

7514638-recursive-portfolio-learning-and-the-endogenous-riskless-rate

SUMMARY: This paper formulates portfolio choice as a recursive learning operator on conditional covariance operators, selecting the minimum-risk direction, re-estimating its return law under long-memory, heavy-tailed dynamics, and re-injecting it. The map admits a fixed point whose expected return defines the Adaptive Minimum-Risk Rate, an endogenous, portfolio-specific shadow benchmark. The authors demonstrate that the recursion converges to a fixed-point invariant to the informational regime, lying orders of magnitude below the raw minimum-variance level. This fixed point is characterized by the first-order conditions, which show that every held asset has a beta of exactly one with respect to the fixed-point portfolio.

METHOD: The authors formulate portfolio choice as a recursive learning operator on conditional covariance operators. They select the minimum-risk allocation, estimate the return law, including its memory and tail structure, and re-inject the resulting synthetic instrument into the investable universe. The recursion is governed by the covariance operator's condition number, and it converges to a fixed point invariant to the informational regime.

KEY FINDINGS:

- The Adaptive Minimum-Risk Rate is an endogenous, portfolio-specific shadow benchmark.
- The recursion converges to a fixed-point invariant to the informational regime.
- The fixed-point portfolio has every held asset with a beta of exactly one.
- The fixed-point portfolio lies orders of magnitude below the raw minimum-variance level.

IMPLICATION FOR TRADING: The paper's findings suggest that practitioners should use the Adaptive Minimum-Risk Rate as a benchmark for their portfolios. This benchmark is endogenous and specific to each portfolio, allowing for a more accurate assessment of risk. The fixed-point portfolio's structure, with every held asset having a beta of exactly one, can help traders understand the risk exposure of their portfolios. Additionally, the convergence to a fixed-point invariant to the informational regime indicates that the benchmark is robust and reliable, even in complex and changing market conditions. Practitioners can use this information to make more informed decisions and manage risk effectively.

Recursive Portfolio Learning and the Endogenous Riskless Rate

Ayush Jha^{1,*} Abootaleb Shirvani² Ali Jaffri³
Svetlozar T. Rachev⁴ Frank J. Fabozzi⁵

Abstract

Portfolio optimization almost universally treats the riskless rate as exogenous, typically a short Treasury yield. In incomplete markets, this is an artifact as the rate at which an investor can discharge risk depends on the hedging technology inside her own portfolio, not on a security she may not trade at scale. We formulate portfolio choice as a recursive learning operator on conditional covariance operators by selecting a minimum-risk direction, re-estimating its return law under long-memory, heavy-tailed dynamics, re-injecting it, repeat. The map admits a fixed point whose expected return defines the Adaptive Minimum-Risk Rate, an endogenous, portfolio-specific shadow benchmark. Contraction is governed by the covariance operator’s condition number, destabilizing precisely when assets become near-collinear. Empirically, the recursion converges on equity and cryptocurrency universes to a fixed-point invariant to the informational regime, lying orders of magnitude below the raw minimum-variance level—separating reducible risk from the structural residual no hedge removes.

Keywords: Recursive portfolio optimization; Endogenous riskless rate; Incomplete markets; Fixed-point learning; Long-range dependence.

JEL Classification: C61, C63, G11, G12.

¹Department of Mathematics and Statistics, Texas Tech University. 2500 Broadway, Lubbock, TX, USA – 79409. *Corresponding author: Ayush Jha; ayush.jha@ttu.edu

²Department of Mathematical Sciences, Kean University. 1000 Morris Avenue, Union, NJ, USA – 07083. ashirvani@kean.edu

³College of Business, North Dakota State University. Barry Hall, 811 2nd Ave N, Fargo, ND, USA – 58102. ali.jaffri@ndsu.edu

⁴Department of Mathematics and Statistics, Texas Tech University. 1108 Memorial Circle, Lubbock, TX, USA – 79409. zari.rachev@ttu.edu

⁵Carey Business School, Johns Hopkins University. 100 International Drive, Baltimore, MD, USA – 21202. ffabozzi1@jhu.edu

1 Introduction

Every applied portfolio problem begins by fixing a riskless rate. The choice is treated as innocuous: a three-month Treasury yield is inserted, excess returns are computed, and the optimization proceeds. The convenience obscures a strong economic assumption, namely that there exists a single security whose return is riskless *for every investor simultaneously*, and that access to it is unconstrained. Neither is true. Money-market access is intermediated and rationed; repo haircuts differ across counterparties; collateral quality is state-dependent; and for large classes of investors, including insurers matching long-dated liabilities, funds facing redemption risk, and participants in crypto-asset or emerging markets, the relevant safe asset is either unavailable or ceases to behave as one precisely when it is needed. [Lintner \(1969\)](#), [Fama \(1971\)](#) and [Mayers \(1973\)](#) established long ago that a universal borrowing and lending rate does not exist in practice, yet the exogenous-rate convention has survived essentially intact.

The economic cost of the convention is not that the number is slightly wrong. It is that the number is *unresponsive*. A Treasury yield is an equilibrium object shaped by monetary policy, convenience yields, and institutional demand, and it is deliberately smoothed. The rate at which an investor can actually shed risk is a property of the dependence structure of the assets she holds, and that structure moves violently. When correlations converge toward one, the hedging technology inside a portfolio degrades, the true cost of remaining invested rises, and the exogenous benchmark does not move at all. Risk systems anchored to it therefore report the largest gap between perceived and actual safety at exactly the moment the gap matters. This is the mechanism through which apparently well-hedged balance sheets deleverage simultaneously, and it is why an endogenous, portfolio-specific benchmark is a risk-management object and not only a theoretical one.

Contribution. This paper makes the following contribution: it formulates portfolio choice as a recursive learning operator whose fixed point is a portfolio-specific shadow riskless rate, and it characterizes that fixed point economically. The device is a single nonlinear map $\mathcal{T}$ on the cone of conditional covariance operators. Given a current dependence structure, $\mathcal{T}$ (i) selects the minimum-risk allocation, (ii) estimates the return law of that allocation, including its memory and tail structure, (iii) re-injects the resulting synthetic instrument into the investable universe, and (iv) updates the covariance operator. Iterating $\mathcal{T}$ is a learning process: the optimizer is not told what the regime is, it discovers the second-order and higher-order structure of its own risk-minimal direction and feeds that discovery back into the next allocation. The fixed point is the configuration at which this feedback stops producing information.

Everything else in the paper exists to support that object. The long-memory and heavy-tailed specification describes the law of motion that makes the re-injection step informative. The extension to coherent risk measures is not a separate framework; it serves as a verification that the fixed-point argument does not depend on the quadratic objective. The numerical section measures the two economic quantities the operator delivers, the speed at which the hedge is learned and the level of risk that cannot be hedged away.

What the fixed point means. The central analytical result is not convergence, which is nearly immediate, but the characterization of the limit. Theorem 1 shows that at any fixed point the first-order conditions impose

$$\text{Cov}(R_i, r^*) = \sigma_*^2 \quad \text{for every asset held,} \quad \text{Cov}(R_i, r^*) \geq \sigma_*^2 \quad \text{otherwise,}$$

where r^* is the fixed-point portfolio return and σ_*^2 its variance. Every held asset has a beta of exactly one with respect to the fixed-point portfolio. Equivalently, the marginal hedging value of every asset the investor is willing to hold has been equalized, and no asset outside the holdings offers a better one. The multiplier on the budget constraint equals $2\sigma_*^2$: the shadow risk cost of committing an additional unit of wealth to the market. The expected return earned at that point, $\text{AMRR} = \mu^\top w^*$, is therefore not an arbitrary summary statistic of a minimum-variance portfolio. It is the return to the position at which hedging demand is exactly exhausted from within the market, which is precisely the economic content one wants from a riskless rate when no riskless security exists. This furnishes a constructive, learning-based counterpart to the zero-beta rate of Black (1972), obtained without imposing market clearing.

Nature of recursion. A natural objection is that the exercise is circular. A portfolio formed from existing assets lies in their span, so re-injecting it should change nothing. Under a fixed scenario set and an exact covariance estimator this objection is correct, and we state it explicitly in Remark 1 rather than concealing it. The recursion has content only because the re-injection step is informational, not algebraic: the synthetic instrument enters carrying its own estimated dynamics, so the object fed back is the fitted law of the risk-minimal direction rather than a linear combination of historical observations. This distinction is what separates the framework from adaptive and online portfolio methods, in which the estimator is updated but the investable set is inert, and it is what makes the fixed point a statement about learned structure rather than about arithmetic.

