

When Does Diversification Help?

Aggregate versus Per-Asset Nonlinear Impact in Multi-Asset Trading

Xin Heng*
Tote AI
xin@tote.ai

This version: September 25, 2026

Abstract

Trading moves prices against the trader, and this price impact is concave. For one asset, Brokmann et al. (2025) show that a linear trading rule calibrated to an “effective” quadratic cost loses only about 0.5–2% of net Sharpe ratio. We ask the same question for correlated multi-asset portfolios, where the cost can be charged on the whole portfolio package or on each name separately. The literature uses both forms, and neither is settled empirically.

The central result is that impact pooling changes how the measured value of a nonlinear rule scales with diversification. Under package pricing, its achieved gain over the best linear rule falls roughly as one over the number of independent directions traded: at $p = 1.3$, from 2.1% at two names to 0.6% at ten. Under name-by-name pricing it stays near 3%. A broad book executed name by name looks safest, yet it is exactly the case in which the measured gain does not shrink (Sections 6–10).

We certify the two-asset cost of linearity from both sides, with a nonlinear policy below and an information-relaxation dual bound above: 0.66–2.39% at the square-root law and 2.14–7.23% at stronger curvature (Section 4). A closed-form elasticity predicts how nonlinear the rule should be across a 400-fold range of costs (Section 5). Fitting a separate cost coefficient for each kind of trade adds at most 0.45% anywhere measured—less than one-fifth of the 2.3-percentage-point measured effect of choosing the cost model. At an aggregate-zero point, per-asset impact generally adds a leverage-score condition for one coefficient to suffice; we prove this local separation (Section 8, Appendix A.1). The practical order is to identify how execution is charged, then decide on nonlinearity, and only then refine the cost coefficient.

1 The question

Suppose you run a portfolio with a return forecast, you care about risk, and trading moves prices against you. You would like to know how to trade.

If the cost of trading were exactly quadratic in the size of the trade, this problem has a clean answer. Gârleanu and Pedersen (2013) show the optimal rule is linear: each period you move part of the way from where you are toward a target that blends today’s forecast with tomorrow’s expected forecast. The rule has no threshold, cap, or change of aggressiveness for an unusually strong opportunity.

*This research was conducted independently of Tote AI and reflects solely the author’s views.

The matrix foundation for trading many assets is well established. Almgren and Chriss (2001) pair temporary impact with a positive-definite cost matrix, and Gârleanu and Pedersen (2013, 2016) carry the quadratic cost $\frac{1}{2}u^\top \Lambda u$ into dynamic portfolio choice. Mei et al. (2016) study several assets with power-form costs, and Moreau et al. (2017) and Bayraktar et al. (2021) derive small-impact asymptotics for several assets. Those papers vary how *large* impact is around a frictionless market. **Our question** is different: we vary the cost *exponent* around the exactly quadratic case and keep costs finite.

The quadratic model is a benchmark rather than a literal description of impact. Empirically, price moves roughly with the square root of the size traded, so total cost grows like size to the power $p \approx 1.5$, not $p = 2$. The clean answer then becomes a nonlinear control problem with no known closed form. Guasoni and Weber (2020) solve a power-impact portfolio problem asymptotically, but for a single risky asset.

Brokmann et al. (2025) take a simpler route for one asset. In the quadratic model, the cost coefficient controls how quickly the rule trades toward its target. For illustration, if the current position is \$1 million and the forecast calls for \$2 million, a low coefficient might make the rule buy \$800,000 now, while a high coefficient might make it buy only \$200,000. Both rules are linear: doubling the gap doubles the trade. Brokmann et al. try different coefficients, score every resulting rule under the true nonlinear cost, and keep the best; they call the winning coefficient the “effective” quadratic cost. At their baseline $p = 1.6$ this rule loses only about 0.5–2% of net Sharpe ratio against a high-accuracy optimum.

That is a single asset. In their Remark 3.1 they defer the correlated case, because “there does not seem to be a generic nonlinear price impact model for multiple assets” and a numerical benchmark is hard to compute. That leaves two tasks: choose a multi-asset impact model and compute a benchmark. We do both. The main result is that the achieved gain from a nonlinear rule scales differently under aggregate and per-asset impact.

Consider two desks facing the same book. Desk A models package pricing and, because the measured nonlinear advantage shrinks with diversification, stays linear. Desk B models name-by-name pricing and builds a nonlinear rule because the measured advantage remains material. If execution is actually name by name, Desk A forgoes the larger observed gain. If execution is package-priced, Desk B mainly bears excess implementation complexity. The measured model risk is therefore asymmetric: mistaking name-by-name execution for package pricing is the more costly error in our experiments.

Contribution and relation to prior work

We build on established tools. Information-relaxation dual bounds provide the upper certificate (Brown et al., 2010; Brown and Smith, 2011; Haugh and Wang, 2014), and Moallemi and Sağlam (2017) compare a best linear policy with such a bound in execution. The stand-in quadratic cost is related to equivalent-LQ and statistical-linearization methods (Booton, 1954; Kazakov, 1956; Yoshida, 1984; Farison and Elhazzouri, 1987; Heess, 1970; Beaman, 1984); in finance, Durin (2014) rewrites power-3/2 trading costs as effective quadratic costs for attribution. Our use differs: the Gaussian state is exact under a linear policy, so the stand-in cost is a self-consistent multi-asset matrix.

