the standard CAPM with harmonic-mean risk aversion $\bar{\eta} := (m_U/\eta_U + m_C/\eta_C)^{-1}$ and pricing equation $\mu - R_f p = \bar{\eta} \Sigma \bar{\theta}$. The analysis in Section 4 characterizes how prices depart from this benchmark when the requirement binds.

4 Equilibrium Characterization

This section derives the equilibrium in closed form. Proofs of all results are collected in the supplementary appendix. Throughout, $\pi := \mu - R_f p \in \mathbb{R}^N$ denotes the vector of risk premia in payoff space (expected payoff in excess of the payoff of p_i dollars invested at the riskless rate). The focus is on equilibria in which the capital requirement binds with $\mathcal{K}(\theta^C) > 0$. Corollary 1 below gives the condition on bank capital W_0^C that separates the binding and slack regimes.²

4.1 Individual demands

Lemma 1 (Optimal portfolios). *Fix a price vector p and suppose $W_0^C > 0$.*

- (i) *The traditional investor's problem (14) has the unique solution $\theta^U = \frac{1}{\eta_U} \Sigma^{-1} \pi$.*
- (ii) *The constrained problem (17) has a unique solution θ^C , and there exists a multiplier $\lambda \geq 0$ satisfying $\lambda [W_0^C - p' \theta^C - \mathcal{K}(\theta^C)] = 0$ such that, whenever $\kappa := \mathcal{K}(\theta^C) > 0$,*

$$\theta^C = H(\lambda, \kappa)^{-1} [\mu - (R_f + \lambda) p], \quad H(\lambda, \kappa) := \eta_C \Sigma + \frac{\lambda}{\kappa} \Sigma_{\mathcal{K}} \succ 0. \quad (21)$$

Equation (21) restates the two channels of Section 3.4 as a demand function: a binding requirement makes the bank discount at the shadow-augmented rate $R_f + \lambda$ and price risk with the blended matrix $H(\lambda, \kappa)$, a convex-cone combination of the true covariance matrix and the regulator's.

4.2 Market clearing and equilibrium prices

Market clearing requires that the two sectors together absorb the supply of every risky asset,

$$m_U \theta^U + m_C \theta^C = \bar{\theta}, \quad (22)$$

while the riskless market clears automatically through its elastic supply at R_f . Substituting the demands of Lemma 1 into (22) and rearranging yields the paper's central pricing equation.

Proposition 2 (Equilibrium pricing). *In any equilibrium in which the capital requirement binds*

²Equilibria in which the requirement binds but $\mathcal{K}(\theta^C) = 0$ would require the constrained sector's entire book to lie in the zero-charge cone $\{\theta : \Sigma_{\mathcal{K}} \theta = 0\}$. This case is nongeneric and is set aside.

with shadow cost $\lambda > 0$ and equilibrium charge $\kappa = \mathcal{K}(\theta^C) > 0$, risk premia satisfy

$$\pi = \underbrace{\bar{\eta} \Sigma \bar{\theta}}_{\text{fundamental risk}} + \underbrace{\chi \lambda p}_{\text{funding wedge}} + \underbrace{\frac{\chi \lambda}{\kappa} \Sigma_{\mathcal{K}} \theta^C}_{\text{regulatory risk}}, \quad \chi := \frac{m_C \bar{\eta}}{\eta_C} = \frac{m_C / \eta_C}{m_U / \eta_U + m_C / \eta_C} \in (0, 1), \quad (23)$$

where $\bar{\eta} = (m_U / \eta_U + m_C / \eta_C)^{-1}$. Equivalently, equilibrium prices are

$$p = \frac{1}{R_f + \chi \lambda} \left[\mu - \bar{\eta} \Sigma \bar{\theta} - \frac{\chi \lambda}{\kappa} \Sigma_{\mathcal{K}} \theta^C \right]. \quad (24)$$

If instead the requirement is slack, then $\lambda = 0$ and (23)–(24) reduce to the standard CAPM: $\pi = \bar{\eta} \Sigma \bar{\theta}$ and $p = R_f^{-1}(\mu - \bar{\eta} \Sigma \bar{\theta})$.

Two features of (23) deserve comment.

First, the constraint enters prices only through the product $\chi \lambda$ and the gradient of the charge.

Using $\nabla \mathcal{K}(\theta^C) = \Sigma_{\mathcal{K}} \theta^C / \kappa$ from (8), the pricing equation can be written compactly as

$$\pi = \bar{\eta} \Sigma \bar{\theta} + \chi \lambda [p + \nabla \mathcal{K}(\theta^C)]. \quad (25)$$

The bracketed vector is asset i 's total claim on the bank's balance sheet: p_i dollars of funding plus $\partial \mathcal{K} / \partial \theta_i$ dollars of marginal regulatory capital. In equilibrium every asset bears an extra required return in proportion to that claim, at the uniform rate $\chi \lambda$: the shadow cost of the requirement, λ , scaled by the constrained sector's share χ of the economy's aggregate risk-bearing capacity. The scalar $\chi \lambda$ is thus the model's sufficient statistic for regulatory tightness, rising as bank capital falls, as risk weights rise, or as the constrained sector looms larger in aggregate risk tolerance.

Second, fundamental risk and regulatory risk are priced side by side, by different matrices, while the funding wedge prices scale. The first term of (23) compensates covariance with the market portfolio under the true covariance matrix Σ , and the third compensates covariance with the bank sector's book under the regulator's matrix $\Sigma_{\mathcal{K}}$: assets whose regulatory covariance with the bank book is large relative to their true covariance with the market must pay investors for the capital they consume, not only for the risk they carry. On top of both, equation (24) discounts expected payoffs at $R_f + \chi \lambda > R_f$, so a binding capital requirement anywhere in the economy raises the effective discount rate on *every* risky asset, including assets that carry no regulatory charge at all.

4.3 The two-beta CAPM and the cross-section

To express Proposition 2 in the language of returns, suppose equilibrium prices are strictly positive and let $R_i := X_i/p_i$ denote gross returns, $R_M := \bar{\theta}'X/p_M$ the market return with $p_M := p'\bar{\theta}$, and $R_C := \theta^C X/p_C$ the return on the constrained sector's book with $p_C := p'\theta^C$, assumed positive. For any portfolio returns, write $\text{Cov}_\mathcal{K}(\cdot, \cdot)$ and $\text{Var}_\mathcal{K}(\cdot)$ for the covariance and variance computed under the regulator's matrix $\Sigma_\mathcal{K}$ in place of Σ (so $\text{Cov}_\mathcal{K}(R_i, R_C) = e_i' \Sigma_\mathcal{K} \theta^C / (p_i p_C)$, the covariance the regulator's risk model *assigns*, whether or not it obtains).

Proposition 3 (Two-beta CAPM). *Under the conditions of Proposition 2, expected returns satisfy*

$$\mathbb{E}[R_i] - R_f = \chi\lambda + \beta_i^M \pi_M + \beta_i^K \pi_K, \quad i = 1, \dots, N, \quad (26)$$

where

$$\beta_i^M := \frac{\text{Cov}(R_i, R_M)}{\text{Var}(R_M)}, \quad \beta_i^K := \frac{\text{Cov}_\mathcal{K}(R_i, R_C)}{\text{Var}_\mathcal{K}(R_C)}, \quad \pi_M := \bar{\eta} p_M \text{Var}(R_M), \quad \pi_K := \frac{\chi\lambda \kappa}{p_C}.$$

Expected excess returns are spanned by two betas and an intercept. The *regulatory beta* β_i^K measures comovement with the bank sector's book as the regulator's risk model sees it, and its price π_K is the tightness statistic $\chi\lambda$ scaled by the regulatory riskiness of the bank book per dollar of value. The intercept $\chi\lambda$ is a zero-beta spread in the sense of Black (1972). Even an asset with both betas zero earns $R_f + \chi\lambda$, because holding it consumes constrained balance sheet. All three departures from the frictionless CAPM are proportional to $\chi\lambda$ and vanish together as the constraint slackens.

Corollary 1 (Binding threshold). *Let $p^0 := R_f^{-1}(\mu - \bar{\eta} \Sigma \bar{\theta})$ denote frictionless CAPM prices. An equilibrium with slack requirement exists if and only if*

$$W_0^C \geq \frac{\bar{\eta}}{\eta_C} [p^{0'} \bar{\theta} + \mathcal{K}(\bar{\theta})], \quad (27)$$

in which case it is the CAPM equilibrium of Proposition 2 with $\lambda = 0$ and $\theta^C = (\bar{\eta}/\eta_C) \bar{\theta}$. If (27) fails, the requirement binds ($\lambda > 0$) in every equilibrium.

The threshold has a transparent reading. In the frictionless CAPM the bank sector holds the fraction $\bar{\eta}/\eta_C$ of the market portfolio, and (27) asks whether bank equity covers the funding cost

plus the capital charge of that book.

The regulatory term in (23) mixes compensation for systematic risk that the regulator measures with compensation for the wedge between regulatory and true risk. Substituting $\Sigma_K = \phi^2(\Sigma - \Sigma_\varepsilon) + \Delta$, with Δ the wedge of (9), into (26) splits the regulatory premium of asset i into

$$\beta_i^K \pi_K = \underbrace{\frac{\chi\lambda\phi^2}{\kappa} \frac{[(\Sigma - \Sigma_\varepsilon)\theta^C]_i}{p_i}}_{\text{systematic repricing}} + \underbrace{\frac{\chi\lambda}{\kappa} \frac{(\Delta\theta^C)_i}{p_i}}_{\text{wedge premium}}. \quad (28)$$

The first term repackages true systematic risk at a higher price and changes only the level of risk compensation. The second is the cross-sectional signature of the wedge. An asset whose risk the covered-factor calibration overstates, and which banks hold long, carries a positive wedge premium, and its price is depressed relative to fundamentals. The coverage wedge enters Δ with a negative sign, so exposures that run through the unregulated factors command less expected return than their true risk warrants.

To understand the pricing implications of the policy choice, consider two extremes. At one extreme, the regulator's matrix is right. If the regulated factors span all payoff risk and are neutrally calibrated, so that $\Sigma_K = \phi^2 \Sigma$ (a counterfactual benchmark requiring $M_u = 0$, $\Sigma_\varepsilon = 0$, and hence $K \geq N$), then expected returns obey an exact one-factor CAPM in true market beta with a strictly positive zero-beta intercept, $\mathbb{E}[R_i] - R_f = b + a p_M \text{Cov}(R_i, R_M)$ with $a > \bar{\eta}$ and $b > 0$ given in closed form in the supplementary appendix. A fully covering, perfectly calibrated sensitivities-based requirement raises required returns in proportion to the size of the position, as in Black (1972) and Frazzini and Pedersen (2014), shifting and tilting the security market line but leaving relative valuations undistorted. The opposite extreme is exposure the model misses entirely.

Corollary 2 (Zero-charge portfolios). *Let $z \in \mathbb{R}^N$ satisfy $M_1' z = 0$ and $p' z \neq 0$. Then the portfolio z has $\beta_z^K = 0$ and*

$$\mathbb{E}[R_z] - R_f = \chi\lambda + \beta_z^M \pi_M.$$

Zero-charge portfolios earn no regulatory premium, priced by fundamental beta plus the funding wedge alone, regardless of how much true risk they carry, idiosyncratic or systematic through the unregulated factors.

Such positions are the cheapest risk on a constrained balance sheet, priced as if the requirement

did not exist except for the funding wedge. Everything in between the two extremes is governed by (28), and Section 5.1 develops the portfolio-composition counterpart.

4.4 Equilibrium computation

Proposition 2 characterizes prices given the equilibrium values of $(\lambda, \kappa, \theta^C)$, which are themselves endogenous. The fixed point has a convenient structure. For *given* (λ, κ) the model is linear, so the entire equilibrium reduces to two scalar consistency conditions.

Proposition 4 (Scalar reduction). *For each $(\lambda, \kappa) \in [0, \infty) \times (0, \infty)$, the linear system consisting of the demands of Lemma 1 and market clearing (22) has a unique solution $(p(\lambda, \kappa), \theta^U(\lambda, \kappa), \theta^C(\lambda, \kappa))$, given in closed form in the supplementary appendix. A price system with binding requirement is an equilibrium if and only if (λ, κ) solves the two scalar equations*

$$\mathcal{K}(\theta^C(\lambda, \kappa)) = \kappa, \quad p(\lambda, \kappa)' \theta^C(\lambda, \kappa) + \kappa = W_0^C. \quad (29)$$

The first equation makes the conjectured charge consistent with the implied book. The second is the binding capital requirement itself. Numerically, computing the equilibrium is a two-dimensional root-finding problem regardless of N and K . This is the basis of the calibrated illustration in Section 6.

4.5 Existence

Theorem 1 (Existence). *Under Assumption 1, for any $W_0^C > 0$, $m_U \in (0, 1)$, and $\bar{\theta} \in \mathbb{R}^N$, an equilibrium exists. If (27) holds, it is the CAPM allocation with slack requirement. If (27) fails, an equilibrium exists and in every equilibrium the requirement binds, $\lambda > 0$.*

The proof, in the supplementary appendix, exploits a feature of the CARA-normal structure that does the heavy lifting. Since traditional investors are unconstrained, their first-order condition holds in any equilibrium, and market clearing then expresses risk premia as an affine function of the bank sector's book alone:

$$\pi(\theta) = \frac{\eta_U}{m_U} \Sigma (\bar{\theta} - m_C \theta), \quad (30)$$

so an equilibrium is a fixed point of the bank's best response to the prices implied by its own book.