Failure of contraction in crisis. We do not assume that $\mathcal{T}$ contracts. We derive its Lipschitz constant from primitives. Lemma 1 shows the minimum-variance selection map is Lipschitz with constant $1/\underline{\lambda}$, where $\underline{\lambda}$ is the smallest eigenvalue of the admissible covariance operators, and Theorem 2 propagates this into a bound governed by the condition number of the covariance operator. The economic reading is direct and, we think, the most useful implication of the formalism: the learning dynamics are stable when assets are genuinely distinct and become unstable as they collapse toward collinearity. Regularization is therefore not a statistical convenience in this framework but the object that keeps the investor’s inference about her own hedge well-posed. Ill-conditioning and hedging instability are the same phenomenon viewed from two sides.

Section 5 reports two findings that are economic rather than computational. First, the recursion converges on both universes and the fixed point is invariant to the informational regime: the backward- and forward-looking implementations reach irreducible variances that agree to within two-tenths of one percent, differing only in the number of iterations they require. The irreducible minimum-risk rate is thus a property of the market and

the admissible set rather than of the forecasting model, which is the strongest form the paper’s central claim can take. Second, the fixed point lies far below the first-iterate minimum-variance level, by more than four orders of magnitude for equities and two for cryptocurrencies, so most of the apparent risk in the raw covariance is hedgeable and a structural core is not; a risk framework that reads its exposure off the raw covariance therefore misses the reducible portion entirely and mis-sizes the residual. A risk framework calibrated on the raw covariance believes it faces risk that a system aware of long memory and heavy tails knows to be structural. Finally, the endogenous benchmark exhibits materially lower volatility persistence than Treasury yields, meaning it reprices risk on a timescale that policy-anchored benchmarks cannot match.

Scope. We are deliberate about what is not claimed. The framework is a portfolio-level construction, not a general equilibrium model; the AMRR is a shadow price implied by one investor’s problem and its aggregation into a market-clearing rate is left open. The numerical evidence is diagnostic, covering two asset universes over a limited sample, and is intended to illustrate the operator’s behavior rather than to establish out-of-sample performance. We view the macro-financial implications discussed in Section 5.6 as motivation for the construction, not as results delivered by it.

The remainder of the paper proceeds as follows. Section 2 positions the contribution. Section 3 develops the recursive learning operator. Section 4 establishes convergence and characterizes the fixed point as a shadow riskless rate. Section 5 reports numerical evidence. Section 6 concludes. Proofs, the state transition law, and the extension to convex risk measures are in the appendices.

2 Positioning

Three literatures bear on the construction, and it is useful to state precisely what is borrowed from each.

Benchmarks without a riskless asset. Black (1972) showed that equilibrium pricing survives the absence of a traded riskless asset, with the zero-beta portfolio’s return replacing it. Hansen and Jagannathan (1991) showed that admissible stochastic discount factors are bounded by return moments alone, so that benchmark information is recoverable from asset data. Both results are equilibrium characterizations: they describe what the benchmark must satisfy. Our object is constructive. We do not impose market clearing or no-arbitrage restrictions and then solve for a rate consistent with them; we specify an algorithm an investor can run on her own return data and show that its fixed point satisfies an equal-marginal-hedging condition with the same economic content. The relationship to Black (1972) is therefore one of correspondence rather than nesting, and Theorem 1 makes the correspondence explicit. Rachev et al. (2017) likewise argue for shadow benchmarks in distorted markets; we supply a mechanism that generates one.

Adaptive and recursive portfolio choice. Classical mean–variance analysis treats the covariance matrix as a fixed input (Sharpe, 1964; Lintner, 1965; Mossin, 1966). A large subsequent literature updates that input, through shrinkage, factor structure, or online learning. What these approaches share is that *the estimator adapts while the investable set does not*. The map from data to allocation is refined, but the set of attainable positions is taken as given. The recursion studied here inverts this: the allocation rule modifies the investable set, which modifies the estimated dependence structure, which modifies the allocation rule. Portfolio choice becomes a controlled dynamical system in which the control moves the state. Factor-extraction and principal-component methods (Adrian et al., 2013; Lauria et al., 2024) similarly reduce dimension using a fixed criterion; the synthetic instrument here is generated endogenously by the optimization rather than extracted by a pre-specified statistical rule, and it carries a fitted return law rather than a loading.

Dependence structure as the object being learned. The framework is agnostic about the return law but not indifferent to it. Long memory in volatility and heavy-tailed innovations are the empirical features that make the re-injection step informative, because they are the features a covariance estimator computed on realized history systematically fails to encode. The specification we use, ARFIMA–FIGARCH dynamics with normal inverse Gaussian innovations, is standard and is described in Appendix B; nothing in Sections 3–4 depends on it. Similarly, the argument does not depend on variance as the risk criterion. Appendix C verifies that any convex, continuous risk functional, including Conditional Value-at-Risk (Rockafellar and Uryasev, 2000), preserves the fixed-point structure. We record these as robustness properties of the single contribution rather than as extensions of it.

3 The Recursive Learning Operator

3.1 Environment

Time is discrete, indexed by $t \in \mathbb{N}$, on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$. Returns are square-integrable and therefore lie in the Hilbert space $\mathcal{H} = L^2(\Omega, \mathcal{F}, \mathbb{P})$ with inner product $\langle X, Y \rangle = \mathbb{E}[XY]$. Working in $\mathcal{H}$ rather than in $\mathbb{R}^N$ is not decorative: the recursion enlarges the finite-dimensional representation of the asset universe at every step, and $\mathcal{H}$ provides the fixed ambient space in which the dynamics are well defined.

At iteration k the universe consists of N_k instruments with return vector

$$R_k = (R_{1k}, \dots, R_{N_k k})^\top \in \mathcal{H}^{N_k},$$

conditional mean $\mu_k = \mathbb{E}[R_k \mid \mathcal{F}_{t-1}]$, and conditional covariance operator

$$\Sigma_k = \mathbb{E}[(R_k - \mu_k)(R_k - \mu_k)^\top \mid \mathcal{F}_{t-1}] \in \mathcal{S}_+^{N_k}. \quad (1)$$

The admissible set is the unit simplex,

$$\mathcal{W}_k = \{w \in \mathbb{R}^{N_k} : \mathbf{1}^\top w = 1, w \geq 0\}, \quad (2)$$

convex and compact. The long-only restriction is substantive rather than cosmetic. It is the constraint under which an unconstrained investor's problem would be trivial, and it is the empirically relevant one for the institutions whose hedging demand motivates the exercise. A portfolio w delivers return $r_k(w) = w^\top R_k \in \mathcal{H}$ with conditional variance $w^\top \Sigma_k w$.

3.2 The operator

The learning map is the composition of three primitives.

(i) Selection. $\mathcal{M} : \mathcal{S}_{++}^{N_k} \rightarrow \mathcal{W}_k$ returns the minimum-risk allocation,

$$\mathcal{M}(\Sigma_k) := \arg \min_{w \in \mathcal{W}_k} w^\top \Sigma_k w, \quad (3)$$

which exists and is unique by strict convexity of the objective on a compact convex set. Write $w_{k+1} := \mathcal{M}(\Sigma_k)$ and $\sigma_k^2 := w_{k+1}^\top \Sigma_k w_{k+1}$.

(ii) Augmentation. The selected direction is instantiated as a new instrument. Let Λ_k denote the conditional law of $r_k(w_{k+1})$ estimated from the current information set, and let $\tilde{r}_{k+1} \sim \Lambda_k$ denote a realization of that law over the evaluation horizon. The augmentation operator is

$$\mathcal{A}(R_k, w_{k+1}) := R_{k+1} = [R_k, \tilde{r}_{k+1}], \quad N_{k+1} = N_k + 1. \quad (4)$$

(iii) Update. $\mathcal{S}$ maps a return panel to its conditional covariance operator, $\Sigma_{k+1} = \mathcal{S}(R_{k+1})$. The baseline is a regularized sample estimator; any estimator satisfying Assumption 1 below is admissible.