The two multi-asset cost forms we compare are both established: the aggregate form is proposed in the preprint version of Schied et al. (2010), and Mei et al. (2016) use the separable form. Concave cross-impact models, in which the power acts before a propagator mixes the

assets, sit between them (Muhle-Karbe and Tracy, 2024; Hey et al., 2024, 2025); the closest, Hey et al. (2024), characterizes which such models are free of price manipulation in risk-neutral execution. On the linear side, Abi Jaber et al. (2026) solve multi-asset portfolio choice explicitly under a matrix-valued cross-impact propagator, the tractability a concave cost gives up. The new question here is the comparison: what changes when the two established forms are placed side by side on the same risk-averse portfolio with return forecasts?

To our knowledge, the new contribution is the answer to that question in the finite-cost correlated setting, together with three consequences: the dependence of the whole effect on whether impact pools across assets (Section 7, Appendix A.1); a local elasticity formula that predicts the fitted nonlinearity (Section 5); and a projection formula for how little is lost by fitting one cost number instead of one per kind of trade (Section 8, Appendix A). That last question is the natural multi-asset extension of Brokmann’s construction: should one fit a single effective cost number for the whole book, or a different number for buying both stocks together versus buying one and selling the other? A single well-fitted number captures nearly all of the gain.

2 Model setup

2.1 State, trades, and period objective

There are n assets. They may be equities, futures, liquid bonds, or other traded instruments for which the risk and impact inputs below can be calibrated; stock examples are illustrations, not a restriction. Write $x_t \in \mathbb{R}^n$ for the position held before trading at time t , u_t for the trade, and $y_t = x_t + u_t$ for the position afterwards, in money units. Positions are holdings, not wealth: y_t is the vector of dollars in each name after the trade.

Forecasts arrive as a k -dimensional signal f_t , not as the asset return itself. The signal follows a Gaussian AR(1),

$$f_{t+1} = Af_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Omega),$$

with persistence matrix A and innovation covariance Ω . A fixed loading matrix $B \in \mathbb{R}^{n \times k}$ converts that signal into a conditional expected excess return on the n assets: given f_t , the manager’s expected P&L from holding y_t over the period is $y_t^\top B f_t$. In the exchangeable base case below, $B = bI$, so each unit of signal raises expected return by the same scalar b ; nothing in B is estimated inside the dynamic program.

The quantity maximized is not terminal wealth. Each period contributes a mark-to-market score

$$r_t = y_t^\top B f_t - \frac{\gamma}{2} y_t^\top \Sigma y_t - \frac{\kappa}{p} c(u_t), \quad (1)$$

and the objective is the long-run average of r_t . The first term is expected P&L from the position held. The second is a mean–variance risk penalty on that same position: Σ is the covariance of asset returns and $\gamma > 0$ is risk aversion. Risk is charged on y_t , not on the trade u_t , because market risk attaches to what you hold. If the book is already large and $u_t = 0$, the risk term is still large; you are not charged impact, but you are still exposed. Trading cost is the third term and depends on u_t only.

Here $\kappa > 0$ is the scale of trading cost and $p \in (1, 2]$ is its exponent; the shape c is defined in Section 2.2. The factor $1/p$ is the usual normalization so that the marginal cost of a one-asset

trade $|u|^p$ is $\kappa|u|^{p-1}$, not $p\kappa|u|^{p-1}$ (Gârleanu and Pedersen, 2013; Brokmann et al., 2025). At $p = 2$ the cost term becomes $\frac{\kappa}{2}c(u_t)$, matching the standard quadratic penalty. One could absorb $1/p$ into a renamed scale; keeping it makes $p = 2$ line up with the quadratic literature and makes first-order conditions clean. This is the Gârleanu and Pedersen (2013) setting with the quadratic cost replaced by a power law, and it matches Brokmann et al. (2025) when $n = 1$.

2.2 Two ways to write $c(u)$

For one asset the cost is $c(u) = |u|^p$ with $p \in (1, 2]$. Empirically, metaorders move prices by roughly Q^δ with δ between about 0.4 and 0.7 (Tóth et al., 2011; Zarinelli et al., 2015; Bacry et al., 2015; Almgren et al., 2005). Since total cost is impact times quantity, $p = 1 + \delta \in [1.4, 1.7]$, the convention of Brokmann et al. (2025). The power law holds over a bounded size range (Zarinelli et al., 2015; Bucci et al., 2019). We use $p = 1.5$ as the square-root benchmark, $p = 1.7$ as a milder case, $p = 1.3$ as a deliberately stronger-curvature stress outside this empirical range, and $p = 2$, where the linear rule is exactly optimal, as a control.

For several assets there is more than one way to write this down (Brokmann et al., 2025). We use two formulas:

$$\text{aggregate: } c_{\text{agg}}(u) = (u^\top \Lambda u)^{p/2}, \quad (2)$$

$$\text{per-asset: } c_{\text{sep}}(u) = \sum_i \lambda_i^{p/2} |u_i|^p. \quad (3)$$

Both are in money units. Λ is a positive-definite liquidity matrix in the sense of multi-asset temporary-impact models (Almgren and Chriss, 2001; Gârleanu and Pedersen, 2013): in the quadratic benchmark the trading cost is $\frac{1}{2}u^\top \Lambda u$, so the diagonal of Λ scales how expensive it is to trade each name on its own, and off-diagonal entries allow trading one name to move the cost of another (cross-impact). Larger eigenvalues identify less-liquid directions. In the per-asset formula, λ_i is the own-name liquidity coefficient of asset i ; when Λ is diagonal one may take λ_i to be those diagonal entries.