Because $\theta = 0$ is always feasible ($W_0^C > 0$) and delivers zero excess utility, any fixed point generates nonnegative surplus at the prices (30) it induces, yielding the a priori bound

$$\|\theta^C\|_\Sigma \leq \rho^* := \frac{(\eta_U/m_U) \|\bar{\theta}\|_\Sigma}{\eta_C/2 + \eta_U m_C/m_U}, \quad \|x\|_\Sigma := \sqrt{x' \Sigma x}, \quad (31)$$

on *every* candidate equilibrium book, for every configuration of risk weights. Existence then follows from Brouwer's theorem applied to the best response truncated to this ball. The bound says that the bank's book is limited by fundamentals and preferences alone, however the regulator sets (w, Γ) . Uniqueness is not claimed. In the single-factor economy of Proposition 6 the equilibrium is unique, and a numerical probe of the calibrated economy in Section 6 finds no equilibrium other than the reported one.

4.6 Comparative statics

Two results describe how the equilibrium moves with the model's policy-relevant parameters: a partial-equilibrium result on the shadow cost of capital that holds in full generality, and a complete general-equilibrium characterization in the single-factor economy, in which the coverage gap is shut down by construction and the statics isolate the funding and scale channels.

Proposition 5 (The shadow value of bank capital). *Fix prices p and let $v(W)$ denote the value of the constrained problem (17) with equity $W > 0$. Then v is nondecreasing and concave, and any KKT multiplier $\lambda(W)$ is a supergradient of v at W :*

$$v(W') \leq v(W) + \lambda(W) (W' - W) \quad \text{for all } W' > 0.$$

Consequently multipliers are nonincreasing in equity, meaning $W < W'$ implies $\lambda(W) \geq \lambda(W')$, and $v'(W) = \lambda(W)$ wherever v is differentiable.

Holding prices fixed, the shadow cost of the requirement falls as bank capital rises, because λ is the marginal value of equity and the constraint is convex. The single-factor economy adds the price channels and delivers global, strict comparative statics.

Proposition 6 (The single-factor economy). *Let $N = 1$ with a single, fully covered risk factor ($K = L = 1$, $M = 1$), $\text{Var}(X) = \sigma^2 > 0$, risk weight $w > 0$, supply $\bar{\theta} > 0$, and bank equity*

$W := W_0^C > 0$, so that $\mathcal{K}(\theta) = w|\theta|$ and the threshold (27) reads $W \geq \bar{W} := (\bar{\eta}/\eta_C)(p^0 + w)\bar{\theta}$ with $p^0 = R_f^{-1}(\mu - \bar{\eta}\sigma^2\bar{\theta})$. Suppose $W < \bar{W}$. Then:

(i) (Uniqueness.) *There is a unique equilibrium. The price p^* is the unique root in $(-w, \infty)$ of the market-clearing equation*

$$\frac{m_U}{\eta_U \sigma^2} (\mu - R_f p) + \frac{m_C W}{p + w} = \bar{\theta}, \quad (32)$$

with $\theta^C = W/(p^* + w) > 0$, $\theta^U = (\mu - R_f p^*)/(\eta_U \sigma^2) > 0$, and

$$\lambda = \frac{\sigma^2}{p^* + w} \left[\frac{\eta_U}{m_U} \bar{\theta} - \left(\frac{\eta_U m_C}{m_U} + \eta_C \right) \theta^C \right] > 0, \quad (33)$$

which is positive because $\theta^C < (\bar{\eta}/\eta_C)\bar{\theta}$. The bank holds less than its frictionless CAPM book.

(ii) (Bank equity.) p^* and θ^C are strictly increasing in W . The risk premium $\mu - R_f p^*$ and the shadow cost λ are strictly decreasing in W .

(iii) (Risk weights.) The risk premium $\mu - R_f p^*$ is strictly increasing in w , with

$$\lim_{w \rightarrow \infty} (\mu - R_f p^*) = \frac{\eta_U \sigma^2 \bar{\theta}}{m_U},$$

the premium of an economy in which the constrained sector holds no risky assets. The shadow cost is not monotone: $\lambda \rightarrow 0$ as $w \rightarrow \infty$, and if $\bar{w} := \eta_C W / (\bar{\eta} \bar{\theta}) - p^0 > 0$ then also $\lambda \rightarrow 0$ as $w \downarrow \bar{w}$, so λ is hump-shaped on the binding region $(\bar{w}, \infty)$ while the premium distortion increases monotonically throughout.

Remark 4 (Tightness is not the shadow cost). Part (iii) separates two objects that are easily conflated. As risk weights rise, the *price* distortion, the gap between the equilibrium premium and its frictionless level $\bar{\eta}\sigma^2\bar{\theta}$, rises monotonically, converging to full disintermediation of the banking sector. The *shadow cost* λ , by contrast, first rises and then falls back to zero, because with punitive risk weights the bank holds almost nothing, so a marginal dollar of equity is worth little, even though prices are then maximally distorted. The right diagnostic in the model is quantities, not shadow prices: the gap between the bank sector's book and its frictionless benchmark $(\bar{\eta}/\eta_C)\bar{\theta}$.

Section 6 explores the general case numerically. Under both a proportional change in risk weights and a change in bank capital, the single-factor comparative statics carry over to the calibrated

multi-asset economy.

5 Asset Pricing Implications

5.1 Who holds what: the tilt toward unregulated risk

Proposition 2 characterizes prices. This section characterizes the quantities held by banks and how the capital requirement reallocates risk between the two sectors. It is shown that the entire reallocation can be characterized by a position in a single portfolio.

Proposition 7 (The tilt portfolio). *In any equilibrium in which the requirement binds ($\lambda > 0$, $\kappa = \mathcal{K}(\theta^C) > 0$), the two sectors' books deviate from their frictionless CAPM benchmarks by opposite positions in a single tilt portfolio τ :*

$$\theta^C = \frac{\bar{\eta}}{\eta_C} \bar{\theta} - \frac{(1-\chi)\lambda}{\eta_C} \tau, \quad \theta^U = \frac{\bar{\eta}}{\eta_U} \bar{\theta} + \frac{\chi\lambda}{\eta_U} \tau, \quad \tau := \Sigma^{-1} [p + \nabla \mathcal{K}(\theta^C)], \quad (34)$$

and the two deviations aggregate to zero. What the banks shed, the traditional investors absorb. The tilt decomposes as $\tau = \tau^{\text{fund}} + \tau^{\text{reg}}$ with

$$\tau^{\text{fund}} := \Sigma^{-1} p, \quad \tau^{\text{reg}} := \frac{1}{\kappa} \Sigma^{-1} M_r \Omega s(\theta^C) \in \Sigma^{-1} \text{col}(M_r),$$

so the regulatory component lies in the K -dimensional subspace $\Sigma^{-1} \text{col}(M_r)$: the constraint induces banks to shed combinations of regulated-factor exposure only. It never directly sheds exposure to the unregulated factors f_u .

The tilt is a mean-variance portfolio in balance-sheet cost. It helps to read the definition of τ in (34) against the frictionless holding it echoes. An unconstrained mean-variance investor holds $\eta^{-1} \Sigma^{-1} \pi$, the inverse covariance matrix applied to the vector of *risk premia*, which is the per-asset reward for bearing risk. The tilt portfolio has the identical form, $\tau = \Sigma^{-1} [p + \nabla \mathcal{K}(\theta^C)]$, with a single substitution. The reward vector π is replaced by $p + \nabla \mathcal{K}(\theta^C)$, the per-asset *claim on the balance sheet* that appears in the pricing equation (25): p_i dollars of funding plus $\partial \mathcal{K} / \partial \theta_i$ dollars of marginal capital. Where an ordinary investor tilts toward high expected return per unit of risk, the bank tilts *away* from high balance-sheet cost per unit of risk, shedding the portfolio τ that buys the most constraint relief for the least true variance. The pre-multiplication by Σ^{-1} carries the

same meaning in both problems. It nets correlated positions and rewards hedges, so the bank does not shed each asset in proportion to its own charge but sheds the variance-efficient combination, reducing the balance-sheet claim wherever doing so disturbs its risk position least. The two pieces of the decomposition are the two halves of that claim priced on their own: $\tau^{\text{fund}} = \Sigma^{-1}p$ is the efficient portfolio in the funding claim, and $\tau^{\text{reg}} = \Sigma^{-1}\nabla\mathcal{K}(\theta^C)$ the efficient portfolio in the marginal-capital claim.

Two remarks connect (34) to what came before. First, substituting the equilibrium θ^U into the traditional investors' first-order condition recovers the pricing equation (25), in which the wedge and regulatory premia are the compensation traditional investors demand for holding the tilt. Second, because τ^{reg} is confined to a K -dimensional subspace built from the regulator's factor list, unregulated risk becomes the cheap input on a constrained balance sheet.

Equal risk aversion. It is easier to interpret the tilt portfolio when the two sectors have the same risk aversion. Set $\eta_U = \eta_C = \eta$. Because the masses sum to one, $\bar{\eta} = \eta$ and $\chi = m_C$, and (34) collapses to

$$\theta^C = \bar{\theta} - \frac{m_U \lambda}{\eta} \tau, \quad \theta^U = \bar{\theta} + \frac{m_C \lambda}{\eta} \tau. \quad (35)$$

The benchmark is now the market portfolio itself. Without the requirement, both sectors would hold $\bar{\theta}$ per capita, so every difference between the two books is the requirement's doing, and that difference is a position in the single portfolio τ . The uniform rate in the pricing equation simplifies alongside, to $\chi\lambda = m_C\lambda$: the shadow cost of bank capital, diluted by the share of the economy that is not subject to the requirement.

Three features of (35) carry over to the general case. First, each sector's deviation is scaled by the *other* sector's mass, which is what makes the two deviations aggregate to zero. Banks shed $m_U\lambda/\eta$ units of the tilt and traditional investors absorb $m_C\lambda/\eta$ units, and market clearing holds term by term. Second, the smaller the unregulated sector, the larger the per-capita position it must hold. As $m_C \rightarrow 1$ the traditional investors' book diverges from the market portfolio without bound in the direction of τ , which is the quantity counterpart of the rising premia they charge for the service. Third, the distance between the two books, $\theta^U - \theta^C = (\lambda/\eta) \tau$, carries no masses at all. The composition of the economy governs how the tilt is divided between the sectors but not how

far apart their books end up. That separation is the shadow cost of capital measured in units of risk tolerance, and it is the object the sensitivity gap of Section 5.3 recovers from reported data.

The sharpest expression of the tilt holds fundamentals fixed and varies only the regulatory label, comparing two assets that are statistically identical but for the factor list.

Corollary 3 (Twin assets: on-list versus off-list risk). *Let assets a and b have identical fundamentals (expected payoff μ_0 , payoff variance σ^2 , per-capita supply $\bar{\theta}_0 > 0$, mutually uncorrelated and uncorrelated with all other assets) and differ only in regulatory treatment: a loads on a regulated factor with risk weight $w > 0$, uncorrelated under the regulator's correlation matrix Γ with every other regulated exposure in the economy, while b loads on an unregulated factor. Suppose the bank's equilibrium position in b is strictly positive. Then, in any binding equilibrium,*

$$\pi_a - \pi_b = \frac{m_C \sigma^2 \theta_b^C \frac{\zeta}{1+\zeta}}{\frac{m_U}{\eta_U} + \frac{m_C}{\eta_C} \frac{1+\lambda/R_f}{1+\zeta}} > 0, \quad \zeta := \frac{\lambda w^2}{\kappa \eta_C \sigma^2}, \quad (36)$$

equivalently $p_a = p_b - (\pi_a - \pi_b)/R_f < p_b$. The bank holds strictly less of the on-list twin ($\theta_a^C < \theta_b^C$) and traditional investors strictly more ($\theta_a^U > \theta_b^U$). Holding the economy-wide statistics $(\lambda, \kappa, \theta_b^C)$ fixed, which is exact when the pair is a vanishing share of the aggregate book, the gap (36) is strictly increasing in the risk weight w and converges, as $w \rightarrow \infty$, to $\eta_U m_C \sigma^2 \theta_b^C / m_U$: the premium impact of full bank exit from asset a .

The twins show the coverage distortion in its simplest form, two assets with *identical* payoff distributions trading at strictly different prices because only one of them appears on the regulator's factor list. The gap (36) is a *coverage basis* that vanishes with λ , w , and m_C , widens with regulatory tightness, and predicts a clientele pattern with a testable signature. Off-list risk concentrates on bank balance sheets and earns less than its fundamental risk warrants, the on-list twin migrates to unconstrained investors and earns more, and the spread between them comoves with dealer capital tightness.