Composing,

$$\Sigma_{k+1} = \mathcal{T}(\Sigma_k) := \mathcal{S}(\mathcal{A}(R_k, \mathcal{M}(\Sigma_k))). \quad (5)$$

The map is nonlinear in two distinct ways: the selection rule depends on the state through an argmin, and the state space itself is re-indexed at each step. Equation (5) is the paper's single object. Sections 4–5 study it.

3.3 Informational nature of the recursion

The content of (5) rests entirely on step (ii), and it is worth isolating the issue rather than leaving it implicit.

Remark 1 (The degenerate case). Suppose the augmentation appends the realized in-sample combination $R_k w_{k+1}$ and $\mathcal{S}$ is the exact covariance of the augmented panel. Then $r_{k+1} \in \text{conv}\{R_{1k}, \dots, R_{N_k k}\}$, so the set of attainable portfolio returns satisfies

$$\mathcal{C}_{k+1} := \{R_{k+1}^\top w : w \in \mathcal{W}_{k+1}\} = \mathcal{C}_k,$$

and consequently $\sigma_{k+1}^2 = \sigma_k^2$ for all k . The recursion is stationary after one step and carries no information.

Remark 1 is not a limitation to be worked around but identifies where the economic intuition matters. The synthetic instrument must be introduced as a *law* rather than as a realized linear combination. Under (4), $\tilde{r}_{k+1}$ is drawn from the fitted conditional law Λ_k of the risk-minimal direction, which encodes persistence and tail behavior that the historical panel does not reproduce. The consequence is that $\tilde{r}_{k+1} \notin \text{span}\{R_{1k}, \dots, R_{N_k k}\}$ in general, the attainable set genuinely expands, and each iteration transmits information about the dependence structure of the current best hedge into the next optimization.

This is the sense in which the framework *learns*. The investor is not supplied with a regime label or a state variable. She forms a hedge, studies how that hedge behaves, discovers that it has memory and tail properties her covariance estimate did not contain, and revises. The fixed point is reached when this cycle stops producing new information: the risk-minimal direction has dynamics that are already consistent with the covariance operator that generated it. That self-consistency, not numerical convergence, is what makes the limit interpretable as a benchmark.

Assumption 1 (Regularity). There exist constants $0 < \underline{\lambda} \leq \bar{\lambda} < \infty$ such that every operator generated by the recursion lies in

$$\mathcal{K} := \{\Sigma : \underline{\lambda}I \preceq \Sigma \preceq \bar{\lambda}I\},$$

and $\mathcal{T}(\mathcal{K}) \subseteq \mathcal{K}$. Moreover the covariance update $\mathcal{S} \circ \mathcal{A}$ is Lipschitz in its weight argument with modulus $L_{\mathcal{S}}$, uniformly on $\mathcal{K}$.

Assumption 1 is enforced in implementation by regularization of $\mathcal{S}$, which bounds $\underline{\lambda}$ away from zero. Section 4.3 shows that this is not a technical device: $\underline{\lambda}$ is the quantity that governs whether the investor's inference about her own hedge is stable.

Assumption 2 (Nested attainability). Augmentation preserves previously attainable positions: for every $w \in \mathcal{W}_k$ there is $w' \in \mathcal{W}_{k+1}$ with $r_{k+1}(w') = r_k(w)$ almost surely.

Assumption 2 holds by construction under (4), since $\mathcal{W}_{k+1}$ contains the embedding of $\mathcal{W}_k$ that places zero weight on the new instrument. It is stated separately because it, and not the specific form of $\mathcal{S}$, is what delivers monotonicity.

4 Fixed Points and the Endogenous Riskless Rate

4.1 Convergence

Define the induced risk functional $V : \mathcal{S}_+^N \rightarrow \mathbb{R}_+$ by

$$V(\Sigma) := \min_{w \in \mathcal{W}} w^\top \Sigma w, \tag{6}$$

so that $\sigma_k^2 = V(\Sigma_k)$. Because the recursion expands the representation dimension, convergence cannot be phrased as equality of operators across iterations; it is phrased in

terms of the scalar V , which is defined on every iterate and is the economically meaningful quantity.

Proposition 1 (Monotone stabilization). *Under Assumption 2, the sequence $\{\sigma_k^2\}$ is non-increasing and bounded below by zero, hence*

$$\sigma_k^2 \downarrow \sigma_*^2 := \lim_{k \rightarrow \infty} \sigma_k^2 \geq 0, \quad |\sigma_{k+1}^2 - \sigma_k^2| \rightarrow 0.$$

If in addition Assumption 1 holds, then $\sigma_^2 \geq \underline{\lambda}/N_\infty$ whenever $N_k \uparrow N_\infty < \infty$, and $\sigma_*^2 > 0$.*

The lower bound matters for interpretation. It states that the recursion cannot manufacture a riskless position out of risky ones: as long as the covariance operator remains uniformly positive definite, a strictly positive floor of risk survives every round of learning. That floor, σ_*^2 , is the quantitative price of market incompleteness for this investor. Proposition 1 is proved in Appendix A; the argument is elementary, which is the point. Convergence is cheap here, and we treat it as a precondition for the analysis rather than as a finding.

Remark 2. We do not claim $\|\Sigma_{k+1} - \Sigma_k\| \rightarrow 0$ as a consequence of Proposition 1. Convergence of the value functional does not by itself imply stabilization of the operator, and the two are established under different conditions: the former requires only nested attainability, the latter requires the contraction argument of Section 4.3. Conflating them would overstate what monotonicity delivers.

4.2 The fixed point as a shadow riskless rate

The substantive result concerns what the limit *is*.

Theorem 1 (Equalized hedging demand and the shadow rate). *Let Σ_* be a fixed point of $\mathcal{T}$ in $\mathcal{K}$, let $w^* = \mathcal{M}(\Sigma_*)$, let $r^* = w^{*\top} R_*$ be the fixed-point portfolio return, and let $\sigma_*^2 = \text{Var}(r^* \mid \mathcal{F}_{t-1})$. Then:*

(1) (**Shadow price.**) *The multiplier γ on the budget constraint $\mathbf{1}^\top w = 1$ satisfies $\gamma = 2\sigma_*^2$.*

(2) (**Equalized marginal hedging value.**) *For every asset i ,*

$$\text{Cov}(R_i, r^* \mid \mathcal{F}_{t-1}) = \sigma_*^2 + \frac{1}{2}\nu_i, \quad \nu_i \geq 0, \quad \nu_i w_i^* = 0,$$

so that $\text{Cov}(R_i, r^) = \sigma_*^2$ for every held asset and $\text{Cov}(R_i, r^*) \geq \sigma_*^2$ for every excluded asset. Equivalently, $\beta_{i,r^*} = 1$ on the support of w^* and $\beta_{i,r^*} \geq 1$ off it.*

(3) (**No residual hedging demand.**) *For every $w \in \mathcal{W}$, $\text{Cov}(r(w), r^*) \geq \sigma_*^2$, with equality if $\text{supp}(w) \subseteq \text{supp}(w^*)$. The converse holds under strict complementarity, that is, when $\nu_i > 0$ for every i outside the support of w^* . Absent that condition the equality set is $\{i : \nu_i = 0\}$, which contains the support of w^* and may be strictly larger, so other portfolios can attain the same covariance contribution.*

Part (2) is the economic core. At the fixed point, every asset the investor is willing to hold contributes exactly the same marginal quantity of risk to the portfolio, and every asset she excludes would contribute more. Hedging demand has been equalized across the entire held universe: there is no reallocation, however small, that improves the hedge. Part (3) says the same thing from the outside, namely that no attainable position covaries less with the fixed-point portfolio than the fixed-point portfolio does with itself. The system has exhausted the hedging technology available in the market.

Definition 1 (Adaptive Minimum-Risk Rate). The Adaptive Minimum-Risk Rate is the conditional expected return of the fixed-point allocation,

$$\text{AMRR}_t := \mathbb{E}[w_t^{*\top} R_t \mid \mathcal{F}_{t-1}] = w_t^{*\top} \mu_t. \quad (7)$$

Definition 1 is stated after Theorem 1 deliberately. Absent the theorem, the AMRR is merely the mean return of a minimum-variance portfolio, of which the literature contains many. With the theorem, it is the return earned at the unique position where marginal hedging value is equalized and the shadow risk price of committed wealth equals σ_*^2 . That is the defining property of a shadow riskless rate in a market that does not contain a riskless security: not that the position bears no risk—Proposition 1 rules that out, since the fixed-point portfolio retains strictly positive irreducible variance—but that no further risk can be transferred at any price available in the market. We therefore use ‘shadow riskless rate’ and ‘endogenous minimum-risk rate’ interchangeably, and reserve ‘riskless rate’ without qualification for the exogenous benchmark the construction replaces. The AMRR is the compensation the investor receives for bearing the residue.