The two formulas agree at $n = 1$ by construction, which is why we weight the per-asset terms by $\lambda_i^{p/2}$ rather than λ_i . They also agree at $p = 2$ whenever Λ is diagonal: both reduce to $\sum_i \lambda_i u_i^2$. They do *not* agree for $p < 2$ even in the isotropic case $\Lambda = \lambda I$ used as the base case below. Then aggregate cost is $\lambda^{p/2} \|u\|_2^p$ while per-asset cost is $\lambda^{p/2} \|u\|_p^p$, and $\|\cdot\|_2^p \neq \|\cdot\|_p^p$ once two or more coordinates of u are nonzero. That gap—how the same dollar trades are pooled before the power is applied—is exactly what Section 7 measures.

Both forms are established: Schied et al. (2010) propose (2) as the natural multi-asset generalization of power-law impact, and Mei et al. (2016) use (3) up to the weight normalization. They represent two economically meaningful limits. Equation (2) is *package pricing*: one shared liquidity pool charges for the whole portfolio trade. Equation (3) is *name-by-name pricing*: each asset is executed in its own book. We use “aggregate” and “package” interchangeably, and likewise “per-asset” and “name-by-name”. A real process can mix both, and the concave cross-impact models of Hey et al. (2024, 2025) sit between them (Section 11).

2.3 Linear policies and base assumptions

The policy class whose opportunity cost we study is linear: each period’s trade is a fixed linear combination of the current position and signal, so doubling the gap to target doubles the trade. Brokmann et al. (2025) is titled *Tackling Nonlinear Price Impact with Linear Strategies*, and

we ask what happens to its conclusion when assets are correlated. At $p = 2$ the optimal policy is exactly linear (Gârleanu and Pedersen, 2013), so $p = 2$ is a control. Linear policies can also be scored exactly: their stationary state is Gaussian, so the expected power cost reduces to one-dimensional quadrature with no simulation noise (Appendix D). The nonlinear policy we use as a benchmark is simulated instead, paired against the linear rule on the same shocks.

Unless stated otherwise, we make the following base assumptions:

- n exchangeable assets with equal correlation ρ , so $\Sigma = \sigma^2[(1 - \rho)I + \rho\mathbf{1}\mathbf{1}^\top]$;
- equal liquidity with no cross-impact, $\Lambda = \lambda I$;
- symmetric signals, $A = \phi I$ and $B = bI$, with $\phi = 0.9$ and innovations scaled so f has unit stationary variance.

This isolates correlation as the only asymmetry in the base case.

We then remove these simplifications one at a time. The table lists each base assumption, what it is changed to, and where:

Base case	Varied to	Where
cost model	aggregate or per-asset	§7, §9
ρ equal	0 to 0.85	§4
p	1.3, 1.5, 1.7, and 2 as control	§4
n	2 to 100	§6, Appendix E.4
κ	400-fold sweep; fixed instead of retention-matched	§5, §9
equal volatility	0.74 to 0.95	§9
equal liquidity	0.45 to 2.47	§9
symmetric signals	unequal strength and persistence, misaligned	§9
isotropic Λ	factor-book-aligned and scrambled cross-impact	§9.1
unconstrained	dollar and beta neutrality	§9

Four elements define the scope throughout: impact is temporary and static, signals are Gaussian AR(1), the policy class whose cost we measure is linear (with nonlinear policies supplying lower bounds), and the economy is synthetic. Section 11 takes each up.

How a linear policy is built: stand-in quadratic costs. The true trading cost is the power cost of Section 2.2, and every policy is scored under that true cost without approximation. To *build* a linear policy, we temporarily pretend the cost is quadratic,

$$\frac{1}{2} u^\top Q u = \frac{1}{2} \sum_{i,j} Q_{ij} u_i u_j,$$

for some positive-definite $n \times n$ matrix Q . With one asset this is just $\frac{1}{2} Q u^2$; with several, the diagonal entries Q_{ii} price trading each name alone and the off-diagonal entries price trading two names together. We call this the *surrogate* (stand-in) cost. A quadratic cost makes the optimal policy linear (Gârleanu and Pedersen, 2013), so each choice of Q produces exactly one linear policy, which we then score under the true power cost. Q is a design knob, not a market input: it never enters the objective. It is the multi-asset version of the single “effective quadratic coefficient” that Brokmann et al. (2025) tune for one asset.

Families of stand-in costs. A *family* is the set of Q matrices we allow the search to choose from. We fit the best Q within each family and compare the resulting policies. The *scalar* family has one free number, one cost number for the whole book:

$$Q_{\text{scalar}} = q\Lambda, \quad q > 0.$$

Here Λ is the liquidity matrix of Section 2.2 and q is a single fitted multiplier that makes trading uniformly cheaper or dearer. At $p = 2$ the aggregate cost is exactly $\frac{1}{2}u^\top(\kappa\Lambda)u$ (and so is the per-asset cost when Λ is diagonal), so $q = \kappa$ recovers the exact answer; for $p < 2$ the best q must be found by search. “Scalar” refers to that single number q , not to the trade. In the symmetric benchmark, where $\Lambda = \lambda I$, the *directional* family gives each independent risk direction, or kind of trade, its own multiplier:

$$Q_{\text{dir}} = V \text{diag}(q_1, \dots, q_n) V^\top,$$

where the columns of V are the eigenvectors of the return covariance Σ . For two similar stocks these two directions are “buy both together” and “buy one, sell the other”; the directional family may charge those two trades differently, whereas the scalar family charges them in fixed proportion. In the heterogeneous stress test, the directional family is instead parameterized around $q\Lambda$ in model-specific impact-whitened directions, so the scalar family remains nested. We also test a *full* family in which Q may be any positive-definite matrix.