5.2 The model in FRTB language

The model's regulated factors f_r map to the delta risk factors that MAR21 prescribes for its seven risk classes (MAR21.8–21.14). For general interest rate risk (GIRR) the factors are the risk-free

yield curve of each currency sampled at ten tenor vertices (0.25 to 30 years), together with a flat inflation rate and cross-currency basis factors (MAR21.8). For the credit spread risk (CSR) classes they are issuer credit spread curves. For equity risk they are spot prices and repo rates (MAR21.12), for commodity risk spot prices by grade and delivery location, and for foreign exchange risk each exchange rate against the reporting currency (MAR21.14) ([Basel Committee on Banking Supervision, 2019](#)). The unregulated factors f_u are everything that moves trading-book values and is not on that list, or that is on it only after the netting embedded in the sensitivity calculation has removed it. The framework itself acknowledges the residue. Jump-to-default is charged separately (the default risk capital requirement, MAR22), and a *residual risk add-on* (RRAO) levies flat charges on gross notional (MAR23.8): 1.0% for instruments on exotic underlyings such as longevity, weather, and natural disasters (MAR23.3), and 0.1% for instruments bearing other residual risks such as gap, correlation, and behavioral risk (MAR23.5). The RRAO is the regulation’s own patch for the coverage wedge. Crucially, it is notional-based rather than risk-based, so *within* the RRAO-covered set the marginal charge is again invariant to true risk.³

Corollary 3 then predicts which exposures gravitate to standardized-approach banks: those whose *marginal* charge $\partial\mathcal{K}/\partial\theta_i$ is small relative to their true risk. Reading the SBM’s netting and correlation structure through that lens points to several places where real risk goes largely uncharged.

1. *Same-factor bases carry no charge.* Two instruments mapping to the same risk factor net perfectly: the on-the-run/off-the-run Treasury basis, the cash–futures basis at a common vertex, and any specialness, liquidity, or deliverability premium between instruments the taxonomy treats as identical. These are the model’s twins in pure form. Such bases are likely to concentrate on standardized-approach bank balance sheets, trade tighter than their true risk warrants, and widen when dealer capital tightens.
2. *Cross-curve bases within a currency measure almost zero.* Sensitivities to different curves in the same currency (e.g., OIS versus term-rate curves) and to different instruments of the same issuer (the bond–CDS basis) are aggregated at a prescribed correlation of 99.90% (MAR21.47

³The delta-only model abstracts from the vega and curvature charges and from the default risk capital requirement. The predictions below are therefore for delta-dominated, default-remote books (rates, FX, and linear equity exposure) and for relative comparisons within them.

for different curves in the same currency, MAR21.54 for different curves of the same issuer). A \$1-for-\$1 ten-year cross-curve basis position carries $\sqrt{2(1 - 0.999)} \approx 1/22$ of the capital of an outright ten-year position: the same true basis risk, at four percent of the measured risk. Prediction: dealer concentration in cross-curve and bond–CDS basis trades, with compressed spreads in calm times. This is the model’s rationalization of the balance-sheet-intensive basis trades documented in the intermediary literature of Section 2.

3. *High-risk-weight directional exposure migrates out.* At the other end of the taxonomy, single-name (risk weights 15–70% by bucket), emerging-market, and small-cap equity and commodities carry the largest marginal charges. The model predicts they migrate to unconstrained investors and earn the funding-plus-regulatory premium of Proposition 3. The clientele prediction is testable in holdings data. Bank-to-nonbank holding ratios should decline in marginal capital intensity, conditional on fundamentals.

5.3 A supervisory diagnostic: the sensitivity gap

The tilt portfolio of Proposition 7 is the model’s sufficient statistic for the requirement’s effect on prices, but it is not directly observable, since it is built from equilibrium prices and the true covariance matrix. Its image in the space of *reported sensitivities*, however, is observable, and this subsection defines that diagnostic and shows how a supervisor would build it. Recall that $\chi = m_C \bar{\eta}/\eta_C$ is the constrained sector’s share of aggregate risk-bearing capacity, so that in the frictionless benchmark the banking sector in aggregate holds $\chi \bar{\theta}$, the fraction χ of the market portfolio.

Definition 2 (Sensitivity gap). Let $s^{\text{bank}} := M'_r(m_C \theta^C)$ denote the aggregate net regulated sensitivities of the banking sector and $s^{\text{mkt}} := M'_r \bar{\theta}$ those of the market portfolio. The *sensitivity gap* is the vector of weighted-sensitivity shortfalls

$$g := D_w \left(\chi s^{\text{mkt}} - s^{\text{bank}} \right), \quad G := \|g\|, \quad (37)$$

the factor-by-factor difference, in the standard’s own capital units, between the regulated sensitivities the banking sector would report in the frictionless benchmark and those it actually reports.

Proposition 8 (The gap measures the distortion). *In any binding equilibrium,*

$$g = \chi(1 - \chi) \frac{\lambda}{\bar{\eta}} D_w M'_r \tau, \quad \text{while} \quad \pi - \bar{\eta} \Sigma \bar{\theta} = \chi \lambda \Sigma \tau, \quad (38)$$

where τ is the tilt portfolio of Proposition 7. The gap and the vector of pricing distortions are linear images of the same equilibrium object $\lambda\tau$: the gap vanishes when the requirement is slack and is generically nonzero whenever it binds.

The proof is immediate from Proposition 7, which gives $m_C \theta^C - \chi \bar{\theta} = -\chi(1 - \chi)(\lambda/\bar{\eta}) \tau$, and from the pricing equation (25). The content is in what the two expressions in (38) do *not* share. The pricing distortion is filtered through the unobservable Σ , while the gap is filtered through the regulator's own $D_w M'_r$, which is the map banks already apply to their books when they compute the charge. Unlike the shadow cost, the gap cannot fall while the distortion rises, because both are increasing images of the same tilt. When banks respond to punitive weights by shedding the charged book, λ collapses but the shed quantity, which is what g records, keeps growing. Section 6.5 verifies this in the calibrated economy, where under a proportional increase in risk weights G rises monotonically with the largest pricing hit through the entire region in which λ peaks and falls.

Both ingredients of (37) are computable from data the regime already generates, so the gap is built to be computed rather than estimated. Factor by factor,

$$g_k = w_k (\chi s_k^{\text{mkt}} - s_k^{\text{bank}}), \quad (39)$$

the difference, at regulated factor k , between the weighted sensitivity the banking sector would report in the frictionless benchmark and the weighted sensitivity it actually reports.

The bank term $w_k s_k^{\text{bank}} = \sum_b \text{WS}_k^b$ is the sum, across standardized-approach banks b , of the weighted sensitivities $\text{WS}_k = w_k s_k$ that are line items of the SBM aggregation itself (MAR21.4). Every reporting bank already computes them to produce its own charge, so a supervisor aggregating them by factor across banks is feasible.

The benchmark term $\chi w_k s_k^{\text{mkt}}$ needs the regulated-factor weighted sensitivities of the outstanding market portfolio, recoverable from issuance and holdings statistics through the same factor definitions, together with an estimate of the intermediary share χ . Neither needs to be exact, since both move slowly, so the informative variation in g at supervisory frequencies comes from the

Table 1: Illustrative sensitivity-gap dashboard, reconstructed from the calibrated economy of Section 6. The benchmark column is χ times each factor’s market weighted sensitivity (capital intensity $\times$ supply, Table S.2). The bank column applies the equilibrium bank shares of Table 2. The gap is their difference (Definition 2). Entries are in the standard’s weighted-sensitivity (capital) units and are rounded independently, so the columns need not combine exactly. The norm of the gap column reproduces the baseline $G = 0.53$ of Figure 2. Populated with reported supervisory weighted sensitivities, the same three columns are the operational indicator.

Regulated factor	Bank $\sum_b WS_k^b$	Benchmark χWS_k^{mkt}	Gap g_k
GIRR 10y (duration)	0.80	1.32	0.53
Equity bucket 12 (index)	1.55	1.50	−0.05
Equity bucket 5 (large-cap)	0.82	0.80	−0.02
Norm $G = \ g\ $			0.53

reported bank sensitivities themselves.

The reporting frequency is not a constraint. The Basel III monitoring exercise already collects granular market-risk data on a semiannual cycle (Basel Committee on Banking Supervision, 2025), and the factor-level aggregation the gap requires is an incremental extension of it. At national supervisory-reporting frequency the bank term refreshes automatically.

Table 1 carries out the assembly at the calibrated economy of Section 6. Each row is a regulated factor, and the three numeric columns are what a supervisor populates.

Read across a row, the gap shows where the requirement’s effect on prices falls. At the baseline calibration the duration factor carries essentially the entire gap (the banking sector’s reported weighted sensitivity to it is about 40% below its benchmark share, having shed duration to 20% of supply against a $\chi = 1/3$ benchmark), while the two equity factors sit near benchmark. Under the proportional increase in risk weights of Figure 2 the equity gaps open as banks are priced out of regulated risk, and the norm $G = \|g\|$ climbs from 0.53 to 1.77 through the region in which the shadow cost peaks and falls. The components g_k show where in the cross-section the repricing of Proposition 3 is concentrated, and a companion market-based indicator is the family of coverage-basis spreads of Section 5.2, which price the same tilt from its other side. None of this requires the pricing distortion to be observed, since it is read directly off the sensitivities banks file.

5.4 Calibrating the pricing hit on regulated exposures

How large are these effects? Equation (25) gives the quantity to calibrate. Relative to an otherwise-identical zero-charge exposure, asset i must earn an additional expected return of

$$\text{pricing hit}_i = \chi\lambda \times \underbrace{\frac{\partial \mathcal{K}/\partial \theta_i}{p_i}}_{\text{marginal capital intensity}}, \quad (40)$$

the product of an economy-wide tightness statistic and an asset-specific intensity that can be read off the FRTB risk weights. The two pieces can be disciplined separately.

The tightness statistic $\chi\lambda$. A plausible range for the share χ of risk-bearing capacity held by constrained intermediaries is one quarter to one half in dealer-intermediated fixed-income, FX, and derivatives markets, where standardized-approach banks are the marginal investor, and less in equities. For λ , the annual excess cost of a dollar of capital tied to the requirement, three routes give a range: banks' required equity spread over the riskless rate (roughly 8–12%), the revealed-preference estimates of Kisin and Manela (2016), which imply materially lower values in calm periods, and stress-period wedges such as covered interest parity deviations (Du et al., 2018), which imply higher ones. Taking $\lambda \in [3\%, 12\%]$ and $\chi \in [1/4, 1/2]$ brackets $\chi\lambda$ between roughly 1% and 5%, with 3% a reasonable central value. Since the pricing hits below depend on χ and λ only through their product $\chi\lambda$, the calibration is robust to the split between them. Propositions 5 and 6 imply that the statistic is countercyclical, so the hits below should be read as widening in stress.

Marginal capital intensity. For a directional position, intensity is approximately the risk weight times the relevant duration (GIRR, CSR) or the risk weight itself (equity, FX, commodity), divided by $\sqrt{2}$ where the standard allows it for liquid currencies and pairs. Table S.1 in the supplementary appendix reports intensities for eleven representative exposures using the delta risk weights of the standard, together with the implied pricing hits across the $\chi\lambda$ range. All figures are medium-scenario, delta-only, and ignore the MAR21.6 scenario-max, so they are conservative.

Three observations follow. First, the magnitudes are in the range of documented post-crisis pricing wedges. At the central $\chi\lambda = 3\%$, the model implies roughly 20 basis points on ten-year duration, 50 on five-year investment-grade corporate credit, 30–45 on FX and equity-index exposure, and 40–70 at the long end. These are the same order as covered interest parity deviations (20–

40bp), negative long-end swap spreads, and funding value adjustments (Du et al., 2018; Jermann, 2020; Andersen et al., 2019). The model reads these not as anomalies of particular markets but as one cross-section indexed by marginal capital intensity. Second, the cross-sectional spread is the identifying prediction. The same $\chi\lambda$ that prices ten-year duration at 20bp prices the cross-curve basis *twin* of that duration at less than 1bp. The result is a 22-fold gap in required returns between positions with comparable true risk, which is Corollary 3 in FRTB units. Third, because every entry scales with $\chi\lambda$, the entire table widens and narrows with dealer capital tightness, giving the framework a single time-series factor and a full cross-section of exposures on which to test it.

5.5 The welfare cost of the requirement and its incidence

The closed-form value functions permit a welfare accounting. Welfare for class j is the certainty equivalent of date-1 wealth, $R_f W_0^j + V^j(\theta^j)$, compared between the binding equilibrium and the frictionless benchmark. The supply side enters as well, through the date-0 proceeds $p'\bar{\theta}$ received by the sellers of the fixed supply.

Proposition 9 (Welfare accounting). *Let $(p, \pi, \theta^U, \theta^C)$ be a binding equilibrium with multiplier $\lambda > 0$, let $\pi^0 = \bar{\eta} \Sigma \bar{\theta}$ denote the frictionless premia, and let τ be the tilt portfolio (34). Write $\|x\|_\Sigma^2 := x' \Sigma x$ and $T := \bar{\theta}' [p + \nabla K(\theta^C)]$. Relative to the frictionless equilibrium:*

- (i) *Traditional investors gain $\Delta CE^U = \frac{\chi\lambda}{\eta_U} [\bar{\eta} T + \frac{\chi\lambda}{2} \|\tau\|_\Sigma^2]$, which is positive whenever $T \geq 0$.*
- (ii) *The constrained investors' change is a scarcity gain net of a compliance cost,*

$$\Delta CE^C = \frac{\chi\lambda}{\eta_C} \left[\bar{\eta} T + \frac{\chi\lambda}{2} \|\tau\|_\Sigma^2 \right] - \frac{\lambda^2}{2\eta_C} \|\tau\|_\Sigma^2,$$

and its sign is ambiguous: positive when λ is small relative to T , negative when the requirement is tight.