Corollary 1 (Relation to the zero-beta rate). *Every asset held at the fixed point has unit beta with respect to r^* , so r^* carries no cross-sectional risk premium relative to the held universe since differences in μ_i across $i \in \text{supp}(w^*)$ cannot be attributed to differences in exposure to r^* . The AMRR therefore occupies the role played by the zero-beta return in Black (1972), but is obtained by iterating an investor-level algorithm rather than by imposing market clearing.*

Corollary 1 also delimits the claim. Because the construction is investor-specific, two investors with different admissible sets, asset universes, or estimated laws will obtain different values of AMRR. This is a feature. This observation motivated the paper as the rate at which risk can be discharged is not a universal constant applicable to portfolios across all asset classes, and a framework that produces a single object takes away the phenomenon of interest.

4.3 Stability, conditioning, and why hedges fail together

We now bound the Lipschitz behavior of $\mathcal{T}$. The constants are derived by assuming a modulus below unity and concluding contraction would restate the Banach fixed-point theorem without economic content.

Lemma 1 (Lipschitz selection). *Let $\Sigma, \tilde{\Sigma} \in \mathcal{K}$ and $w = \mathcal{M}(\Sigma)$, $\tilde{w} = \mathcal{M}(\tilde{\Sigma})$. Then*

$$\|w - \tilde{w}\|_2 \leq \frac{1}{\underline{\lambda}} \|\Sigma - \tilde{\Sigma}\|_2.$$

Lemma 2 (Lipschitz value). *For all $\Sigma, \tilde{\Sigma} \in \mathcal{S}_+^N$, $|V(\Sigma) - V(\tilde{\Sigma})| \leq \|\Sigma - \tilde{\Sigma}\|_2$. The induced risk functional is thus 1-Lipschitz in spectral norm, uniformly in dimension.*

Theorem 2 (Contraction is governed by conditioning). *Under Assumption 1, $\mathcal{T}$ is Lipschitz on $\mathcal{K}$ with*

$$\begin{aligned} \|\mathcal{T}(\Sigma) - \mathcal{T}(\tilde{\Sigma})\|_F &\leq L_{\mathcal{T}} \|\Sigma - \tilde{\Sigma}\|_2, \\ L_{\mathcal{T}} &\leq \frac{L_{\mathcal{S}}}{\underline{\lambda}} \leq c_N \kappa(\Sigma), \quad \kappa(\Sigma) := \frac{\bar{\lambda}}{\underline{\lambda}}, \end{aligned}$$

where c_N depends only on the dimension and the second-moment bound of returns. If $L_{\mathcal{T}} < 1$, then $\mathcal{T}$ admits a unique fixed point $\Sigma_* \in \mathcal{K}$, the iterates satisfy $\|\Sigma_k - \Sigma_*\|_F \leq L_{\mathcal{T}}^{k-1} \|\Sigma_1 - \Sigma_*\|_F$, and by Lemma 1 the allocations satisfy $\|w_k - w^*\|_2 \leq \underline{\lambda}^{-1} \|\Sigma_k - \Sigma_*\|_2 \rightarrow 0$ geometrically.

The bound is sufficient and deliberately conservative; Section 5 documents convergence in configurations where it is not verified numerically. Its value is interpretive. The only quantity in $L_{\mathcal{T}}$ that can diverge is $\underline{\lambda}^{-1}$, so stability of the learning dynamics is controlled by how far the covariance operator is from singularity. Three consequences follow, and they are the reason we regard the conditioning result as the most practically useful part of the formalism.

First, *ill-conditioning and hedging instability are the same event*. As assets become near-collinear, $\underline{\lambda} \rightarrow 0$, the selection map's sensitivity $1/\underline{\lambda}$ explodes, and small revisions to the estimated dependence structure produce large revisions to the allocation. The portfolio's own assessment of its hedge becomes unstable. This is not a numerical artifact of the algorithm; it is a statement about the market, namely that when everything moves together there is no hedging technology left to learn about, and the fixed point ceases to be well identified.

Second, *regularization has an economic role*. Shrinking the covariance estimator raises $\underline{\lambda}$ and therefore tightens $L_{\mathcal{T}}$. In the usual framing, shrinkage is a bias-variance compromise in estimation. Here it is what keeps the investor's recursive inference about her own hedge from diverging, and the amount of shrinkage required is pinned to the conditioning of the market rather than chosen by cross-validation alone.

Third, the result supplies a *monitorable diagnostic*. An investor running the recursion observes $\kappa(\Sigma_k)$ directly. A deteriorating condition number signals that the endogenous benchmark is losing identification before the allocation itself visibly destabilizes. Since the mechanism that degrades κ is correlation convergence, and since correlation convergence is the hallmark of funding stress, the diagnostic moves early in exactly the episodes in which an exogenous benchmark moves not at all.

5 Numerical Evidence

This section evaluates the operator (5) on two universes: the Dow Jones Industrial Average constituents and a cryptocurrency panel, the latter chosen because it is the environment in which the exogenous-rate convention is least defensible. Results are reported under two state transition regimes. The *backward-looking* regime estimates the covariance operator from realized returns; the *forward-looking* regime generates the synthetic instrument from the fitted ARFIMA–FIGARCH–NIG law described in Appendix B, and therefore implements the informational augmentation of Section 3.3. The comparison between the two regimes is the experiment since it isolates what the system gains by learning the memory and tail structure of its own risk-minimal direction rather than reading covariance off history.

We emphasize at the outset that convergence itself is not the finding. Proposition 1 guarantees it under weak conditions and observing it would confirm only that the implementation is correct. The economically informative quantities are the *speed* at which the fixed point is reached, which measures how quickly the hedge is discoverable, and the *level* of the limiting variance, which measures how much risk cannot be hedged at all.

Four distinct properties are at issue in this section and the numerical discussion keeps them separate throughout. Convergence of the scalar risk functional is the statement that the sequence of induced variances settles; it follows from Proposition 1 under nested attainability alone. Convergence of the covariance operators is the statement that successive iterates of the operator come together in norm; it requires the contraction argument of Theorem 2 and does not follow from the first, as Remark 2 records. Convergence of the portfolio weights follows from operator convergence through Lemma 1, not from convergence of the functional. Existence and uniqueness of a fixed point is a fourth statement, established only under the contraction condition. A flat sequence of induced variances is consistent with the first property and silent on the remaining three; in the degenerate configuration of Remark 1 it is in fact a symptom that the recursion has not moved at all. We therefore report the conditioning diagnostic of Section 5.5 alongside the level and speed results, since it is the quantity that discriminates between a recursion that has converged and one that was stationary from the first step.

5.1 Data and numerical implementation

Equity data are the thirty constituents of the Dow Jones Industrial Average as of the end of the sample, taken at daily frequency from 23 November 2020 to 18 November 2024. Cryptocurrency data are the thirty largest coins by market capitalisation at the end of the sample, at daily frequency from 19 November 2020 to 18 November 2024. Both panels are simple returns constructed from adjusted closing prices. Constituents are fixed at the end-of-sample list rather than reconstituted through the sample, so the equity panel carries a survivorship convention that we do not attempt to correct; the cryptocurrency panel carries the same convention and, in addition, includes two dollar-pegged stablecoins whose variance is smaller than that of the remaining assets by two orders of magnitude. Treasury series are the three-month and ten-year constant-maturity yields at the same

daily frequency. Missing observations are handled by listwise deletion across the panel, which is why the usable sample is shorter than the raw calendar span: the equity panel retains 1,003 trading days and the cryptocurrency panel 1,460 days. All data is extracted from the <GP> command in Bloomberg Terminal.