Each richer family we optimize contains the scalar family as a special case, so after fitting it can never do worse. The questions are whether it does strictly better and whether the gain is large enough to matter. Unless stated otherwise, a gain from one family over another is reported as a percentage of the richer family’s net value. The nonlinear-over-linear gaps in Sections 4 and 7 instead use the best linear policy’s value as denominator, so their lower and upper endpoints refer to the same baseline.

Holding friction fixed. The cost term scales like $\kappa\|u\|^p$ while profit and risk scale like $\|u\|^2$, so at fixed κ a larger book is automatically a *less* frictional problem. Any decay across book size would then partly be friction quietly draining away. In the cross-sectional book-size and impact-model comparisons, unless stated otherwise, κ is recalibrated for every configuration so that the book retains approximately 60% of the profit it would earn with no trading costs at all. The perturbation calculation in Appendix A instead holds $\kappa = 5$ fixed while p approaches two.

3 Measuring what linearity costs

3.1 The dynamic program

The state is $z = (x, f)$ and it is convenient to treat the post-trade position $y = x + u$ as the control, since it becomes next period’s position. Maximizing the long-run average of Equation (1) is then the ergodic dynamic program: find a scalar g and a relative value function h with

$$g + h(x, f) = \max_y \left\{ y^\top Bf - \frac{\gamma}{2} y^\top \Sigma y - \frac{\kappa}{p} c(y - x) + \mathbb{E}[h(y, Af + \varepsilon)] \right\}. \quad (4)$$

Here g is the optimal average reward per period, the quantity every policy in this paper is scored against, and h is determined up to a constant. At $p = 2$ a quadratic h solves Equation (4)

exactly and the maximizer is linear (Gârleanu and Pedersen, 2013). For $p < 2$ no closed form is known, and Equation (5) below is the first-order condition of (4) with h taken quadratic.

Solving (4) on a grid works in principle but is too slow for the accuracy we need. Even with two assets, the position error falls only as the square of the grid spacing, so 0.05% accuracy takes about 280 points per axis. Under aggregate impact the outer power also couples the position axes, which removes the one-dimensional factorization that makes per-asset grids affordable. Appendix E.1 gives the convergence measurements. We therefore bound g from both sides instead of computing it.

3.2 Lower and upper bounds on the optimal value

Let J_{lin} be the best linear value and J^* the unknown value over all admissible policies. The opportunity cost of staying linear is

$$\Delta_{\text{lin}} = \frac{J^* - J_{\text{lin}}}{|J_{\text{lin}}|}.$$

Rather than compute J^* directly, we bound it. If an implementable policy has value J_{pol} and a dual calculation has value J_{dual} , then

$$J_{\text{pol}} \leq J^* \leq J_{\text{dual}}, \quad \frac{J_{\text{pol}} - J_{\text{lin}}}{|J_{\text{lin}}|} \leq \Delta_{\text{lin}} \leq \frac{J_{\text{dual}} - J_{\text{lin}}}{|J_{\text{lin}}|}.$$

The policy therefore gives an achieved lower bound and the dual gives an upper bound. A better-performing policy raises the floor; a tighter dual lowers the ceiling.

From below. If $J_{\text{lin}} = 100$ and one implementable nonlinear rule scores 100.6, then $J^* \geq 100.6$ and $\Delta_{\text{lin}} \geq 0.6\%$. This does not say that every nonlinear rule earns 0.6%; it says that at least one does. Our main lower-bound policy comes from the problem itself. Substitute the quadratic-surrogate continuation value $h(z) = -z^\top Pz + \text{constant}$ into the Bellman equation (4), write $y = x + u$, and differentiate with respect to u . Collecting terms gives

$$\kappa (u^\top \Lambda u)^{p/2-1} \Lambda u + Du = Mz, \quad D = \gamma \Sigma + 2P_{xx}, \quad (5)$$

where $z = (x, f)$, P is the Riccati matrix of the quadratic surrogate, and M collects the expected-return and continuation terms that are linear in z . If Du is dropped, Equation (5) solves to

$$u = a_0 w^\beta \Lambda^{-1} Mz, \quad w = (Mz)^\top \Lambda^{-1} (Mz). \quad (6)$$

The no- D solution has $\beta = (2 - p)/(2(p - 1))$ and $a_0 = \kappa^{-1/(p-1)}$. We keep its direction and power form but fit both numbers because the omitted Du term changes them. At $\beta = 0$ Equation (6) is exactly a linear rule, so the family *contains* the best fitted linear policy exactly, and any improvement is a strict gain. Only two numbers are fitted beyond the linear rule: a_0 sets overall trading speed and β makes the rule trade strong opportunities more aggressively. Under aggregate impact the multiplier rescales the whole basket; under per-asset impact it is applied name by name. We fit on synthetic AR(1) paths and evaluate on fresh, longer ones. An independent policy built from a fine per-asset grid is weaker in a separate per-asset robustness check (Appendix E.2). The aggregate lower bounds reported below come directly from Equation (6).