- (iii) *The supply side's date-0 proceeds fall by $\chi\lambda T/R_f$.*
- (iv) *Summing (i)–(iii) at population weights and the riskless rate, every price term cancels as a transfer, and total surplus falls by the deadweight loss*

$$\text{DWL} = \frac{\chi(1-\chi)}{2\bar{\eta}} \lambda^2 \|\tau\|_\Sigma^2 \geq 0. \quad (41)$$

Two points stand out. First, the deadweight loss is *second order* in the shadow cost, growing as λ^2 while every pricing effect in Section 4 is first order in $\chi\lambda$, so price distortions are not themselves welfare statements. Second, the incidence is uneven. The unconstrained sector always gains, the constrained sector can gain as well when the scarcity premium outweighs its compliance cost, and the burden falls on the supply side through lower date-0 prices, that is, on issuers' cost of capital.

The loss also sits on a different margin from the one that cost-benefit assessments of capital usually weigh. Those exercises (Basel Committee on Banking Supervision, 2010; Firestone et al., 2019) price the *level* of capital, its funding and output costs against the crises it prevents. Equation (41) is instead a *composition* cost, present at any level of capital and set by the structure of the charge rather than its quantity, so a regulator who has already fixed the level optimally can still reduce it by recalibrating relative weights, at no cost in loss absorbency. The accounting here prices this cost margin alone, and says nothing about the default and stability benefits the model deliberately omits, against which (41) would be weighed.

6 Numerical Illustration

This section solves the model at FRTB parameters, using the scalar reduction of Proposition 4. For each candidate (λ, κ) the economy is linear and solves in closed form, and the equilibrium is the root of the two consistency conditions (29), a two-dimensional problem regardless of the number of assets and factors. At every reported solution, all of the paper's equilibrium identities (market clearing, both first-order conditions, the pricing equation of Proposition 2, and the tilt identity of Proposition 7) are verified numerically to within 10^{-13} .⁴

6.1 Calibration

The economy has $N = 4$ assets and $L = 4$ factors, of which $K = 3$ are regulated. Table S.2 in the supplementary appendix summarizes the calibrated parameters. The first asset is ten-year government duration in a liquid currency (GIRR, vertex risk weight $1.1\%/\sqrt{2}$, duration 8.5, capital intensity 6.61% per dollar). The second is an advanced-economy equity index (equity bucket 12, risk

⁴A multiplicity probe re-solves the baseline and joint-stress equilibria from 200 dispersed starting points in (λ, κ) . Every start that converges to a solution passing all equilibrium checks reaches the reported equilibrium, so no second equilibrium was found.

weight 15%), and the third is a large-cap advanced-economy equity position (bucket 5, risk weight 30%). The fourth asset loads on the one unregulated factor, f_4 , which lies outside the regulator's list. It has 5.5% systematic volatility and is positively correlated with the equity factors, but its capital charge is zero because the SBM does not measure its factor. It is the model's stand-in for the exposures Section 5.2 identifies as outside the factor list. The true index/large-cap factor correlation is 0.85, against a prescribed cross-bucket $\gamma = 0.45$ from MAR21.80, so the calibration wedge is live alongside the coverage wedge. The cross-risk-class correlation in Γ is set to zero.⁵ Investors have equal risk aversion ($\eta_U = \eta_C = 0.05$) and the bank sector has mass $m_C = 1/3$, so $\chi = 1/3$, and $R_f = 1.04$. Expected payoffs are set so that frictionless prices equal one, delivering frictionless premia of 1.05% (duration), 5.0–5.6% (equity), and 0.94% (uncharged asset). The binding threshold of Corollary 1 is $\bar{W} = 111.2$. The baseline endows banks with $W_0^C = 60$, about 55% of it.

6.2 The baseline equilibrium

Table 2 reports the solution. The shadow cost is $\lambda = 0.61\%$ per year, so the tightness statistic is $\chi\lambda = 20.5\text{bp}$, at the bottom of the back-of-the-envelope range of Section 5.4. This is a first general-equilibrium lesson, developed below. The funding wedge, common to all assets, is 20.5bp. The regulatory components on top of it are 0.5bp (duration), 2.7bp (index), 4.2bp (large-cap), and zero (uncharged asset), ordered by equilibrium marginal capital intensity as (40) requires.

Four features of the baseline deserve comment.

The coverage gap lets banks short the uncharged asset for funding. The bank sector's share of the uncharged asset is -121% of supply. Rather than holding the uncharged asset, banks *short* it, using the position both as charge-free funding (shorting raises b without raising $\mathcal{K}$) and as a charge-free hedge of their equity book, with which the unregulated factor is positively correlated. This is the general case anticipated after Proposition 7. When the unregulated factor is correlated with charged exposure (the calibration deliberately builds in this correlation, stepping outside the twin-asset independence of Corollary 3), the cheapest use of the charge-free room is to reduce

⁵The actual standard *sums* the risk-class charges, which corresponds to a higher effective cross-class correlation for same-sign books. Setting it to zero grants the bank cross-class diversification the rule denies, so the regulatory differentials below are conservative.

Table 2: Baseline equilibrium at $W_0^C = 60$: $\lambda = 0.614\%$, $\chi\lambda = 20.5\text{bp}$, $\kappa = 6.60$. “Total” is the equilibrium expected excess return less its frictionless value, and it splits into three additive parts: a funding wedge $\chi\lambda$ common to every asset, the regulatory term of (40), and a fundamental part that is true risk repriced at the higher equilibrium discount rate $R_f + \chi\lambda$. The bank share is $m_C\theta_i^C/\bar{\theta}_i$, with frictionless value $1/3$. Displayed components need not sum to the total because of rounding.

Asset	$\mathbb{E}[R_i] - R_f$ (%)		Pricing hit (bp)				Bank share
	frictionless	equilibrium	total	funding	regulatory	fundamental	
10y govt bond	1.05	1.26	21.2	20.5	0.5	0.2	0.20
Equity index	5.02	5.26	24.3	20.5	2.7	1.1	0.34
Large-cap equity	5.61	5.87	26.0	20.5	4.2	1.3	0.34
Uncharged asset	0.94	1.15	20.6	20.5	0.0	0.2	-1.21

regulated-factor risk rather than to hold it long. Both patterns, shorting the uncharged asset for funding and holding it as inventory, appear in the comparative statics below.

The calibration wedge is large where the coverage wedge is not. A hedged-equity portfolio (long \$1 large-cap, short \$1 index) has true volatility of 11.7% per dollar long, but a capital charge of 26.8 cents, which is 2.3 times its true risk. The reason is that the prescribed cross-bucket correlation (0.45) refuses most of the offset the true correlation (0.85) delivers. At the mild baseline tightness this costs the hedge only about 1.5bp of equilibrium premium, but the charge ratio is invariant to λ : in stress, the same book pays 2.3 times its risk-proportionate share of any tightening.

At statutory-scale weights, coverage shapes composition, not measured tightness. Recomputing the equilibrium with the coverage gap closed, meaning the unregulated factor is charged neutrally at its true volatility so that $K = L$, moves λ from 0.614% only to 0.619%. The coverage gap’s first-order effect at baseline is on *who holds what* and on relative prices, not on the shadow cost a supervisor would measure.

The welfare cost is an order of magnitude smaller than the effect on prices, and banks bear none of it. Evaluating Proposition 9 at the baseline gives a deadweight loss of 5.0 basis points of aggregate market value per year, an order of magnitude below the 21–26bp pricing hits, as the second-order formula (41) requires. Traditional investors gain 16.2bp, the supply side loses 21.8bp through lower date-0 prices, and the bank sector gains 0.6bp, because at baseline tightness the scarcity gain exceeds the compliance cost, so the regulated sector is a net beneficiary of its own requirement. Closing the coverage gap moves the loss only from 5.00 to 5.02bp, the welfare




Figure 1: A proportional change in all risk weights, each scaled by the common factor c . Panel (a): the shadow cost is hump-shaped (peaks marked), at both baseline and high bank equity. Panel (b): pricing hits at baseline equity rise monotonically throughout. Past the peak of λ , prices keep deteriorating while the measured shadow cost falls.

counterpart of the composition finding above.

6.3 Comparative statics at FRTB parameters

Lowering bank capital (Figure S.1 in the supplementary appendix) raises the shadow cost from zero at $\bar{W}$ to 1.17% at $W_0^C = 15$, with every pricing hit scaling alongside it. At the most stressed point the hits reach 40–50bp and the cross-sectional spread between the most and least capital-intensive assets widens from 5.4bp at baseline to 10.7bp, the model’s prediction that these price gaps comove with dealer capital, now in general-equilibrium units.

Figure 1 scales every risk weight by a common factor c and gives the paper’s clearest illustration. The shadow cost is *hump-shaped*, peaking at $c \approx 17$ at baseline equity and declining thereafter, and peaking near $c \approx 23$ at high bank equity ($W_0^C = 104$, just enough to fund the frictionless book with nothing left for the charge). Yet the pricing hits rise *monotonically* throughout, reaching roughly 200–230bp on the equity assets at $c = 30$ while duration and the uncharged asset stabilize near 35–50bp (their residual hit is mostly the funding wedge). This is the general-equilibrium, multi-asset confirmation of Proposition 6(iii) and Remark 4. The quantity side tells the same story, with the bank’s share of the charged assets collapsing (from 20–34% of supply at $c = 1$ to 2–10% at $c = 30$) and its position in the uncharged asset swinging from –121% of supply to +11%. The bank’s short position in the uncharged asset at the baseline turns into a long position once punitive weights

price banks out of regulated risk entirely, as anticipated in Section 5.1.

The exercise yields two lessons. First, the equilibrium tightness statistic is modest at plausible parameters ($\chi\lambda$ stays in the 20–40bp range as risk weights are scaled up) because prices adjust, the bank sector shrinks its book, and the taxonomy’s uncharged margins absorb part of the burden. Reaching the large end of the back-of-the-envelope range of Section 5.4 requires impairing the *unconstrained* sector’s capacity as well, the experiment of Section 6.4. Second, the cross-section, not the level, is where the framework’s distinctive predictions are clearest. The ordering of the regulatory components, the hedge-penalty ratio, and the switch of the uncharged asset from a short to a long position are sign-definite implications that survive at every point of both the proportional change in risk weights and the change in bank capital.

6.4 Stress: impairing the unconstrained sector

The distortions above are computed in an economy whose unconstrained investors stand ready to absorb whatever banks shed. Stress episodes are the opposite configuration. Dealer capital falls at the same time as the risk-bearing capacity of the nonbank sector, which is the pattern of March 2020 in dealer-intermediated markets. [Becker and Benmelech \(2025\)](#) document the behavior of the United States corporate bond market in that episode. Table 3 recomputes the equilibrium under two impairments, separately and together: bank equity halved ($W_0^C = 30$) and unconstrained risk-bearing capacity cut to a third (η_U raised from 0.05 to 0.15). Fundamentals (μ , Σ , supplies, risk weights) are held at their baseline values, and each scenario’s pricing hits are measured against that scenario’s own frictionless benchmark, so the reported distortions isolate the contribution of the capital requirement under the stressed conditions.⁶

Halving bank equity alone moves the tightness statistic from 20 to 33 basis points: painful but modest, because the deep unconstrained sector keeps absorbing the tilt at little price concession. Impairing the unconstrained sector changes the arithmetic on both margins of $\chi\lambda$ at once. The banking sector’s share of the economy’s remaining risk-bearing capacity χ rises from one third to 0.60, and the shadow cost λ jumps to 2.8% because the residual investors now demand far more

⁶The stressed equilibria are computed by continuation. The solver walks W_0^C down from just below the scenario’s binding threshold, warm-starting each step, which keeps it on the economically relevant branch. All equilibrium identities are verified at the reported solutions.

Table 3: Stress scenarios. Each scenario is solved at baseline fundamentals with the indicated preferences and bank equity. Hits are measured against the scenario’s own frictionless benchmark. The coverage spread is the pricing-hit gap between the most capital-intensive asset (large-cap equity) and the uncharged asset. The last column is the bank sector’s position in the uncharged asset as a share of its supply.

Scenario	η_U	W_0^C	$\bar{W}$	λ	$\chi\lambda$ (bp)	Largest hit (bp)	Coverage spread (bp)	Uncharged share
Baseline	0.05	60	111.2	0.61%	20	26	5.4	−1.21
Bank stress	0.05	30	111.2	0.98%	33	42	8.9	−2.13
Nonbank impairment	0.15	60	196.5	2.81%	168	228	56.5	−3.57
Joint stress	0.15	30	196.5	3.46%	208	282	70.8	−4.52

compensation to hold what banks shed. The tightness statistic reaches 168 basis points, and 208 in the joint scenario. Pricing hits on the charged assets reach 214–282 basis points.

Three features of the stressed equilibria deserve emphasis. First, the coverage spread widens from 5.4 basis points at baseline to 70.8 under joint stress, a factor of 13. The model’s most identifiable prediction, that on-list/off-list spreads between comparably risky exposures comove with intermediation capacity, is most visible in stress, just when the corresponding empirical objects (sudden widening of such spreads, dealer balance-sheet congestion) are observed. Second, the bank’s short position in the uncharged asset deepens rather than closing, growing from −121% of supply to −452%, because charge-free funding is most valuable when capital is scarcest. The sensitivity gap of Section 5.3 would register both movements directly, while the headroom-based indicators a supervisor conventionally watches would register only the moderate movement in λ . Third, the welfare accounting changes sign for the banking sector. The deadweight loss (41) rises from 5 basis points of market value at baseline to 13 under bank stress and 91 under joint stress, about a fifth of the aggregate risk premium, and in both scenarios with $W_0^C = 30$ the compliance cost overtakes the scarcity gain, so banks lose on net even as the unconstrained sector’s gain grows.