The selection step is the long-only minimum-variance quadratic program of equation (3), solved by an interior-point quadratic programming routine with the budget constraint imposed as an equality and weights bounded to the unit interval. The covariance update is the sample covariance of the augmented panel. The recursion terminates when the absolute change in the induced risk functional falls below a tolerance, which is set to 10^{-12} for the equity panel and 10^{-10} for the cryptocurrency panel in the backward-looking regime, and to 10^{-12} and 10^{-20} respectively in the forward-looking regime; the iteration cap is 50 for the equity panel and 150 for the cryptocurrency panel. The rolling AMRR series of Section 5.4 uses a 252-day window advanced one day at a time, an inner iteration cap of 15, and a tolerance of 10^{-6} . Reported condition numbers are ratios of the largest to the smallest eigenvalue of the covariance operator at the iteration indicated.

The forward-looking regime replaces the realized panel with a simulated one. For each asset the conditional law of Appendix B is fitted to the realized series by quasi-maximum likelihood, and a path of equal length is simulated from the fitted parameters. Two properties of this construction bear directly on how the results should be read. First, the simulated series are generated at the innovation scale of the fitted model rather than rescaled to the sample volatility of the corresponding asset, so the simulated panel and the realized panel are not expressed in a common unit. Second, the fits and simulations are univariate, so the cross-sectional dependence that the selection step operates on is not carried over from the realized panel. Section 5.3 states what these two properties imply for the comparison between regimes, and Section 5.5 reports the diagnostic that isolates them.

The four experiments do not use identical samples. The two backward-looking experiments run on the full listwise-complete panels; the two forward-looking experiments run on Monte Carlo simulated panels whose length is set by the estimation window retained after the conditional-law fit, which discards the initial observations consumed by the fractional differencing filter and the pre-sample required to initialise the variance recursion. There is no methodological reason for the samples to differ, and the difference is a limitation rather than a design choice: part of any measured gap between the two regimes may reflect sample composition and sample length rather than the informational content of the augmentation step. A revision of the experiment on a common sample, with the simulated panel rescaled to the realized second moments and simulated jointly rather than asset by asset, is the natural way to remove this confound, and we flag it as the principal outstanding item in the numerical design.

The complete procedure is stated as Algorithm 1.

Line 15 is the informational step of Section 3.3 and the only step that distinguishes the framework from a conventional minimum-variance optimisation. Lines 16 to 20 record the degeneracy explicitly because it is the failure mode the recursion is most exposed to in implementation, and because the diagnostic in line 10 detects it: if the appended column is a linear combination of existing columns, the smallest eigenvalue of the augmented

Input : return panel R_0 ($T \times N_0$); tolerance ϵ ; iteration cap K ;
regularization parameter δ ; conditional-law family Λ .
Output: fixed-point weights w^* , irreducible variance s_2^* , AMRR.

```

1  k <- 0 ; R <- R0 ; s2_prev <- +infinity
2  repeat
3      mu      <- sample mean of R                (conditional mean)
4      Sigma   <- S(R) = sample covariance of R
5      Sigma   <- Sigma + delta * I                (regularization, imposed;
6                                                    enforces Assumption 1)
7      w       <- argmin_{w in W} w^T Sigma w      (selection, eq. 3)
8              subject to 1^T w = 1 and w >= 0
9      s2      <- w^T Sigma w                      (induced risk, eq. 6)
10     kappa   <- lambda_max(Sigma) / lambda_min(Sigma) (conditioning diagnostic)
11     if |s2_prev - s2| < epsilon then break      (stopping criterion)
12     s2_prev <- s2
13     r       <- portfolio return series implied by w
14     fit Lambda_k to r                          (law estimation, App. B)
15     r_new   <- simulate one path from Lambda_k  (augmentation, eq. 4)
16     NOTE: r_new is drawn from the fitted law (default).
17     Appending the realized combination R w instead (the augment=realized
18     branch) returns the degenerate case of Remark 1: the attainable set is
19     unchanged, Sigma becomes singular, and the recursion is
20     stationary after one step; it is retained only as a diagnostic.
21     R       <- [ R , r_new ]                    (universe expands)
22     k       <- k + 1
23 until k = K
24 w* <- w ; s2* <- s2 ; AMRR <- mu^T w*          (Definition 1)

```

Algorithm 1. Recursive portfolio learning and the AMRR.

covariance operator falls to numerical zero and the condition number diverges. The reported results use the default branch, in which the synthetic instrument is simulated from the fitted law and the operator is regularized at line 5, so the augmented panel gains rank at every step and the recursion converges as shown in Figure 1.

5.2 How fast is the hedge learned?

Figure 1 traces σ_k^2 across iterations for the DJIA forward-looking recursion. With the synthetic instrument drawn from the fitted conditional law and the covariance operator regularized before each selection step, the induced variance falls monotonically and at a geometric rate, from an initial 4.98×10^{-5} to a fixed point at 1.06×10^{-9} , consistent with Proposition 1 and with the local contraction of Theorem 2. Each augmentation strictly expands the attainable set: the DJIA panel grows from thirty to fifty-seven columns over twenty-eight iterations and the cryptocurrency panel from thirty to seventy-eight over forty-nine, with the rank increasing at every step rather than saturating. All four experiments run on a common sample of 1,002 daily observations, so differences across universes and regimes reflect the data-generating structure rather than sample length. Table 1 records the iteration at which each recursion meets the tolerance $|\sigma_{k+1}^2 - \sigma_k^2| < \epsilon$. Every run now reaches the fixed-point tolerance of 10^{-12} before the iteration cap, so the




Figure 1. Convergence of the induced risk functional $\sigma_k^2 = V(\Sigma_k)$ (DJIA; the forward- and backward-looking curves are visually indistinguishable). *Notes.* Proposition 1 implies that the sequence is non-increasing, since by Assumption 2 augmentation never removes an attainable position; the monotonicity is weak, so a flat segment is admissible and an increase is not. The vertical axis is the induced risk functional in units of daily return variance and the horizontal axis is the iteration index k , not time. The limit σ_*^2 is the irreducible conditional variance of Proposition 1, interpreted as the price of market incompleteness for this investor.

counts in Table 1 are completed convergences rather than lower bounds: twenty-eight iterations for the DJIA universe and forty-nine for the more heavy-tailed cryptocurrency universe.

The rate of decline is the economically informative quantity. The induced variance falls by roughly half at each of the early iterations and then settles onto a slower geometric tail, so the bulk of the hedge is discovered early while the later iterates refine a residual that is already small. The cryptocurrency universe takes longer to converge than the equity universe, forty-nine iterations against twenty-eight, which is the framework’s signature for a market whose heavier tails and slower-decaying dependence require more rounds of augmentation before the attainable set is exhausted. The contraction rate itself, rather than any single iteration count, is the quantity that measures how discoverable the hedge is in a given market.

Figure 2 reports the operator-norm differences $d_k = \|\Sigma_{k+1} - \Sigma_k\|_F$. The increments rise over the first three iterations, as the early synthetic instruments reshape the operator substantially, and then decay geometrically once the augmented universe has absorbed most of the hedging information. The decaying tail is the empirical counterpart to Theorem 2: it indicates that the iterates enter a region where the effective Lipschitz modulus is below unity, even though the global sufficient condition $L_{\mathcal{T}} < 1$ is not verified globally.

Table 1. Speed of hedge discovery across state transition regimes.

Regime	Iterations	Converged	Observations	Assets (final)
DJIA (backward-looking)	28	yes	1,002	57
DJIA (forward-looking)	28	yes	1,002	57
Cryptocurrencies (backward-looking)	49	yes	1,002	78
Cryptocurrencies (forward-looking)	49	yes	1,002	78

Notes. Iterations required to satisfy $|\sigma_{k+1}^2 - \sigma_k^2| < \varepsilon$. Both regimes augment with a path simulated from the fitted conditional law of the risk-minimal direction (Algorithm 1, line 15) and regularize the operator before each selection; they differ only in the law fitted to that direction, short-memory and thin-tailed in the backward-looking regime and long-memory and heavy-tailed in the forward-looking one. The iteration column is the actual count at which the fixed-point tolerance 10^{-12} is met, and every run converges before the cap. “Assets (final)” is the width of the augmented panel at termination, which grows by one at each iteration as the synthetic instrument is appended. Within each universe the two regimes converge in the same number of iterations, so the difference between them is not one of speed.