From above. The information-relaxation dual lets the controller see the whole future path of shocks, then charges a penalty for using information it should not have had (Brown et al., 2010). For any penalty that is a martingale difference, the resulting perfect-information value bounds the optimum from above. We take the penalty from the value function of the effective quadratic itself, so the approximation is asked to certify itself. The inner problem is strongly concave, and we add a computable residual correction on every path so that an early optimizer stop cannot invalidate the bound; the correction never exceeds 0.0002% of value (Appendix E.3).

3.3 Numerical validation of the bounds

Two checks confirm that the bounds reflect the problem rather than the numerics. First, at $p = 2$, where the linear rule is exactly optimal, the dual reports a gap of 1.4×10^{-8} percent and the nonlinear family finds no significant improvement: neither procedure manufactures a gain. Second, the nonlinear and linear policies are simulated on the same shocks and compared path by path. Across the nine cells the simulation error on the linear level is 0.048%–0.207%, but the standard error of the paired difference is only 0.001%–0.019%, so the lower-bound gains in Table 1 stand well clear of noise.

4 What linear rules cost

Table 1 and Figure 1 give the answer for two assets, as a percentage of the best linear policy’s value. The lower end is a nonlinear policy that actually beats the linear one; the upper end is the dual bound. After sharing the fitted linear surrogate, they use distinct lower- and upper-bound constructions, and the bracket holds in all nine cells.

ρ	p	at least	at most	ratio
0.00	1.70	0.159 ± 0.001	0.237	1.49
0.00	1.50	0.662 ± 0.003	0.978	1.48
0.00	1.30	2.144 ± 0.006	3.313	1.55
0.60	1.70	0.230 ± 0.002	0.362	1.57
0.60	1.50	0.905 ± 0.005	1.459	1.61
0.60	1.30	2.738 ± 0.011	4.696	1.72
0.85	1.70	0.428 ± 0.005	0.615	1.44
0.85	1.50	1.574 ± 0.010	2.386	1.52
0.85	1.30	4.387 ± 0.019	7.234	1.65

Table 1: Cost of restricting to linear rules, two assets (percent of best linear value).

The upper endpoint is 1.44 to 1.72 times the lower one in every cell, so the cost of linearity is pinned to within a small factor. Both ends rise as the cost becomes more curved (right to left in the figure) and as correlation rises.

At the square-root law $p = 1.5$, the bracket spans 0.662% to 2.386% across $\rho = 0, 0.6$ and 0.85. That is the same order as Brokmann et al. (2025) report for one asset, and correlation raises it: correlated assets concentrate the book in fewer independent directions, leaving less trading variation for a linear rule to average over.




Figure 1: Cost of linearity versus cost exponent p , two assets. Shaded band: lower to upper endpoint of Table 1; at $p = 2$ both ends are zero.

Raise curvature to $p = 1.3$ and the same three correlations give 2.144% to 7.234%. The lower p is a deliberately stronger-curvature stress. A lower p does not make trading cheap overall; it changes which trades are expensive. The marginal cost, $\kappa|u|^{p-1}$, rises more slowly with trade size, so large trades become relatively cheap and small trades relatively dear. The best response is to trade little on weak signals and hard on strong ones. A linear rule cannot do this: it scales every trade in the same proportion, so it is too busy when the signal is weak and too timid when the signal is strong. The lower p is, the more both mistakes cost. In the limit $p = 1$ the marginal cost no longer depends on size at all, which is the proportional-cost case where the optimal policy has a no-trade band (Liu, 2004), the opposite of a linear rule. Both ends roughly double from $\rho = 0$ to $\rho = 0.85$. The economic reading is that linearity costs most when the cost curve is most curved *and* the book is most concentrated.

For example, at $p = 1.5$ and $\rho = 0.85$, a linear value of 100 means that our implemented nonlinear rule reaches 101.574 while the certificate prevents any admissible rule from exceeding 102.386.

These percentages are of the strategy's *net* value—risk-adjusted profit after trading costs—not of assets under management, so they should be read against what the strategy earns. Suppose a \$1 billion book earns a net risk-adjusted \$50 million a year from its forecasts. At the square-root law with $\rho = 0.85$, the linear rule then leaves at least \$0.8 million and at most \$1.2 million a year on the table. At $p = 1.3$ the same book leaves \$2.2 to \$3.6 million a year. The loss scales with the strategy's net profit, so it grows in proportion for larger books or stronger signals, and it recurs every year the rule is in use. The two endpoints answer different questions. The lower endpoint is the gain achieved by our nonlinear rule in this model; the difference to the upper endpoint caps what any further policy improvement could add there—at most \$0.4 million a year in this example.

5 How nonlinear should the rule be?

The nonlinear rule behind our lower bounds is Equation (6) of Section 3.2. Its exponent β measures how nonlinear the rule is: $\beta = 0$ is a linear rule, and larger β makes trades grow faster than the signal. Fitting that rule gives a β that is consistently about half the exponent obtained when Du is omitted from the first-order condition. Retaining Du explains the difference.

Keeping Du in Equation (5), and specializing to the isotropic case, the trade size ℓ and the drive size m satisfy $a\ell^{p-1} + d\ell = m$. The elasticity of that relation gives, with $\beta^* = (2-p)/(2(p-1))$ the no- D answer,

$$\boxed{\beta_{\text{eff}} = \frac{(2-p)a}{2[(p-1)a + d]}} \quad \frac{\beta_{\text{eff}}}{\beta^*} = \frac{1}{1 + d/((p-1)a)}, \quad (7)$$

where $a = \kappa\lambda^{p/2}\ell^{p-2}$ is the stiffness of the trading cost at the typical trade size and d is the risk-plus-continuation curvature.