6.5 The sensitivity gap as risk weights rise, and a recalibration counterfactual

Two final exercises connect the model’s supervisory content to the calibrated economy: the sensitivity gap of Section 5.3 computed as risk weights are scaled up, and the recalibration counterfactual promised by the wedge decomposition (10).

The sensitivity gap tracks the distortion where the shadow cost does not. Figure 2 recomputes




Figure 2: The sensitivity gap under the proportional change in risk weights of Figure 1. Panel (a): the shadow cost λ (solid, left axis) is hump-shaped with its peak at $c \approx 17$. The sensitivity gap G of Definition 2 (dashed, right axis) rises monotonically throughout. Panel (b): the largest pricing hit plotted against G . The dot marks the λ peak, which is invisible in the relationship. To keep units comparable across c , G is evaluated at the statutory ($c = 1$) weights. Evaluating it at the scaled weights $c D_w$ only steepens the profile. All identities verified at every grid point.

the risk-weight scaling of Figure 1 and reports the gap G alongside λ and the largest pricing hit. The picture is Proposition 8 in operation. The shadow cost peaks at $c \approx 17$ and declines thereafter, while G rises monotonically from 0.53 at statutory weights to 1.77 at $c = 30$, tracking the pricing hits one for one. Panel (b) shows the hit as a strictly increasing function of G along the entire range of risk weights, with the λ peak an unremarkable interior point. A supervisor plotting reported weighted sensitivities against a market benchmark would see the requirement's effect on prices growing throughout the region in which every headroom-based indicator reports it fading.

Recalibration repairs the calibration wedge, at essentially zero cost in measured tightness, and cannot touch the coverage wedge. The counterfactual replaces the statutory parameters with the neutral calibration of Remark 1: risk weights proportional to the true volatilities of the three regulated factors, prescribed correlations equal to the true ones, and the scale ϕ chosen so that the charge on the frictionless bank book is unchanged ($\phi = 1.01$, delivering weights of 6.0%, 16.2%, and 19.2% against the statutory 6.6%, 15%, and 30%). At baseline equity, measured tightness barely moves: λ goes from 0.614% to 0.617% and $\chi\lambda$ from 20.5 to 20.6 basis points. The composition changes are what matter. The charge on the beta-matched hedged-equity book of Section 6.2 falls from 2.3 times its true risk to 0.86 (slightly below parity because the neutral weights charge systematic risk only), the regulatory premia of the two equity assets converge toward their true risk

shares (4.2 and 2.7 basis points become 3.4 and 3.1), and the cross-sectional spread of regulatory components among the charged assets narrows from 3.7 to 3.0 basis points. Under the joint-stress scenario of Table 3, where composition matters most, the largest pricing hit falls from 282 to 272 basis points, the coverage spread from 70.8 to 60.4, and the regulatory spread narrows by roughly a fifth (45.8 to 36.8 basis points). What does not move, in any scenario, is the coverage margin. The bank sector’s position in the uncharged asset (-121% of supply at baseline, -452% under joint stress) and the baseline deadweight loss (5.00 versus 5.02 basis points of market value) are unchanged with either set of weights. Reweighting removes the distortions the weights create and leaves untouched the ones the factor list creates.

7 Policy Implications

This paper does not argue against risk-sensitive capital requirements. The question the model answers is: given that trading books will be capitalized through a sensitivity-based formula, what does the formula’s design do to markets, and what should the officials who calibrate and supervise it watch? Some suggestions include the following.

Monitor quantities along with shadow prices. The natural indicators of regulatory tightness, such as internal shadow prices on capital, track λ . Proposition 6 and Figure 1 show that λ can be misleading since it peaks and then falls as risk weights tighten, because banks respond by shrinking the charged book, and it is smallest where its effect on prices is largest. The model suggests measuring quantities along with shadow prices. The sensitivity gap of Definition 2, which by Proposition 8 is a linear image of the same equilibrium object that drives every pricing distortion, and which Figure 2 shows rising monotonically through the entire region in which the shadow cost peaks and falls. Generating these data has a zero marginal cost since the weighted sensitivities WS_k that enter the gap are the inputs to the SBM calculation itself, reported at each risk factor, currency, and bucket, so a supervisor can aggregate them across standardized-approach banks today. A complementary market-based indicator is the family of coverage-basis spreads (on-/off-the-run, cash-futures, bond-CDS), which the model predicts widen with dealer capital tightness and which, unlike shadow prices, are observable daily.

Risk-weight calibration alters asset prices. A change in a delta risk weight is conventionally scored by its effect on capital ratios. Equation (40) says that scoring by its effect on the cross-section of expected returns is also helpful. Every reweighting moves relative prices economy-wide in proportion to $\chi\lambda$, and the incidence falls on issuers and unconstrained investors as well as on banks. The decomposition of the regulatory wedge (10) gives calibration reviews a concrete objective. The calibration component, meaning deviations of prescribed weights and correlations from neutral relative risk, is removable without changing the aggregate stringency of the regime. Scaling all weights up or down moves only the Black-type level effect, while the relative structure determines who holds what and at what price. The hedged-equity example of Section 6 quantifies what is at stake in one prescribed parameter. At the standard’s cross-bucket correlation of 0.45, against a true correlation of 0.85, a beta-matched long–short equity book is charged 2.3 times its true risk.

The coverage margin, and the design of the residual risk add-on. The coverage wedge poses more of a regulatory challenge. Whatever the weights, risk outside the prescribed factor list is free on constrained balance sheets, and Section 5.1 predicts that it will migrate there. The taxonomy of the SBM is therefore a decision about which risks the regulated sector will be induced to accumulate. This has at least two consequences. First, the residual risk add-on, the framework’s existing patch, levies flat percentages of portfolio value, so within its coverage the marginal charge is again invariant to risk and the model predicts selection toward the most true risk per unit of notional within asset classes. Making the add-on sensitive to a risk proxy, even crudely, would blunt it. Second, taxonomy reviews could be profitably informed by holdings. Sustained concentration of standardized-approach banks in instruments with low marginal charges relative to their volatility is a sign that a coverage gap is being exploited, and could be detectable in the same supervisory sensitivity data described above.

What the framework does not say. This model does not address the optimal level of capital, or default and its externalities. None of the results here weigh against capital adequacy as such. The model shows that holding the level of capital fixed, the structure of the charge, meaning its relative weights, its prescribed correlations, and its coverage, determines an equilibrium pattern of

relative prices and risk locations. That pattern is currently unmonitored, and the data to monitor it already exist inside the regime. The scenario-max layer of MAR21.6 adds a further and different distortion on top of the smooth structure priced here. This points to the conclusion that the functional form of prudence, and not only its quantity, is a policy variable.

8 Conclusion

The sensitivities-based method makes the covariance matrix a regulatory instrument. This paper works out what that choice does to asset prices in a transparent equilibrium model. Its main implication for supervision is that the shadow cost of the requirement is hump-shaped in the rule's stringency while the pricing distortions rise monotonically, so the indicators that track the shadow cost understate the requirement where its effect on prices is largest. The sensitivity gap corrects this, and it can be computed from data the regime already collects.

Several extensions remain. A dynamic version would let the procyclicality of $\chi\lambda$ interact with capital accumulation, and adding the vega, curvature, and default-risk layers would bring the model closer to the standard as written. The sharpest open question is empirical, namely whether on-list and off-list spreads between comparably risky exposures move together with dealer capital, which can be examined with dealer position data and the supervisory sensitivities the FRTB already generates. The broader point is simple. When the marginal investor's capital depends on a risk model the regulator writes, differences between that model and underlying risk will affect asset prices.

Supplementary Appendix

This supplementary appendix collects the proofs of all results in the main text, together with one result stated here in full (Corollary S.1), a supplementary figure, and supplementary tables. Unprefixed numbers for lemmas, propositions, theorems, corollaries, equations, sections, figures, and assumptions refer to the main body. Items numbered with the prefix S (Corollary S.1, Figure S.1, Table S.1, and so on) are internal to this appendix.

Proofs

Proof of Proposition 1. Fix the benchmark scale $\phi > 0$. By Remark 1, $\Sigma - \Sigma_\varepsilon = M\Sigma_f M'$. Let P be the orthogonal projector onto $\text{col}(M_r)$ and $P_\perp = I - P$, and write $M_u^\perp := P_\perp M_u$ for the projection of the unregulated loadings onto the complement of the regulated span. Every admissible regulatory matrix $\Sigma_K = M_r \Omega M_r'$ (with $\Omega = D_w \Gamma D_w \succeq 0$) has row and column spaces contained in $\text{col}(M_r)$, so $P_\perp \Sigma_K = \Sigma_K P_\perp = 0$, and therefore $P_\perp \Delta P_\perp = -\phi^2 P_\perp M \Sigma_f M' P_\perp$. Since $P_\perp M_r = 0$, the three terms of

$$M \Sigma_f M' = M_r \Sigma_{rr} M_r' + M_r \Sigma_{ru} M_u' + M_u \Sigma_{ur} M_r' + M_u \Sigma_{uu} M_u'$$

that carry an outer M_r vanish under the compression, leaving $P_\perp \Delta P_\perp = -\phi^2 M_u^\perp \Sigma_{uu} M_u^{\perp'}$. A projector compression does not increase any unitarily invariant norm ($\|P_\perp\| = 1$), so $\|\Delta\| \geq \|P_\perp \Delta P_\perp\|$, which is (11). Its right-hand side depends on M_u and Σ_{uu} alone, not on (w, Γ) . The matrix $M_u^\perp \Sigma_{uu} M_u^{\perp'}$ is positive semidefinite, and since $\Sigma_{uu} \succ 0$ (a diagonal block of $\Sigma_f \succ 0$) it is nonzero if and only if $M_u^\perp \neq 0$, i.e. $\text{col}(M_u) \not\subseteq \text{col}(M_r)$. In the complementary case $\text{col}(M_u) \subseteq \text{col}(M_r)$ the regulated factors span all systematic risk, and the neutral calibration of Remark 1 attains $\Sigma_K = \phi^2(\Sigma - \Sigma_\varepsilon)$, so $\Delta = 0$. $\square$

Proof of Lemma 1. (i) V^U is strictly concave ($\Sigma \succ 0$) and coercive, so the first-order condition $\mu - R_f p - \eta_U \Sigma \theta = 0$ is necessary and sufficient, giving $\theta^U = \eta_U^{-1} \Sigma^{-1} \pi$.

(ii) The feasible set $C := \{\theta : p'\theta + \mathcal{K}(\theta) \leq W_0^C\}$ is closed and convex ($\mathcal{K}$ is a seminorm, hence convex and continuous) and contains $\theta = 0$ with strict inequality since $W_0^C > 0$, so Slater's condition holds. V^C is strictly concave and coercive, so a unique maximizer $\theta^C \in C$ exists, and by Slater there is $\lambda \geq 0$ such that θ^C maximizes the Lagrangian $V^C(\theta) - \lambda(p'\theta + \mathcal{K}(\theta) - W_0^C)$ over $\mathbb{R}^N$.

and complementary slackness holds. The Lagrangian is concave, and its supergradient condition is

$$0 \in \pi - \eta_C \Sigma \theta^C - \lambda p - \lambda \partial \mathcal{K}(\theta^C).$$

Wherever $\mathcal{K}(\theta^C) = \kappa > 0$ the charge is differentiable with $\partial \mathcal{K}(\theta^C) = \{\Sigma_{\mathcal{K}} \theta^C / \kappa\}$ by (8), so

$$\pi - \lambda p = \left(\eta_C \Sigma + \frac{\lambda}{\kappa} \Sigma_{\mathcal{K}} \right) \theta^C = H(\lambda, \kappa) \theta^C.$$

Since $\Sigma \succ 0$, $\Sigma_{\mathcal{K}} \succeq 0$, $\eta_C > 0$, and $\lambda/\kappa \geq 0$, it follows that $H(\lambda, \kappa) \succ 0$. Inverting and using $\pi - \lambda p = \mu - (R_f + \lambda)p$ gives (21). $\square$

Proof of Proposition 2. From the constrained first-order condition, $\eta_C \Sigma \theta^C = \pi - \lambda p - (\lambda/\kappa) \Sigma_{\mathcal{K}} \theta^C$, hence

$$\theta^C = \frac{1}{\eta_C} \Sigma^{-1} \left[\pi - \lambda p - \frac{\lambda}{\kappa} \Sigma_{\mathcal{K}} \theta^C \right].$$

Substitute this and $\theta^U = \eta_U^{-1} \Sigma^{-1} \pi$ into market clearing (22):

$$\left(\frac{m_U}{\eta_U} + \frac{m_C}{\eta_C} \right) \Sigma^{-1} \pi - \frac{m_C}{\eta_C} \Sigma^{-1} \left[\lambda p + \frac{\lambda}{\kappa} \Sigma_{\mathcal{K}} \theta^C \right] = \bar{\theta}.$$

Multiplying through by $\bar{\eta} \Sigma$ and using $\chi = m_C \bar{\eta} / \eta_C$ yields (23). Since $\bar{\eta}^{-1} = m_U / \eta_U + m_C / \eta_C > m_C / \eta_C$, it follows that $\chi \in (0, 1)$. Equation (24) follows by substituting $\pi = \mu - R_f p$ into (23) and solving for p (the coefficient $R_f + \chi \lambda$ is positive). When $\lambda = 0$ the two distortion terms vanish. $\square$