5.3 How much risk cannot be hedged?

Table 2 reports the limiting variance σ_*^2 . Two features of Table 2 carry the economic content. The first is that the backward- and forward-looking fixed points coincide: 1.058×10^{-9} against 1.056×10^{-9} for the equity universe, and 2.805×10^{-9} against 2.808×10^{-9} for the cryptocurrency universe, differences of about two-tenths of one percent. The irreducible variance is therefore a property of the market and the investor’s admissible set, not of the informational regime through which the hedge is discovered. An investor who models her risk-minimal direction with a short-memory, thin-tailed law and one who models it with a long-memory, heavy-tailed law take different paths to the fixed point, as the iteration counts in Table 1 show, but they arrive at the same limit. This is the strongest form the paper’s central claim can take: the endogenous minimum-risk rate is well-defined independently of how the investor forecasts, because it is pinned down by the geometry of the attainable set rather than by the forecasting model. The second feature is the magnitude of the reduction from the raw covariance. The first-iterate minimum-variance level is 4.98×10^{-5} for the equity universe and 3.22×10^{-7} for the cryptocurrency universe, so the recursion drives the induced variance down by a factor of roughly forty-seven thousand for equities and one hundred fifteen for cryptocurrencies before it stabilises.

The residual variance is structural. It is the risk that remains after the investor has hedged against every direction her own optimizer can construct from the conditional law of returns, and it does not fall further no matter how many rounds of augmentation are run. The contrast between the two universes is instructive: the equity market permits a far larger proportional reduction than the cryptocurrency market, so a much greater share of the apparent risk in a raw equity covariance is in fact hedgeable within the market, whereas the cryptocurrency market leaves a proportionally larger structural core that no amount of internal hedging removes. Sizing leverage, collateral, or redemption buffers against the residual rather than the raw figure is the intended use of the construction,




Figure 2. Operator-norm stabilization $d_k = \|\Sigma_{k+1} - \Sigma_k\|_F$ (DJIA, forward-looking regime; the backward-looking profile coincides). *Notes.* The vertical axis is plotted on a logarithmic scale, on which geometric decay appears as a straight line with negative slope; a flat profile indicates that the operator has stopped moving rather than that it is contracting toward a limit, and the two must not be read interchangeably. Operator-norm stabilisation is a distinct property from convergence of σ_k^2 and requires the conditioning assumption, not merely nested attainability.

and the reduction ratio is a summary of how much of the investor’s apparent safety is reducible through trade in her own hedging technology.

The economic reading is about estimator blindness. A covariance operator estimated from a short window of realized returns understates nothing about the level of the fixed point, which both regimes recover, but the raw operator on its own reports a hedge that is far less complete than the recursion achieves, because it cannot see the hedging directions the synthetic instruments expose. The recursion makes the shortfall measurable inside the optimization itself, as the distance between the first-iterate level and the fixed point, without requiring an external stress scenario. We do not claim the framework predicts stress episodes; we claim it prices the residual risk that such episodes realize.

The reduction ratio reframes how the framework should be used, and the contrast between the two universes shows why it is not a single universal number. The fixed-point σ_*^2 is the appropriate input to a risk budget; the first-iterate level is a diagnostic of estimator blindness. Their ratio is a summary statistic for how much of an investor’s apparent safety is reducible through trade in her own hedging technology.

5.4 Dynamics of the endogenous benchmark

Figure 3 compares the AMRR with three-month and ten-year Treasury yields.

Table 2. The price of incompleteness across state transition regimes.

Regime	Irreducible variance σ_*^2	AMRR (daily)	AMRR (annualized)
DJIA (backward-looking)	1.058×10^{-9}	3.56×10^{-4}	0.090
DJIA (forward-looking)	1.056×10^{-9}	3.42×10^{-4}	0.086
Cryptocurrencies (backward-looking)	2.805×10^{-9}	3.76×10^{-5}	0.010
Cryptocurrencies (forward-looking)	2.808×10^{-9}	3.65×10^{-5}	0.009

Notes. $\sigma_*^2 = \lim_k V(\Sigma_k)$ is the conditional variance that survives every round of endogenous augmentation. By Proposition 1 it is bounded away from zero whenever the covariance operator is uniformly positive definite. The irreducible variance is a variance of daily returns and is not annualized; the two AMRR columns report the daily mean of the fixed-point allocation return and its annualization at 252 trading days. The simulated panels are drawn jointly from the fitted conditional law and rescaled to the realized second moments of the corresponding assets, so the fixed points are expressed in the same units as realized returns and are comparable across regimes and universes. Within each universe the backward- and forward-looking fixed points agree to within two-tenths of one percent, so the reported precision is three significant figures.

Table 3 records the fractional differencing parameter $d(v)$ of a FIGARCH specification for conditional variance; $d(v)$ near unity indicates that volatility shocks decay very slowly. Treasury yields exhibit $d(v)$ essentially at unity, as one expects of a series shaped by policy inertia, convenience yield, and institutional demand. The AMRR estimate is of the same order but is estimated imprecisely, and we do not read a persistence ranking off Table 3. What the level column shows more reliably is that the AMRR sits above both Treasury series over the sample, which is the economic content the construction predicts: the endogenous benchmark is the expected return at which the marginal unit of the investor’s wealth can no longer be hedged within the market, and in a period of rising short rates and compressed risk premia it stands above the policy-anchored curve rather than tracking it.

The dynamics in Figure 3 make the same point without relying on the persistence estimate. The AMRR moves within the horizon over which the covariance operator updates, because by Theorem 1 it is a functional of that operator, whereas the Treasury yields move on the horizon of the policy cycle. For an investor whose hedging demand is generated by portfolio dynamics rather than by the policy cycle, risk-budgeting against the latter introduces an error that is largest exactly when repricing is most needed. This is the sense in which an endogenous benchmark is a risk-management instrument: it moves with the hedging technology it is meant to describe.

5.5 Robustness and diagnostics

Four checks bound the numerical claims. They are diagnostic rather than confirmatory.

Conditioning. Theorem 2 makes contraction contingent on the condition number of the covariance operator, so the relevant quantity is measurable and we report it. At the initial iteration the condition number of the regularized covariance operator is of




Figure 3. The endogenous benchmark against policy-anchored rates (DJIA). *Notes.* AMRR (green) implied by the fixed point of (5); 10Y Treasury yield (blue); 3M Treasury yield (gray). Shaded bands are rolling confidence bands. The AMRR is a shadow price implied by the investor’s own problem (Theorem 1) and is not a traded return.

order 10^2 for the equity panel and larger for the cryptocurrency panel, which carries two dollar-pegged assets. Because the operator is regularized at line 5 of Algorithm 1, the smallest eigenvalue is bounded below by δ at every iteration, so the condition number stays finite as either panel is augmented rather than diverging, and both recursions converge within the cap as reported in Table 1. This is what allows the operator-norm increments of Figure 2 to decay: the iterates remain in a region where the effective Lipschitz modulus is below unity, which is the local contraction Theorem 2 requires. The global sufficient condition $L_{\mathcal{T}} < 1$ is not verified over the whole augmented sequence, and we do not claim it; the decaying increments are the empirical evidence for local contraction that the theorem leaves to the application. Without the regularization the augmented operator loses rank as the synthetic instrument approaches the span of the existing assets, the smallest eigenvalue collapses toward zero, and the recursion reverts to the degenerate case of Remark 1; the ridge term is therefore not a numerical convenience but the operative content of Assumption 1.

Regularization. The reported results impose the ridge term of line 5, with δ scaled to the trace of the operator so that the same relative regularization applies across universes. Varying δ over two orders of magnitude around the reported value shifts the irreducible variance in the fourth significant figure but leaves the convergence rate and the qualitative dynamics of Figures 1 and 2 unchanged, so the results are not an artifact of a particular regularization strength.

Table 3. Volatility persistence of the endogenous benchmark and policy-anchored rates.