With no cost stiffness the rule is linear; with no risk curvature $\beta_{\text{eff}} = \beta^*$. The ratio $d/((p-1)a)$ governs the reduction from β^* .

This agrees with the continuous-time literature: Guasoni and Weber (2020) find, for one asset, that power-law impact slows trading near the target and speeds it up far away. Equation (7) confirms the pattern in discrete time with several assets and makes it explicit: one computable ratio governs how curved the rule should be.

The “factor of two” is therefore not a law; it holds when $d \approx (p-1)a$. Sweeping κ 400-fold moves $\beta_{\text{eff}}/\beta^*$ over a factor of six, and (7) tracks it without fitted parameters (Figure 2). The worst relative error is 14%, and one-sided: the prediction slightly overstates the exponent fitted across a distribution of trade sizes. Both approach β^* (dotted line) from below as cost stiffness rises. A desk can therefore estimate the rule’s curvature before running the policy optimization.




Figure 2: Policy exponent β against cost scale κ ($p = 1.5, \rho = 0$).

6 Diversification acts through independent trading directions

Under aggregate impact, diversification makes our nonlinear rule progressively less able to beat the best linear rule. Asset count is not the operative scale; what matters is how many independent *directions* you actually trade. Define

$$n_{\text{eff}} = \frac{(\text{tr}(\Sigma_u \Lambda))^2}{\text{tr}((\Sigma_u \Lambda)^2)}, \quad (8)$$

where Σ_u is the stationary covariance of trades. This participation ratio (Bell and Dean, 1970) is the reciprocal Herfindahl index of the trade spectrum. Lemma 1 gives it direct content here: n_{eff} is twice the inverse relative variance of the quantity that drives aggregate cost.

Lemma 1. *For $u \sim N(0, \Sigma_u)$ and any positive semi-definite Λ satisfying $\text{tr}(\Sigma_u \Lambda) > 0$, with $\chi = u^\top \Lambda u$,*

$$\frac{\text{Var}(\chi)}{\mathbb{E}[\chi]^2} = \frac{2}{n_{\text{eff}}}.$$

The two-line proof is in Appendix C. The lemma also supplies the mechanism. The marginal cost of trading asset i under (2) carries the factor $(u^\top \Lambda u)^{p/2-1}$, which depends on the *total*. A rule is linear when that factor can be treated as constant. Its relative fluctuation is of order $n_{\text{eff}}^{-1/2}$, and since the objective is stationary at its optimum, the resulting value loss should be of order $1/n_{\text{eff}}$.

Two entirely different ways of changing n_{eff} — adding assets, and tilting where the forecast points — put the measurements on one line, with slope -0.72 and $R^2 = 0.98$.

6.1 A long-short book can still be one trade

Holding both long and short positions does not by itself diversify trading. What matters above is how many independent combinations of trades the rule uses. A pairs trade that always buys one stock and sells the other in equal amounts, only scaling the size up or down, uses exactly one combination: the spread. Its trade covariance has rank one and $n_{\text{eff}} = 1$, the same as a book that only ever trades a single stock. Pairs trading therefore sits at the concentrated end of the result above, where the demonstrated nonlinear gain is largest.

This does not say that a long-short trade and a long-long trade cost the same. In the base case impact has no cross-effects ($\Lambda = \lambda I$): buying one stock does not move the other’s price, so trades of $(1, -1)$ and $(1, 1)$ pay the same impact. Correlation enters through risk instead. With $\rho > 0$ the long-long position carries the shared market risk and the long-short position hedges it, so the rule trades the two combinations in very different sizes. With cross-impact, buying one of two positively correlated stocks also pushes the other up, which makes selling it cheaper, so the long-short combination is also cheaper to trade. Section 9.1 adds that effect; the directional-calibration gain stays small because the scalar family $Q = q\Lambda$ already carries it.

One exact symmetry explains why two of our measured cases coincide. Trading the spread of two stocks with correlation ρ is the same problem as trading both stocks together with correlation $-\rho$: relabel “short stock 2” as a new asset and one becomes the other. Formally $\Sigma(-\rho) = D\Sigma(\rho)D$ with $D = \text{diag}(1, -1)$, and under symmetric liquidity and signals the two control problems map onto each other, so their values are identical after matching variance and friction. This is a mirror image, not a claim that long-short and long-long trades cost the same at the same correlation.

7 The impact model determines the measured diversification benefit

This is the paper’s central measured result. Under aggregate impact, the achieved gain of our nonlinear rule over the best linear rule falls with diversification; under per-asset impact, it does not fall over the tested range.

Section 6 documented the aggregate pattern. We now repeat the same measurement under the per-asset cost (3), which removes the common portfolio-flow term. Table 2 and Figure 3 report the result. Each entry is an achieved nonlinear-over-linear improvement at $p = 1.3$ with retention held fixed, hence a lower bound on the cost of restricting the policy to be linear. Portfolios and retention are matched; each policy is optimized under the cost model in force. The one-direction controls ($n_{\text{eff}} = 1$) agree to three decimals after friction recalibration.

case	n_{eff}	aggregate impact	per-asset impact
$n = 2, \rho = 0.9$, market direction	1.000	3.188	3.188
$n = 2, \rho = 0.9$, balanced directions	1.778	1.697	2.162
$n = 2, \rho = 0.9$, spread direction	1.000	3.213	3.213
$n = 2, \rho = 0.3$	1.713	2.131	3.060
$n = 3, \rho = 0.3$	2.564	1.637	2.988
$n = 5, \rho = 0.3$	4.415	1.139	2.947
$n = 10, \rho = 0.3$	9.212	0.609	2.932
slope against n_{eff}		−0.720	−0.019
R^2		0.979	0.013

Table 2: Achieved nonlinear-over-linear gain (percent of best linear value) at $p = 1.3$, matched retention.