Proof of Proposition 3. Divide row i of (23) by $p_i > 0$ and use $\mathbb{E}[R_i] - R_f = \pi_i / p_i$. For the fundamental term,

$$\frac{\bar{\eta} (\Sigma \bar{\theta})_i}{p_i} = \bar{\eta} p_M \frac{\text{Cov}(X_i, \bar{\theta}' X)}{p_i p_M} = \bar{\eta} p_M \text{Cov}(R_i, R_M) = \beta_i^M \cdot \bar{\eta} p_M \text{Var}(R_M) = \beta_i^M \pi_M.$$

For the funding term, $\chi \lambda p_i / p_i = \chi \lambda$. For the regulatory term, by the definitions of $\text{Cov}_{\mathcal{K}}$ and $\text{Var}_{\mathcal{K}}$,

$$\frac{\chi \lambda}{\kappa} \frac{(\Sigma_{\mathcal{K}} \theta^C)_i}{p_i} = \frac{\chi \lambda}{\kappa} p_C \text{Cov}_{\mathcal{K}}(R_i, R_C) = \beta_i^{\mathcal{K}} \cdot \frac{\chi \lambda}{\kappa} p_C \text{Var}_{\mathcal{K}}(R_C) = \beta_i^{\mathcal{K}} \frac{\chi \lambda \kappa}{p_C},$$

where the last equality uses $\text{Var}_{\mathcal{K}}(R_C) = \theta^{C'} \Sigma_{\mathcal{K}} \theta^C / p_C^2 = \kappa^2 / p_C^2$. $\square$

Proof of Corollary 1. In any equilibrium with $\lambda = 0$, both sectors' demands are proportional to $\Sigma^{-1} \pi$, so clearing pins down $\pi = \bar{\eta} \Sigma \bar{\theta}$, prices $p^0 = R_f^{-1}(\mu - \bar{\eta} \Sigma \bar{\theta})$, and $\theta^C = (\bar{\eta} / \eta_C) \bar{\theta}$. This

candidate is an equilibrium if and only if θ^C is feasible at $\lambda = 0$, i.e. $p^{0r}\theta^C + \mathcal{K}(\theta^C) \leq W_0^C$. By homogeneity of degree one of $\mathcal{K}$, $\mathcal{K}((\bar{\eta}/\eta_C)\bar{\theta}) = (\bar{\eta}/\eta_C)\mathcal{K}(\bar{\theta})$, which gives (27). Conversely, if (27) fails, no equilibrium can have $\lambda = 0$ (the unique $\lambda = 0$ candidate is infeasible), so every equilibrium has $\lambda > 0$. $\square$

The next result, referenced in Section 4.3 of the main text, states the Black-type benchmark in closed form.

Corollary S.1 (A covering, neutrally calibrated rule raises returns in proportion to size). *Suppose $\Sigma_{\mathcal{K}} = \phi^2 \Sigma$ for some $\phi > 0$. In any binding equilibrium ($\lambda > 0$, $\kappa = \mathcal{K}(\theta^C) > 0$), expected returns obey the exact one-factor relation*

$$\mathbb{E}[R_i] - R_f = b + a p_M \text{Cov}(R_i, R_M), \quad i = 1, \dots, N,$$

where, with $\tilde{a} := \eta_C + \lambda\phi^2/\kappa$ and $\psi := \chi\lambda\phi^2/(\kappa\tilde{a}) \in (0, 1)$,

$$a = \frac{\bar{\eta}}{1 - \psi} > \bar{\eta}, \quad b = \frac{\lambda\chi\eta_C}{\tilde{a}(1 - \psi)} > 0.$$

The requirement raises the zero-beta intercept and steepens the security market line but leaves relative valuations undistorted.

Proof of Corollary S.1. With $\Sigma_{\mathcal{K}} = \phi^2 \Sigma$ and $\kappa > 0$, $H(\lambda, \kappa) = \tilde{a} \Sigma$ with $\tilde{a} := \eta_C + \lambda\phi^2/\kappa$, so (21) gives $\Sigma\theta^C = (\pi - \lambda p)/\tilde{a}$. Substituting into (23),

$$\pi = \bar{\eta} \Sigma \bar{\theta} + \chi \lambda p + \psi (\pi - \lambda p), \quad \psi := \frac{\chi \lambda \phi^2}{\kappa \tilde{a}}.$$

Since $\chi < 1$ and $\lambda\phi^2/\kappa < \tilde{a}$, it follows that $\psi \in (0, 1)$. Collecting terms,

$$\pi = a \Sigma \bar{\theta} + b p, \quad a := \frac{\bar{\eta}}{1 - \psi} > \bar{\eta}, \quad b := \frac{\lambda(\chi - \psi)}{1 - \psi} = \frac{\lambda\chi\eta_C}{\tilde{a}(1 - \psi)} > 0,$$

using $\chi - \psi = \chi(1 - \lambda\phi^2/(\kappa\tilde{a})) = \chi\eta_C/\tilde{a}$. Dividing row i by p_i and using $(\Sigma\bar{\theta})_i/p_i = p_M \text{Cov}(R_i, R_M)$ gives the return-space statement. $\square$

Proof of Corollary 2. $M_r' z = 0$ implies $\Sigma_{\mathcal{K}} z = M_r \Omega M_r' z = 0$, and by symmetry $z' \Sigma_{\mathcal{K}} \theta^C = (\Sigma_{\mathcal{K}} z)' \theta^C = 0$, so $\beta_z^{\mathcal{K}} = 0$. Premultiplying (23) by z' gives $z' \pi = \bar{\eta} z' \Sigma \bar{\theta} + \chi \lambda z' p$. Dividing by $p' z \neq 0$ and writing the result in return form yields the claim. $\square$

Proof of Proposition 4. Fix $(\lambda, \kappa) \in [0, \infty) \times (0, \infty)$ and write $H = H(\lambda, \kappa) \succ 0$. Solving clearing (22) for θ^C and substituting into the constrained first-order condition $H\theta^C = \pi - \lambda p = (1 + \lambda/R_f) \pi - (\lambda/R_f) \mu$ (using $p = (\mu - \pi)/R_f$) gives

$$A(\lambda, \kappa) \pi = H \bar{\theta} + \frac{m_C \lambda}{R_f} \mu, \quad A(\lambda, \kappa) := m_C \left(1 + \frac{\lambda}{R_f}\right) I + \frac{m_U}{\eta_U} H \Sigma^{-1}.$$

The matrix $H \Sigma^{-1}$ is similar to $\Sigma^{-1/2} H \Sigma^{-1/2} \succ 0$, so its eigenvalues are real and positive. Hence every eigenvalue of A is at least $m_C(1 + \lambda/R_f) > 0$ and A is invertible. Thus

$$\pi(\lambda, \kappa) = A(\lambda, \kappa)^{-1} \left[H \bar{\theta} + \frac{m_C \lambda}{R_f} \mu \right], \quad p = \frac{\mu - \pi}{R_f}, \quad \theta^U = \frac{\Sigma^{-1} \pi}{\eta_U}, \quad \theta^C = \frac{\bar{\theta} - m_U \theta^U}{m_C},$$

uniquely. Given this map, a binding-constraint equilibrium requires that (a) the conjectured κ equal the realized charge, $\mathcal{K}(\theta^C(\lambda, \kappa)) = \kappa$, so that (21) is the true optimality condition, and (b) the constraint hold with equality, $p' \theta^C + \kappa = W_0^C$, consistent with $\lambda > 0$ and complementary slackness. Conversely any (λ, κ) solving (29) generates portfolios and prices satisfying optimality (Lemma 1) and market clearing, hence an equilibrium. $\square$

Proof of Theorem 1. Step 1: reformulation as a fixed point. For $\theta \in \mathbb{R}^N$ define

$$\pi(\theta) := \frac{\eta_U}{m_U} \Sigma (\bar{\theta} - m_C \theta), \quad p(\theta) := R_f^{-1} [\mu - \pi(\theta)],$$

and let $C(\theta) := \{\theta' \in \mathbb{R}^N : p(\theta)' \theta' + \mathcal{K}(\theta') \leq W_0^C\}$ and $V(\theta'; \theta) := \theta'' \pi(\theta) - \frac{\eta_C}{2} \theta'' \Sigma \theta'$. Since $V(\cdot; \theta)$ is strictly concave and coercive and $C(\theta)$ is closed, convex, and nonempty ($0 \in C(\theta)$), the best response

$$\text{BR}(\theta) := \arg \max_{\theta' \in C(\theta)} V(\theta'; \theta)$$

is well defined and single-valued. The claim is that (p, θ^U, θ^C) is an equilibrium if and only if $\theta^C = \text{BR}(\theta^C)$, with $p = p(\theta^C)$ and $\theta^U = (\bar{\theta} - m_C \theta^C)/m_U$. Indeed, in any equilibrium the traditional investors' first-order condition $\pi = \eta_U \Sigma \theta^U$ combined with clearing $\theta^U = (\bar{\theta} - m_C \theta^C)/m_U$ gives $\pi = \pi(\theta^C)$, so the bank optimizes against $p(\theta^C)$ and $\theta^C = \text{BR}(\theta^C)$. Conversely, a fixed point delivers bank optimality by construction, traditional-investor optimality since $\eta_U \Sigma \theta^U = \pi(\theta^C)$, and clearing by the definition of θ^U .

Step 2: a priori bound. Suppose $\theta^* = \text{BR}(\theta^*)$, or more generally that θ^* maximizes $V(\cdot; \theta^*)$

over any subset of $C(\theta^*)$ containing 0. Then $V(\theta^*; \theta^*) \geq V(0; \theta^*) = 0$, i.e.

$$\frac{\eta_C}{2} \|\theta^*\|_\Sigma^2 \leq \theta^{*\prime} \pi(\theta^*) = \frac{\eta_U}{m_U} \left[\theta^{*\prime} \Sigma \bar{\theta} - m_C \|\theta^*\|_\Sigma^2 \right].$$

Collecting terms and applying the Cauchy–Schwarz inequality in the inner product $\langle x, y \rangle_\Sigma := x' \Sigma y$,

$$\left(\frac{\eta_C}{2} + \frac{\eta_U m_C}{m_U} \right) \|\theta^*\|_\Sigma^2 \leq \frac{\eta_U}{m_U} \|\theta^*\|_\Sigma \|\bar{\theta}\|_\Sigma,$$

which gives $\|\theta^*\|_\Sigma \leq \rho^*$ as in (31).

Step 3: Brouwer on the truncated best response. Let $\rho := \rho^* + 1$ and $B := \{\theta : \|\theta\|_\Sigma \leq \rho\}$, a compact convex set with 0 in its interior. Define $\text{BR}_\rho(\theta) := \arg \max_{\theta' \in C(\theta) \cap B} V(\theta'; \theta)$ for $\theta \in B$; this is single-valued (strict concavity, nonempty compact convex feasible set) and maps B into B . The correspondence $\theta \mapsto C(\theta) \cap B$ is compact- and convex-valued and continuous: upper hemicontinuity follows from the closed graph of $\{(\theta', \theta) : g(\theta', \theta) \leq 0\}$, $g(\theta', \theta) := p(\theta)' \theta' + \mathcal{K}(\theta') - W_0^C$, intersected with the compact B . For lower hemicontinuity, take $\theta_n \rightarrow \theta$ and $\theta' \in C(\theta) \cap B$, and set $\theta'_n := t_n \theta'$ with $t_n := \min\{1, W_0^C / (W_0^C + \max\{g(\theta', \theta_n), 0\})\}$. By convexity of $g(\cdot, \theta_n)$ and $g(0, \theta_n) = -W_0^C < 0$,

$$g(t_n \theta', \theta_n) \leq t_n g(\theta', \theta_n) - (1 - t_n) W_0^C \leq 0,$$

so $\theta'_n \in C(\theta_n) \cap B$ (the ball is preserved since $t_n \leq 1$), and $t_n \rightarrow 1$ because $g(\theta', \theta_n) \rightarrow g(\theta', \theta) \leq 0$. Hence $\theta'_n \rightarrow \theta'$. Since V is jointly continuous, Berge's maximum theorem implies BR_ρ is continuous, and Brouwer's fixed-point theorem yields $\theta^* = \text{BR}_\rho(\theta^*) \in B$.

Step 4: the truncation does not bind. The fixed point θ^* maximizes $V(\cdot; \theta^*)$ over $C(\theta^*) \cap B \ni 0$, so Step 2 gives $\|\theta^*\|_\Sigma \leq \rho^* < \rho$: θ^* is interior to B . If some $\theta' \in C(\theta^*)$ had $V(\theta'; \theta^*) > V(\theta^*; \theta^*)$, then points on the segment $[\theta^*, \theta']$ near θ^* would lie in $C(\theta^*) \cap B$ (convexity of C , interiority in B) and, by concavity of V , have strictly higher value than θ^* , a contradiction. Hence $\theta^* = \text{BR}(\theta^*)$ and, by Step 1, an equilibrium exists. Finally, by Corollary 1, if (27) holds the slack CAPM allocation is an equilibrium, and if it fails every equilibrium, including the one just constructed, has $\lambda > 0$. $\square$

Proof of Proposition 5. Monotonicity is immediate since the feasible set $C(W) := \{\theta : p' \theta + \mathcal{K}(\theta) \leq W\}$ expands with W . For concavity, let θ_i be optimal at W_i ($i = 1, 2$) and $t \in [0, 1]$. By convexity of $\theta \mapsto p' \theta + \mathcal{K}(\theta)$, the portfolio $t\theta_1 + (1 - t)\theta_2$ is feasible at $tW_1 + (1 - t)W_2$, and concavity of V^C

gives $v(tW_1 + (1-t)W_2) \geq tv(W_1) + (1-t)v(W_2)$.