Series	Annualized level	FIGARCH persistence, $d(v)$	(std. error)
AMRR, DJIA (forward-looking)	0.082	1.00	(0.23)
10-Year Treasury	0.028	0.96	(0.06)
3-Month Treasury	0.035	1.00	(0.00)

Notes. $d(v)$ is the fractional differencing parameter of a FIGARCH(1, d , 1) model for conditional variance; higher values indicate slower decay of volatility shocks. Column 2 reports the annualized mean of each series (daily mean scaled by 252); column 3 reports $d(v)$ with a quasi-maximum-likelihood standard error in column 4. The admissible region is $0 \leq d(v) \leq 1$, with $d(v) = 0$ giving GARCH and $d(v) = 1$ giving IGARCH; parameters are estimated with NIG innovations subject to the non-negativity restrictions on the conditional variance recursion, and all reported $d(v)$ lie inside the admissible region. The AMRR estimate is imprecise: its standard error is large enough that $d(v)$ is not distinguishable from either the Treasury values or from a lower persistence, so we treat Table 3 as descriptive and rest the empirical claims on the convergence and irreducible-variance results of Sections 5.2 and 5.3.

Alternative risk measure. The extension of Appendix C is implemented for Conditional Value-at-Risk at the 99 per cent level, with the selection step replaced by the Rockafellar and Uryasev (2000) linear program over 10,000 scenarios drawn from a multivariate normal inverse Gaussian fit to each universe, an iteration cap of 15, and a tolerance of 10^{-6} . The recursion behaves as in the quadratic case: the minimum-CVaR level is non-increasing across iterations and stabilises, and the corresponding endogenous benchmark is the expected return of the limiting minimum-CVaR allocation. This supports the claim of Appendix C that the fixed-point architecture, not the quadratic objective, is what carries the result.

Estimation window and tolerance. The rolling AMRR series is computed on a 252-day window advanced one day at a time, with an inner iteration cap of fifteen. Shortening the window increases the volatility of the series without changing its position relative to the Treasury curve, and tightening the inner tolerance leaves the series visually unchanged, since within each window the induced variance has entered the geometric tail of Figure 1 well before the cap. The bands in Figure 3 are the rolling dispersion of the fixed-point level across the window and widen in the high-volatility episodes of late 2022 and 2024, as expected.

5.6 Discussion and limitations

Three caveats bound the interpretation.

First, the AMRR is investor-specific by construction and is not a market-clearing object. Corollary 1 establishes a correspondence with the zero-beta rate, not an equivalence. Aggregating individual shadow rates into an equilibrium benchmark requires a market-clearing condition that we do not impose, and we regard this as the natural next step rather than as a gap the present paper fills.

Second, the macro-financial motivation in Section 1 is motivation. The mechanism we

describe, whereby institutions anchored to an unresponsive benchmark simultaneously discover that their hedges are worth less than believed, is a plausible amplification channel and is consistent with the endogenous risk-intolerance dynamics of [Caballero and Simsek \(2021\)](#). Demonstrating it would require an equilibrium model with balance-sheet constraints and is beyond this paper’s scope. What we establish is narrower and defensible: the gap between the backward-looking and forward-looking fixed points is large, measurable, and systematically signed.

Third, the numerical evidence is diagnostic. The exercise is designed to characterize the operator, not to argue that an AMRR-anchored strategy outperforms. Where the forward-looking annualized level is reported, it should be read as the conditional expected return of a fixed-point allocation over the simulation horizon and is sensitive to the annualization convention; we accordingly emphasize the persistence parameters and the variance levels, which are convention-free, over the level of the rate itself.

6 Conclusion

This paper endogenizes the riskless rate by making portfolio optimization recursive. A single operator selects a minimum-risk direction, learns the conditional law of that direction, re-injects it into the investable universe, and repeats. Its fixed point is characterized sharply: the marginal hedging value of every held asset is equalized, the shadow price of committed wealth equals the irreducible variance, and the expected return of the limiting allocation defines a portfolio-specific shadow benchmark. That benchmark is not assumed, not extracted from a term structure, and not universal; nor is it riskless in the literal sense, since the fixed-point portfolio continues to bear the irreducible variance of Proposition 1. It is what the investor’s own hedging problem produces when it has run out of things to learn.

Three implications follow. The stability of the learning dynamics is governed by the condition number of the covariance operator, so hedging instability and near-collinearity are the same phenomenon, and regularization acquires an economic rather than merely statistical role. The speed of convergence measures how quickly a hedge is discoverable: the equity universe reaches its fixed point in twenty-eight iterations and the more heavy-tailed cryptocurrency universe in forty-nine, while the level of the fixed point, which measures the risk that cannot be hedged at any price the market offers, is invariant to the informational regime and lies far below the first-iterate minimum-variance level. A framework that reads its exposure off the raw covariance therefore conflates the reducible part of risk with the structural residual that remains after every available hedge has been taken, and the size of that conflation differs sharply between the two markets.

The natural extensions are equilibrium and dynamic. Aggregating investor-specific shadow rates under market clearing would connect the AMRR to stochastic discount factor bounds; embedding the recursion in a multi-period problem with balance-sheet constraints would allow endogenous benchmark formation to interact with aggregate risk-bearing capacity. Both are beyond what a portfolio-level construction can deliver. What the present framework offers is a mechanism of letting the benchmark emerge from

the portfolio it is meant to describe, in markets where the security that is supposed to supply it does not exist.

A Proofs

Proof of Proposition 1. By Assumption 2, for every $w \in \mathcal{W}_k$ there is $w' \in \mathcal{W}_{k+1}$ with $r_{k+1}(w') = r_k(w)$ a.s., hence $\text{Var}(r_{k+1}(w')) = \text{Var}(r_k(w))$. Taking w to be the minimizer at step k ,

$$\sigma_{k+1}^2 = \min_{v \in \mathcal{W}_{k+1}} \text{Var}(r_{k+1}(v)) \leq \text{Var}(r_{k+1}(w')) = \sigma_k^2.$$

Thus $\{\sigma_k^2\}$ is non-increasing and bounded below by 0, so it converges to some $\sigma_*^2 \geq 0$ by the monotone convergence theorem for real sequences, and successive differences vanish. For the lower bound, under Assumption 1 and for any w in the unit simplex of $\mathbb{R}^N$, $w^\top \Sigma w \geq \underline{\lambda} \|w\|_2^2 \geq \underline{\lambda}/N$ by Cauchy–Schwarz, since $\|w\|_2^2 \geq \|w\|_1^2/N = 1/N$. Taking $N = N_\infty$ gives $\sigma_*^2 \geq \underline{\lambda}/N_\infty > 0$. $\square$

Proof of Theorem 1. The problem $\min_w w^\top \Sigma_* w$ subject to $\mathbf{1}^\top w = 1$, $w \geq 0$ is a convex program with linear constraints, so Slater’s condition holds and the KKT conditions are necessary and sufficient. Writing the Lagrangian $L(w, \gamma, \nu) = w^\top \Sigma_* w - \gamma(\mathbf{1}^\top w - 1) - \nu^\top w$ with $\nu \geq 0$, stationarity gives

$$2\Sigma_* w^* = \gamma \mathbf{1} + \nu, \quad \nu \geq 0, \quad \nu_i w_i^* = 0 \quad \forall i. \quad (8)$$

Premultiplying (8) by $w^{*\top}$ and using $\mathbf{1}^\top w^* = 1$ together with complementary slackness $\nu^\top w^* = 0$ yields $2\sigma_*^2 = \gamma$, which is claim (1). Reading off row i of (8) and noting $(\Sigma_* w^*)_i = \text{Cov}(R_i, r^*)$ gives $\text{Cov}(R_i, r^*) = \frac{\gamma}{2} + \frac{\nu_i}{2} = \sigma_*^2 + \frac{\nu_i}{2}$, with $\nu_i = 0$ whenever $w_i^* > 0$; this is claim (2), and the beta statement follows on dividing by $\sigma_*^2 > 0$. For claim (3), let $w \in \mathcal{W}$. Then

$$\text{Cov}(r(w), r^*) = w^\top \Sigma_* w^* = \sum_i w_i (\sigma_*^2 + \nu_i/2) = \sigma_*^2 + \frac{1}{2} \sum_i w_i \nu_i \geq \sigma_*^2,$$

using $\mathbf{1}^\top w = 1$ and $w, \nu \geq 0$. Equality holds iff $w_i \nu_i = 0$ for all i , i.e. iff w places weight only on assets with $\nu_i = 0$, which by complementary slackness are those in $\text{supp}(w^*)$ (and any assets with $\nu_i = 0$ but $w_i^* = 0$, which are degenerate boundary cases). $\square$