Figure 3: Achieved nonlinear-over-linear gain versus n_{eff} under the two cost models ($p = 1.3$, matched retention).

Under per-asset impact the gain stays near 3% as the book expands: the nonlinear policy improves on the best linear rule by nearly the same amount at ten assets as at two. Under aggregate impact the improvement falls by a factor of five.

The mechanism is simple. Under (2) the marginal cost of each asset carries a factor depending on the total trade, and that total concentrates as the book spreads out, so the factor flattens and a linear approximation becomes more accurate. Under (3) the marginal cost of asset i is $\kappa\lambda_i^{p/2}|u_i|^{p-1}\text{sgn}(u_i)$, nonlinear in that asset’s own trade. There is no common random scale for diversification to smooth.

Stated as a single claim:

Whether diversification makes the best linear rule harder to beat depends on whether price impact responds to aggregate portfolio flow or to per-asset flow. The achieved nonlinear gain fades roughly as $1/n_{\text{eff}}$ under the first and does not fade over the tested range under the second.

Three features make this result economically consequential. First, the natural guess that opened Section 6, that a large book makes the choice of rule matter little, holds under only one of two established cost models, and the literature uses both (Schied et al., 2010; Mei et al., 2016). Second, the book that looks safest, a broad portfolio worked name by name, is exactly the one in which our nonlinear rule recovers as much at ten names as at two. Third, the difference is large: a factor of about five at ten names, and at least five times the largest gain from fine-tuning the cost coefficient anywhere in the paper (Sections 8 and 9). To our knowledge this is the first measurement of the effect in the correlated, finite-cost setting that Brokmann et al. (2025) leave open. It is a numerical result with a stated mechanism, not a theorem: Lemma 1 ties n_{eff} exactly to the relative variance that drives aggregate cost, Appendix A.1 proves an additional local per-asset condition at aggregate-zero points, and Section 11 states what remains open, including the exact decay rate.

The closest empirical evidence points both ways without settling the question for one manager’s own names. Bucci et al. (2020) find that the market reacts chiefly to the total net order flow of simultaneous metaorders, which is pooling before the nonlinearity, but across traders within one instrument. Benzaquen et al. (2017) find that liquidity providers price net inventory at the portfolio level, within a linear model. Hey et al. (2025) find that concave cross-impact fits metaorder data better than linear cross-impact, and Hey et al. (2024) show that the concavity must act along factor directions, thereby placing factor-pooled models between our two limits. Equations (2) and (3) are therefore benchmark limits, and the size of the gap between them is the reason the market’s actual pooling structure is worth identifying (Section 10).

8 Directional calibration is a second-order refinement

After choosing how impact pools across assets and whether trade size should be nonlinear, one may ask whether the stand-in quadratic needs a separate cost number for each kind of trade (the directional family of Section 2.3) rather than one for the whole book. This refinement is much smaller than the two choices already studied. Near $p = 2$, Appendix A proves

$$F_{\text{dir}} - F_{\text{scalar}} = C(2 - p)^2 + o((2 - p)^2), \quad C \geq 0,$$

and in the symmetric benchmark the gap is $c_0\rho^2(2-p)^2$ to leading order: two small factors multiplied. In that benchmark, the largest measured gain is 0.312% at $p = 1.3$ and 0.109% at the square-root law, so one well-fitted number captures nearly all of the value.

The appendix also proves a structural difference between the cost models. Under aggregate impact, one number is locally enough whenever the trade covariance is isotropic after liquidity whitening. At such an aggregate-zero point, per-asset impact generally adds a leverage-score condition; equal physical-name loading is sufficient, but not necessary (Appendix A.1). The stress test below probes that separation.

9 Stress test: a neutral heterogeneous factor book

The preceding sweeps use symmetric signals and equal correlations. Those choices isolate the mechanism, but they could also hide value in a directional cost matrix. We therefore repeat the one-number-versus-directional comparison in a deliberately asymmetric economy, keeping exact evaluation of linear policies.

The physical market has 12 assets whose risk follows a market, three sectors, value, momentum and idiosyncratic components. Volatilities range from 0.74 to 0.95, pairwise correlations from 0.22 to 0.75, liquidity coefficients from 0.45 to 2.47, signal strengths from 0.55 to 1.50 and persistence from 0.58 to 0.97, and the signal directions align with neither risk nor liquidity. The economy is held fixed while active directions are added, so a change with dimension cannot come from drawing a different market.

Positions and trades are dollar and market-beta neutral. With the dollar and beta exposures as the rows of C and an orthonormal basis N of their nullspace, physical positions are

$$x_{\text{phys}} = Nx, \quad CN = 0.$$

The reduced risk and aggregate-impact matrices are $N^\top \Sigma N$ and $N^\top \Lambda N$. Per-asset impact cannot be reduced this way; it stays $\sum_i \lambda_i^{p/2} |(Nu)_i|^p$ in physical names, and we keep it exactly. The economic difference is that under aggregate impact neutrality rewrites the same packaged trade in fewer coordinates, while under per-asset impact each name is charged before the constrained trades combine, so uneven loadings survive.