For the supergradient property, fix $W > 0$ with optimum $\theta(W)$ and multiplier $\lambda(W)$. Slater's condition ($\theta = 0$) implies the saddle-point property: for all $\theta' \in \mathbb{R}^N$,

$$V^C(\theta') - \lambda(W) [p'\theta' + \mathcal{K}(\theta') - W] \leq V^C(\theta(W)) - \lambda(W) [p'\theta(W) + \mathcal{K}(\theta(W)) - W] = v(W),$$

using complementary slackness for the equality. If θ' is feasible at W' , then $p'\theta' + \mathcal{K}(\theta') - W \leq W' - W$, so $V^C(\theta') \leq v(W) + \lambda(W)(W' - W)$. Taking the supremum over θ' feasible at W' gives $v(W') \leq v(W) + \lambda(W)(W' - W)$.

Monotonicity of multipliers follows: for $W < W'$, the supergradient inequalities at W and W' give

$$\lambda(W')(W' - W) \leq v(W') - v(W) \leq \lambda(W)(W' - W),$$

hence $\lambda(W') \leq \lambda(W)$. A concave function is differentiable almost everywhere, and at points of differentiability the supergradient is unique, so $v'(W) = \lambda(W)$ there. $\square$

Proof of Proposition 6. Throughout, “premium” means $\mu - R_f p$. Note $W < \bar{W}$ requires $\bar{W} > 0$, hence $p^0 + w > 0$.

Step 0: any equilibrium has $\theta^C > 0$, $p + w > 0$, and a positive premium. In a slack equilibrium ($\lambda = 0$), Corollary 1 pins the CAPM allocation, which is feasible if and only if $W \geq \bar{W}$. Since $W < \bar{W}$, every equilibrium is binding, $\lambda > 0$, with the constraint holding with equality: $p\theta^C + w|\theta^C| = W > 0$. This rules out $\theta^C = 0$. If $\theta^C < 0$, then $(p - w)\theta^C = W > 0$ forces $p < w$. Clearing gives $\theta^U = (\bar{\theta} - m_C\theta^C)/m_U > \bar{\theta}/m_U > 0$, so the premium $\eta_U\sigma^2\theta^U$ is strictly positive, and the bank's first-order condition $\mu - R_f p - \eta_C\sigma^2\theta^C = \lambda(p - w)$ has a strictly positive left side and a strictly negative right side, a contradiction. Hence $\theta^C > 0$, the binding constraint reads $(p + w)\theta^C = W$, so $p + w > 0$, and the bank's first-order condition $\mu - R_f p = \eta_C\sigma^2\theta^C + \lambda(p + w) > 0$ gives a positive premium and $\theta^U > 0$.

Step 1: uniqueness and (33). By Step 0, any equilibrium satisfies $\theta^C = W/(p + w)$ and $\theta^U = (\mu - R_f p)/(\eta_U\sigma^2)$, so clearing is (32). On $(-w, \infty)$ the left side of (32) is strictly decreasing (both terms are), tends to $+\infty$ as $p \downarrow -w$, and to $-\infty$ as $p \rightarrow \infty$. Hence there is a unique root p^* .

Substituting clearing into the bank's first-order condition,

$$\lambda(p^* + w) = \mu - R_f p^* - \eta_C \sigma^2 \theta^C = \frac{\eta_U \sigma^2}{m_U} (\bar{\theta} - m_C \theta^C) - \eta_C \sigma^2 \theta^C = \sigma^2 \left[\frac{\eta_U}{m_U} \bar{\theta} - \left(\frac{\eta_U m_C}{m_U} + \eta_C \right) \theta^C \right],$$

which is (33). Since $(\eta_U/m_U)/(\eta_U m_C/m_U + \eta_C) = \bar{\eta}/\eta_C$, it follows that $\lambda > 0$ if and only if $\theta^C < (\bar{\eta}/\eta_C) \bar{\theta}$. To verify this inequality, note that at $W = \bar{W}$ the pair $(p, \theta^C) = (p^0, (\bar{\eta}/\eta_C) \bar{\theta})$ solves (32) (direct substitution, using $\mu - R_f p^0 = \bar{\eta} \sigma^2 \bar{\theta}$ and $1/\bar{\eta} = m_U/\eta_U + m_C/\eta_C$). By Step 2 below, θ^C is strictly increasing in W . Since $W < \bar{W}$, it follows that $\theta^C < (\bar{\eta}/\eta_C) \bar{\theta}$ and $\lambda > 0$. The candidate then satisfies the sufficient KKT conditions of the concave bank problem, so it is an equilibrium, and by the uniqueness of p^* it is the only one.

Step 2: comparative statics in W . Raising W raises the left side of (32) pointwise on $(-w, \infty)$. Since that function is strictly decreasing in p , the root p^* strictly increases. Hence $\theta^U = (\mu - R_f p^*)/(\eta_U \sigma^2)$ and the premium strictly decrease, and $\theta^C = (\bar{\theta} - m_U \theta^U)/m_C$ strictly increases. In (33), the bracket is positive and strictly decreasing in W (via θ^C), and the factor $1/(p^* + w)$ is strictly decreasing in W (via p^*). Hence λ is strictly decreasing.

Step 3: comparative statics in w . At fixed p , raising w lowers the left side of (32), so p^* strictly decreases, θ^U and the premium strictly increase, and θ^C strictly decreases (clearing). For the limit, Step 0 bounds the premium by $\eta_U \sigma^2 \bar{\theta}/m_U$ from above (since $\theta^C > 0$ implies $\theta^U < \bar{\theta}/m_U$), so p^* is bounded below and $p^* + w \rightarrow \infty$ as $w \rightarrow \infty$. Then $\theta^C = W/(p^* + w) \rightarrow 0$, $\theta^U \rightarrow \bar{\theta}/m_U$, and the premium converges to $\eta_U \sigma^2 \bar{\theta}/m_U$. From (33), $0 < \lambda \leq \sigma^2 (\eta_U/m_U) \bar{\theta}/(p^* + w) \rightarrow 0$. Finally, $\bar{W}$ is strictly increasing in w , so the binding region is $\{w > \bar{w}\}$ with $\bar{W} = W$ at $w = \bar{w}$. The root $p^*(w)$ varies continuously (a strictly monotone, continuous family of clearing equations), and at $w = \bar{w}$ Step 1 gives $\theta^C = (\bar{\eta}/\eta_C) \bar{\theta}$, so $\lambda \rightarrow 0$ as $w \downarrow \bar{w}$. Since $\lambda > 0$ on $(\bar{w}, \infty)$ and $\lambda \rightarrow 0$ at both ends, λ attains an interior maximum and is nonmonotone, while by the first part of this step the premium is strictly increasing throughout. $\square$

Proof of Proposition 7. The traditional first-order condition is $\eta_U \Sigma \theta^U = \pi$. The bank's (Lemma 1, with $\kappa > 0$) is $\eta_C \Sigma \theta^C = \pi - \lambda p - (\lambda/\kappa) \Sigma_K \theta^C$. Substituting the former into the latter and multiplying by Σ^{-1} ,

$$\eta_C \theta^C = \eta_U \theta^U - \lambda \tau, \quad \tau = \Sigma^{-1} [p + \Sigma_K \theta^C / \kappa] = \Sigma^{-1} [p + \nabla \mathcal{K}(\theta^C)].$$

Substituting $\theta^U = (\bar{\theta} - m_C \theta^C)/m_U$ from market clearing and collecting terms,

$$\left(\eta_C + \frac{\eta_U m_C}{m_U} \right) \theta^C = \frac{\eta_U}{m_U} \bar{\theta} - \lambda \tau.$$

Since $\eta_C + \eta_U m_C/m_U = (\eta_C m_U + \eta_U m_C)/m_U = \eta_U \eta_C/(\bar{\eta} m_U)$, this gives $\theta^C = (\bar{\eta}/\eta_C) \bar{\theta} - (\bar{\eta} m_U \lambda/(\eta_U \eta_C)) \tau$, and $\bar{\eta} m_U/\eta_U = 1 - \chi$ yields the first equation in (34). Then, from clearing,

$$\theta^U = \frac{\bar{\theta} - m_C \theta^C}{m_U} = \frac{1}{m_U} \left(1 - \frac{m_C \bar{\eta}}{\eta_C} \right) \bar{\theta} + \frac{m_C (1 - \chi) \lambda}{m_U \eta_C} \tau = \frac{\bar{\eta}}{\eta_U} \bar{\theta} + \frac{\chi \lambda}{\eta_U} \tau,$$

using $1 - m_C \bar{\eta}/\eta_C = \bar{\eta} m_U/\eta_U$ and $m_C (1 - \chi)/(\eta_U \eta_C) = m_C \bar{\eta}/(\eta_U \eta_C) = \chi/\eta_U$. The deviations aggregate to zero because $m_C (1 - \chi)/\eta_C = m_U \chi/\eta_U (= m_U m_C \bar{\eta}/(\eta_U \eta_C))$. The decomposition of τ is immediate from $\nabla \mathcal{K}(\theta^C) = M_r \Omega s(\theta^C)/\kappa$, whose Σ^{-1} -image lies in $\Sigma^{-1} \text{col}(M_r)$. $\square$

Proof of Corollary 3. Because a and b are uncorrelated with each other and with every other asset under the true covariance Σ , the a - and b -rows of $\Sigma \theta$ read $(\Sigma \theta)_a = \sigma^2 \theta_a$ and $(\Sigma \theta)_b = \sigma^2 \theta_b$ for any book θ . On the regulatory side, b loads only on an unregulated factor, so its row of M_r vanishes and $(\Sigma_K \theta)_b = 0$; asset a loads on a single regulated factor of risk weight w that is Γ -uncorrelated with every other regulated exposure, so $(\Sigma_K \theta)_a = w^2 \theta_a$. The a - and b -rows of both sectors' first-order conditions therefore decouple from the rest of the book. Writing $\tilde{\pi}_i := \pi_i - \lambda p_i = (1 + \lambda/R_f) \pi_i - (\lambda/R_f) \mu_0$ (using $p_i = (\mu_0 - \pi_i)/R_f$), the traditional condition $\pi = \eta_U \Sigma \theta^U$ and the bank condition (19) give, row by row,

$$\theta_i^U = \frac{\pi_i}{\eta_U \sigma^2}, \quad \theta_a^C = \frac{\tilde{\pi}_a}{\eta_C \sigma^2 (1 + \zeta)}, \quad \theta_b^C = \frac{\tilde{\pi}_b}{\eta_C \sigma^2}, \quad \zeta := \frac{\lambda w^2}{\kappa \eta_C \sigma^2}.$$

Market clearing for each twin is $m_U \theta_i^U + m_C \theta_i^C = \bar{\theta}_0$. Subtracting the two clearing equations and using $\tilde{\pi}_a - \tilde{\pi}_b = (1 + \lambda/R_f)(\pi_a - \pi_b)$,

$$\left[\frac{m_U}{\eta_U} + \frac{m_C}{\eta_C} \frac{1 + \lambda/R_f}{1 + \zeta} \right] \frac{\pi_a - \pi_b}{\sigma^2} = \frac{m_C}{\eta_C \sigma^2} \tilde{\pi}_b \frac{\zeta}{1 + \zeta}.$$

Since $\tilde{\pi}_b = \eta_C \sigma^2 \theta_b^C > 0$ by hypothesis and $\zeta > 0$ (as $\lambda > 0$, $w > 0$), the right side is strictly positive, which gives (36) and $\pi_a > \pi_b$. The inequality $p_a < p_b$ follows from $p_i = (\mu_0 - \pi_i)/R_f$. Then $\theta_a^U > \theta_b^U$ directly, and $\theta_a^C < \theta_b^C$ from clearing. For the comparative statics, write (36) as $F(\zeta) = c_1 \zeta / [c_2(1 + \zeta) + c_3]$ with $c_1 = m_C \sigma^2 \theta_b^C$, $c_2 = m_U/\eta_U$, $c_3 = (m_C/\eta_C)(1 + \lambda/R_f)$. Then $F'(\zeta) = c_1(c_2 + c_3)/[c_2(1 + \zeta) + c_3]^2 > 0$, ζ is strictly increasing in w , and $F(\zeta) \rightarrow c_1/c_2 =$

$\eta_U m_C \sigma^2 \theta_b^C / m_U$ as $\zeta \rightarrow \infty$. □

Proof of Proposition 9. For any θ , completing the square gives

$$V^j(\theta) = \frac{1}{2\eta_j} \pi' \Sigma^{-1} \pi - \frac{\eta_j}{2} \left\| \theta - \frac{1}{\eta_j} \Sigma^{-1} \pi \right\|_{\Sigma}^2.$$