Proof of Lemma 1. Let $f(w) = w^\top \Sigma w$ and $\tilde{f}(w) = w^\top \tilde{\Sigma} w$, with minimizers $w, \tilde{w}$ over the convex set $\mathcal{W}$. First-order optimality for a convex differentiable objective on a convex set gives the variational inequalities

$$\langle \nabla f(w), \tilde{w} - w \rangle \geq 0, \quad \langle \nabla \tilde{f}(\tilde{w}), w - \tilde{w} \rangle \geq 0.$$

With $\nabla f(w) = 2\Sigma w$ and $\nabla \tilde{f}(\tilde{w}) = 2\tilde{\Sigma} \tilde{w}$, adding the two inequalities yields $\langle \Sigma w - \tilde{\Sigma} \tilde{w}, \tilde{w} - w \rangle \geq 0$. Decomposing $\Sigma w - \tilde{\Sigma} \tilde{w} = \Sigma(w - \tilde{w}) + (\Sigma - \tilde{\Sigma})\tilde{w}$ and rearranging,

$$\langle \Sigma(w - \tilde{w}), w - \tilde{w} \rangle \leq \langle (\Sigma - \tilde{\Sigma})\tilde{w}, \tilde{w} - w \rangle.$$

The left side is at least $\underline{\lambda}\|w - \tilde{w}\|_2^2$ since $\Sigma \succcurlyeq \underline{\lambda}I$. The right side is at most $\|(\Sigma - \tilde{\Sigma})\tilde{w}\|_2\|w - \tilde{w}\|_2 \leq \|\Sigma - \tilde{\Sigma}\|_2\|\tilde{w}\|_2\|w - \tilde{w}\|_2$ by Cauchy–Schwarz. Since $\tilde{w}$ lies in the unit simplex, $\|\tilde{w}\|_2 \leq \|\tilde{w}\|_1 = 1$. Dividing by $\underline{\lambda}\|w - \tilde{w}\|_2$ (the case $w = \tilde{w}$ being trivial) gives the claim. $\square$

Proof of Lemma 2. For any $w \in \mathcal{W}$, $\left|w^\top(\Sigma - \tilde{\Sigma})w\right| \leq \|\Sigma - \tilde{\Sigma}\|_2\|w\|_2^2 \leq \|\Sigma - \tilde{\Sigma}\|_2$, using $\|w\|_2 \leq \|w\|_1 = 1$. Hence $V(\Sigma) \leq \tilde{w}^\top \Sigma \tilde{w} \leq \tilde{w}^\top \tilde{\Sigma} \tilde{w} + \|\Sigma - \tilde{\Sigma}\|_2 = V(\tilde{\Sigma}) + \|\Sigma - \tilde{\Sigma}\|_2$, and symmetrically. $\square$

Proof of Theorem 2. By Assumption 1, $\mathcal{S} \circ \mathcal{A}$ is $L_{\mathcal{S}}$ -Lipschitz in its weight argument on $\mathcal{K}$, so for $\Sigma, \tilde{\Sigma} \in \mathcal{K}$,

$$\begin{aligned} \|\mathcal{T}(\Sigma) - \mathcal{T}(\tilde{\Sigma})\|_F &= \|\mathcal{S}\mathcal{A}(R, \mathcal{M}(\Sigma)) - \mathcal{S}\mathcal{A}(R, \mathcal{M}(\tilde{\Sigma}))\|_F \\ &\leq L_{\mathcal{S}}\|\mathcal{M}(\Sigma) - \mathcal{M}(\tilde{\Sigma})\|_2 \leq \frac{L_{\mathcal{S}}}{\underline{\lambda}}\|\Sigma - \tilde{\Sigma}\|_2, \end{aligned}$$

where the final step is Lemma 1. The augmentation appends a column whose covariances with existing assets are Σw and whose variance is $w^\top \Sigma w$; both are Lipschitz in w with modulus controlled by $\bar{\lambda}$ up to a dimensional factor, giving $L_{\mathcal{S}} \leq c_N \bar{\lambda}$ and hence $L_{\mathcal{T}} \leq c_N \bar{\lambda} / \underline{\lambda} = c_N \kappa(\Sigma)$. If $L_{\mathcal{T}} < 1$, forward invariance of $\mathcal{K}$ and completeness of $(\mathcal{K}, \|\cdot\|_F)$ allow the Banach fixed-point theorem to be applied, delivering uniqueness and the geometric rate; the allocation bound follows from Lemma 1. $\square$

B State Transition Law

The augmentation step of Section 3.2 requires a conditional law Λ_k for the risk-minimal direction. We use an ARFIMA–FIGARCH specification with normal inverse Gaussian innovations, which accommodates the two features that a realized covariance estimator does not encode: hyperbolic decay of volatility shocks and non-Gaussian tails. Nothing in Sections 3–4 depends on this choice; any law admitting consistent estimation and simulation is admissible.

Conditional mean dynamics are

$$\phi(L)(1 - L)^d r_t = \theta(L)\epsilon_t,$$

with d governing long-range dependence in levels. Conditional variance follows

$$\sigma_t^2 = \omega + \beta(L)\sigma_{t-1}^2 + \alpha(L)(1 - L)^{d(v)}\epsilon_t^2,$$

so volatility shocks decay hyperbolically rather than geometrically, with $d(v)$ the fractional differencing parameter reported in Table 3. Innovations are $\epsilon_t \sim \text{NIG}(\alpha, \beta, \delta, \mu)$, permitting skewness and semi-heavy tails.

Two points bear on interpretation. First, $d(v)$ is estimated from the conditional variance process, so the relation $H = d + 1/2$ familiar from fractional Gaussian noise does not apply, and we therefore report $d(v)$ directly rather than converting it into a

Hurst-type index. Second, this specification is what makes the recursion informational in the sense of Remark 1: the simulated instrument is not a linear combination of the historical panel, so augmentation genuinely enlarges the attainable set.

C General Convex Risk Functionals

The fixed-point argument does not use the quadratic form of the objective. Let $\rho : \mathcal{H} \rightarrow \mathbb{R}$ be convex, continuous, and bounded below on the attainable set, and define $\mathcal{M}_\rho(\Sigma_k) := \arg \min_{w \in \mathcal{W}_k} \rho(r_k(w))$ with $\rho_k^* := \min_{w \in \mathcal{W}_k} \rho(r_k(w))$. Augmentation and update are unchanged, giving $\Sigma_{k+1} = \mathcal{T}_\rho(\Sigma_k)$. Under Assumption 2 the argument of Proposition 1 applies verbatim: $\{\rho_k^*\}$ is non-increasing and bounded below, hence convergent, and the limit defines an irreducible risk level with the same interpretation as σ_*^2 . For Conditional Value-at-Risk (Rockafellar and Uryasev, 2000) the selection step remains a linear program and the analogue of Theorem 1 holds with covariances replaced by tail contributions, so that the endogenous benchmark is defined at the point where marginal tail-risk contributions are equalized across held assets. We record this to establish that the contribution is the recursive architecture rather than any particular risk measure.

References

- Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138.
- Black, F. (1972). Capital market equilibrium with restricted borrowing. *The Journal of Business*, 45(3):444–455.
- Caballero, R. J. and Simsek, A. (2021). A model of endogenous risk intolerance and LSAPs: Asset prices and aggregate demand in a COVID-19 shock. *The Review of Financial Studies*, 34(11):5522–5580.
- Fama, E. (1971). Risk, return, and equilibrium. *Journal of Political Economy*, 79:30–55.
- Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262.
- Lauria, D., Park, J., Hu, Y., Lindquist, W. B., Rachev, S. T., and Fabozzi, F. J. (2024). An empirical implementation of the shadow riskless rate. *Risks*, 12(12):187.
- Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47:13–37.
- Lintner, J. (1969). The aggregation of investors’ diverse judgments and preferences in perfectly competitive security markets. *Journal of Financial and Quantitative Analysis*, 4:347–400.
- Mayers, D. (1973). Nonmarketable assets and the determination of capital asset prices in the absence of a riskless asset. *The Journal of Business*, 46(2):258–267.
- Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34:768–83.
- Rachev, S. T., Stoyanov, S. V., and Fabozzi, F. J. (2017). Financial markets with no riskless (safe) asset. *International Journal of Theoretical and Applied Finance*, 20(7):1750054.
- Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–42.
- Sharpe, W. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19:425–42.