For the traditional investor the second term is zero at the optimum, so $V^U = \frac{1}{2\eta_U} \pi' \Sigma^{-1} \pi$ in either equilibrium. Proposition 2 gives $\pi = \pi^0 + \chi \lambda \Sigma \tau$, and it follows that

$$\pi' \Sigma^{-1} \pi - \pi^{0'} \Sigma^{-1} \pi^0 = 2\chi \lambda \pi^{0'} \tau + (\chi \lambda)^2 \|\tau\|_{\Sigma}^2 = 2\chi \lambda \bar{\eta} T + (\chi \lambda)^2 \|\tau\|_{\Sigma}^2,$$

using $\pi^0 = \bar{\eta} \Sigma \bar{\theta}$ and $\Sigma \tau = p + \nabla \mathcal{K}(\theta^C)$. Dividing by $2\eta_U$ gives (i), and $T \geq 0$ makes both terms nonnegative. For (ii), the first-order condition (18) rearranges to $\frac{1}{\eta_C} \Sigma^{-1} \pi - \theta^C = \frac{\lambda}{\eta_C} \tau$, so the square-completion term equals $\frac{\lambda^2}{2\eta_C} \|\tau\|_{\Sigma}^2$, the compliance cost. The first term changes by the premium difference above scaled by $\frac{1}{2\eta_C}$, the scarcity gain. For (iii), $R_f(p^0 - p)' \bar{\theta} = (\pi - \pi^0)' \bar{\theta} = \chi \lambda \tau' \Sigma \bar{\theta} = \chi \lambda T$. For (iv), form $m_U \Delta C E^U + m_C \Delta C E^C - \chi \lambda T$. The terms linear in T collect to $\chi \lambda T [\bar{\eta} (m_U / \eta_U + m_C / \eta_C) - 1] = 0$ by the definition of $\bar{\eta}$. The quadratic terms collect to

$$\frac{\lambda^2 \|\tau\|_{\Sigma}^2}{2} \left[\chi^2 \left(\frac{m_U}{\eta_U} + \frac{m_C}{\eta_C} \right) - \frac{m_C}{\eta_C} \right] = \frac{\lambda^2 \|\tau\|_{\Sigma}^2}{2} \cdot \frac{\chi^2 - \chi}{\bar{\eta}} = -\frac{\chi(1-\chi)}{2\bar{\eta}} \lambda^2 \|\tau\|_{\Sigma}^2,$$

using $m_C / \eta_C = \chi / \bar{\eta}$. Since $\chi \in (0, 1)$, the sum is $-\text{DWL} \leq 0$. Equivalently, expected aggregate payoffs are invariant to the allocation of the fixed supply, so the change in total surplus equals the change in aggregate variance penalties, and those penalties are minimized by the proportional-sharing allocation $\theta^j = (\bar{\eta} / \eta_j) \bar{\theta}$, which is the frictionless equilibrium allocation. □

Supplementary figures and tables

Figure S.1, referenced in Section 6.3 of the main text, reports the effect of varying bank capital at FRTB parameters. Table S.1, referenced in Section 5.4 of the main text, reports back-of-the-envelope marginal capital intensities and the implied pricing hits. Table S.2, referenced in Section 6.1 of the main text, reports the calibrated parameters of the numerical economy.

References

Viral V. Acharya, Philipp Schnabl, and Gustavo Suarez. Securitization without risk transfer. *Journal of Financial Economics*, 107(3):515–536, 2013.




Figure S.1: The effect of varying bank capital. Panel (a): shadow cost λ and tightness statistic $\chi\lambda$. The requirement is slack beyond $\overline{W} = 111.2$. Panel (b): pricing hits by asset. All identities verified at every grid point.

- Tobias Adrian and Hyun Song Shin. Proccyclical leverage and value-at-risk. *Review of Financial Studies*, 27(2): 373–403, 2014.
- Leif Andersen, Darrell Duffie, and Yang Song. Funding value adjustments. *Journal of Finance*, 74(1):145–192, 2019.
- Alessandro Avellone, Ilaria Foroni, and Chiara Pederzoli. Minimum capital requirement portfolios according to the new basel framework for market risk. *Financial Markets and Portfolio Management*, 39(2):171–192, 2025.
- Malcolm Baker and Jeffrey Wurgler. Do strict capital requirements raise the cost of capital? Bank regulation, capital structure, and the low-risk anomaly. *American Economic Review*, 105(5):315–320, 2015.
- Suleyman Basak and Benjamin Croitoru. Equilibrium mispricing in a capital market with portfolio constraints. *Review of Financial Studies*, 13(3):715–748, 2000.
- Suleyman Basak and Anna Pavlova. Asset prices and institutional investors. *American Economic Review*, 103(5): 1728–1758, 2013.
- Suleyman Basak and Alexander Shapiro. Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies*, 14(2):371–405, 2001.
- Basel Committee on Banking Supervision. An assessment of the long-term economic impact of stronger capital and liquidity requirements. Technical report, Bank for International Settlements, August 2010.
- Basel Committee on Banking Supervision. Minimum capital requirements for market risk. Technical report, Bank for International Settlements, January 2019. Standardized approach consolidated in the Basel Framework as MAR20–MAR23.
- Basel Committee on Banking Supervision. Basel III monitoring report. Technical report, Bank for International Settlements, 2025.
- Basel Committee on Banking Supervision. Implementation of Basel standards: Jurisdictional implementation dashboard. Technical report, Bank for International Settlements, 2026.
- Bo Becker and Efraim Benmelech. The resilience of the U.S. corporate bond market during financial crises. *International Journal of Central Banking*, 21(4):405–463, 2025.
- Bo Becker and Victoria Ivashina. Reaching for yield in the bond market. *Journal of Finance*, 70(5):1863–1902, 2015.
- Juliane Begenau and Tim Landoigt. Financial regulation in a quantitative model of the modern banking system. *Review of Economic Studies*, 89(4):1748–1784, 2022.
- Markus Behn, Rainer Haselmann, and Vikrant Vig. The limits of model-based regulation. *Journal of Finance*, 77(3): 1635–1684, 2022.
- Fischer Black. Capital market equilibrium with restricted borrowing. *Journal of Business*, 45(3):444–455, 1972.
- Martijn A. Boermans and Bram van der Kroft. Capital regulation induced reaching for systematic yield: Financial instability through fire sales. *Journal of Banking & Finance*, 158:107030, 2024.
- Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial*

Table S.1: Marginal capital intensity and implied pricing hit (annual expected excess return relative to a zero-charge twin), by FRTB delta treatment. *Notes:* Risk weights as encoded in the project’s FRTB parameter file. GIRR tenor weights are 1.7% (0.25y) declining to 1.1% (5y–30y), from MAR21.42, with the MAR21.44 $\sqrt{2}$ divisor for specified currencies. Equity bucket weights follow MAR21.77, and FX is 15% (MAR21.86) with the $\sqrt{2}$ divisor for liquid pairs (MAR21.87). CSR non-securitisation parameters follow MAR21.53–21.54: the corporate row uses bucket 4 (investment-grade industrials, risk weight 3.0%) with spread duration 4.5, and the bond–CDS row uses the 99.90% basis correlation between different curves of the same issuer. Durations: 1.9 (2y), 4.5 (5y), 8.5 (10y), 17.5 (30y). Commodity rows are omitted. All figures are medium-scenario, delta-only, and ignore the MAR21.6 scenario-max, which adds a further loading that depends on book composition. They are therefore conservative.

Exposure (per \$1 market value)	FRTB treatment	Capital intensity	Pricing hit (bp) at $\chi\lambda =$		
			1%	3%	5%
2y government note, liquid ccy	GIRR 2y, RW $1.3\%/\sqrt{2}$	1.7%	2	5	9
10y government bond, liquid ccy	GIRR 10y, RW $1.1\%/\sqrt{2}$	6.6%	7	20	33
30y government bond, liquid ccy	GIRR 30y, RW $1.1\%/\sqrt{2}$	13.6%	14	41	68
5y IG corporate bond	GIRR 5y + CSR b4, RW 3.0%	17.0%	17	51	85
FX position, liquid pair	FX, RW $15\%/\sqrt{2}$	10.6%	11	32	53
Advanced-economy equity index	Equity b12, RW 15%	15.0%	15	45	75
Large-cap advanced equity	Equity b5, RW 30%	30.0%	30	90	150
Large-cap emerging equity	Equity b1, RW 55%	55.0%	55	165	275
Cross-curve 10y basis, same ccy	GIRR, $\rho = 99.90\%$	0.30%	0.3	0.9	1.5
Bond–CDS basis, 5y IG corporate	CSR, $\rho_{\text{basis}} = 99.90\%$	0.60%	0.6	1.8	3.0
On-/off-the-run basis	Same risk factor	0	0	0	0

Table S.2: Calibrated economy of Section 6 of the main text. Factor correlations: $\text{corr}(f_2, f_3) = 0.85$, $\text{corr}(f_2, f_4) = 0.50$, $\text{corr}(f_3, f_4) = 0.45$. Prescribed Γ : $\gamma_{23} = 0.45$ (equity cross-bucket), cross-class 0. $\eta_U = \eta_C = 0.05$, $m_C = 1/3$, $R_f = 1.04$, $\bar{W} = 111.2$, baseline $W_0^C = 60$.

Asset	Factor	Sys. vol	Idio. vol	Supply	FRTB delta treatment	$\mathcal{K}$ per \$1
10y govt bond	f_1 (GIRR 10y)	5.9%	0.5%	60	RW $1.1\%/\sqrt{2}$, $D = 8.5$	6.61%
Equity index	f_2 (equity b12)	16%	1%	30	RW 15%	15%
Large-cap equity	f_3 (equity b5)	19%	6%	8	RW 30%	30%
Basis package	f_4 (off-list)	5.5%	1%	6	none	0

Studies, 22(6):2201–2238, 2009.

Francisco Covas and Shigeru Fujita. Procyclicality of capital requirements in a general equilibrium model of liquidity dependence. *International Journal of Central Banking*, 6(4):137–173, 2010.

Domenico Cuoco. Optimal consumption and equilibrium prices with portfolio constraints and stochastic income. *Journal of Economic Theory*, 72(1):33–73, 1997.

Domenico Cuoco and Hong Liu. An analysis of VaR-based capital requirements. *Journal of Financial Intermediation*, 15(3):362–394, 2006.

Jon Danielsson, Hyun Song Shin, and Jean-Pierre Zigrand. The impact of risk regulation on price dynamics. *Journal of Banking & Finance*, 28(5):1069–1087, 2004.

Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957, 2018.

Simon Dubecq, Benoît Mojon, and Xavier Ragot. Risk shifting with fuzzy capital constraints. *International Journal of Central Banking*, 11(1):71–101, 2015.

- Darrell Duffie. Financial regulatory reform after the crisis: An assessment. *Management Science*, 64(10):4835–4857, 2018.
- Andrew Ellul, Chotibhak Jotikasthira, and Christian T. Lundblad. Regulatory pressure and fire sales in the corporate bond market. *Journal of Financial Economics*, 101(3):596–620, 2011.
- European Commission. Commission delegated regulation amending the own funds requirements for market risk: Targeted multiplier and operational relief measures. Technical report, European Commission, June 2026. Adopted 4 June 2026; applicable 2027–2030.
- Simon Firestone, Amy Lorenc, and Benjamin Ranish. An empirical economic assessment of the costs and benefits of bank capital in the United States. *Federal Reserve Bank of St. Louis Review*, 101(3):203–230, 2019.
- Matthias Fleckenstein and Francis A. Longstaff. Renting balance sheet space: Intermediary balance sheet rental costs and the valuation of derivatives. *Review of Financial Studies*, 33(11):5051–5091, 2020.
- Andrea Frazzini and Lasse Heje Pedersen. Betting against beta. *Journal of Financial Economics*, 111(1):1–25, 2014.
- Leonardo Gambacorta and Sudipto Karmakar. Leverage and risk-weighted capital requirements. *International Journal of Central Banking*, 14(5):153–191, 2018.
- Nicolae Gârleanu and Lasse Heje Pedersen. Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022, 2011.
- Carlos Alexander Grajales and Santiago Medina Hurtado. Sensitivities-based method and expected shortfall for market risk under FRTB: Its impact on options risk capital. *Journal of Economics, Finance and Administrative Science*, 28(55), 2023.
- Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2–3):361–407, 2002.
- Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–770, 2013.
- Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017.
- Giuliano Iannotta, George Pennacchi, and João A. C. Santos. Ratings-based regulation and systematic risk incentives. *Review of Financial Studies*, 32(4):1374–1415, 2019.
- 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.
- Urban J. Jermann. Negative swap spreads and limited arbitrage. *Review of Financial Studies*, 33(1):212–238, 2020.
- Anil K. Kashyap, Natalia Kovrijnykh, Jian Li, and Anna Pavlova. The benchmark inclusion subsidy. *Journal of Financial Economics*, 142(2):756–774, 2021.
- Roni Kisin and Asaf Manela. The shadow cost of bank capital requirements. *Review of Financial Studies*, 29(7):1780–1820, 2016.
- Sven Klingler and Suresh Sundaresan. An explanation of negative swap spreads: Demand for duration from underfunded pension plans. *Journal of Finance*, 74(2):675–710, 2019.
- Ralph S. J. Koijen and Motohiro Yogo. The cost of financial frictions for life insurers. *American Economic Review*, 105(1):445–475, 2015.
- Ralph S. J. Koijen and Motohiro Yogo. The fragility of market risk insurance. *Journal of Finance*, 77(2):815–862, 2022.
- Christian C. Opp, Marcus M. Opp, and Milton Harris. Rating agencies in the face of regulation. *Journal of Financial Economics*, 108(1):46–61, 2013.
- Guillaume Plantin. Shadow banking and bank capital regulation. *Review of Financial Studies*, 28(1):146–175, 2015.
- Matthew C. Plosser and João A. C. Santos. Banks’ incentives and inconsistent risk models. *Review of Financial Studies*, 31(6):2080–2112, 2018.
- Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *Journal of Finance*, 52(1):35–55, 1997.