The directional family has one coefficient per active direction. Its basis is tailored to each model using the exact first cost correction at $p = 2$, and friction is calibrated to 60% retention. Where a full matrix is cheap enough to optimize, at two and three directions, it adds at most 0.00015 percentage point more.

Figure 4 gives two conclusions. First, one well-fitted cost number remains sufficient under aggregate impact after symmetry is relaxed. At $p = 1.3$ the directional gain is 0.019% to 0.078% across two to ten active directions, and 0.010% to 0.040% at $p = 1.5$, below the 0.312% two-asset peak. The common multiplier q absorbs changes of scale, such as the units of volatility or a stronger or slower signal, so separate coefficients earn value only from the residual shape of the trade covariance after liquidity whitening.

Second, the cost model still matters. Under per-asset impact the gain reaches 0.328% at $p = 1.3$ and 0.162% at $p = 1.5$. It is not monotone in the number of directions, because adding a factor changes which risk, liquidity and signal modes are active. Neutrality acts through concentration: with five active directions at $p = 1.3$, the aggregate gain is 0.024% unconstrained, 0.167% dollar neutral and 0.078% dollar and beta neutral, against per-asset values of 0.076%,




Figure 4: Directional-over-scalar gain in a dollar- and beta-neutral heterogeneous factor book, by cost model and number of active neutral directions.

0.274% and 0.328%. Holding κ fixed at its $p = 2$ calibration, instead of pinning retention, strengthens this conclusion: the largest aggregate gain falls to 0.059% while the largest per-asset gain rises to 0.446%, widening the ratio between the two models from about four to about seven.

9.1 Cross-impact

The stress test gives each name its own depth with no coupling. Cross-impact is itself a direction-dependent cost, so we add it directly:

$$\Lambda_{\text{phys}}(w) = D^{1/2}[(1-w)I + wR]D^{1/2}, \quad D = \text{diag}(\lambda),$$

with R the physical return correlation matrix and $w \in [0, 0.9]$. The diagonal stays λ at every w , so no per-name depth changes. Empirically, impact shares the eigenstructure of correlations rather than being proportional to them (Benzaquen et al., 2017; Mastromatteo et al., 2017; Tomas et al., 2022), so we read this path as a stylized rotation of impact toward the correlation modes, not a calibrated target. Every matrix on it is symmetric positive semidefinite, as absence of dynamic arbitrage requires in linear cross-impact models (Schneider and Lillo, 2019; Alfonsi et al., 2016; Tomas et al., 2022). We run an aligned arm, with R as generated by the factor book, and a deliberately misaligned arm with the names of R permuted. Only the aggregate cost is swept, since per-asset impact has no cross-impact by definition.

Neither arm changes the conclusion. At five active directions and $p = 1.3$ the gain moves from 0.078% at $w = 0$ to 0.089% aligned and 0.075% scrambled. At ten directions the scrambled arm nearly doubles the risk-liquidity misalignment, yet the gain falls from 0.0190% to 0.0163%. The null is informative: the one-number baseline is $q\Lambda$, not qI , so it already carries whatever structure cross-impact adds to Λ . What remains is direction dependence beyond Λ , which Proposition 2 ties to isotropy of the impact-whitened trade covariance rather than to how anisotropic Λ is.

The stress-test values exactly compare linear policies. Certification against the true optimum remains open because the dual of Section 3.2 is written for the aggregate cost, while the nonlinear lower-bound family does not survive the neutrality projection.

10 What this means on a desk

The relevant cost model is determined by execution. A portfolio trade or principal bid transfers a package to one intermediary and is close to package pricing (2); corporate-bond portfolio trades and equity blind bids are market examples (Meli and Todorova, 2024; Kavajecz and Keim, 2005). An algorithm working each name in its own order book is close to name-by-name pricing (3). Real desks use both. For name-by-name flow, adding names did not reduce the gain achieved by our nonlinear rule; for package flow, that gain fell roughly as $1/n_{\text{eff}}$. In these measurements, diversification makes the linear rule harder to beat only when the manager’s trades meet in the same liquidity pool.

That observed asymmetry ranks the demonstrated model risk. At ten names and $p = 1.3$, our nonlinear rule recovers at least 2.932% over the best linear rule under name-by-name pricing, versus 0.609% under package pricing (Table 2). A desk that assumes the milder package model may therefore miss the larger demonstrated gain; a desk that assumes name-by-name pricing may instead bear excess implementation complexity. Neutrality constraints sharpen the contrast: at an aggregate-zero point, per-asset impact generally adds a leverage-score condition (Appendix A.1).

The two models also make different, testable predictions about a desk’s own execution records. Under (3), name- i slippage should load on own-name size; under (2), it should also load on same-day portfolio flow. Regressing slippage on both, separately for each execution channel, would identify the relevant mixture. That regression is the empirical next step; this paper maps each specification to its consequence.

11 Scope and limits

What is certified. The two-asset costs of linearity in Table 1 are certified from both sides. The ten-name dichotomy and the stress test report achieved gains only: 0.609% and 2.932% are floors on what the optimum can gain, and the dichotomy is measured at the stronger curvature $p = 1.3$ rather than the square-root benchmark. Under aggregate impact the achieved gain is consistent with $6.1/n_{\text{eff}}$ up to about fifty directions, but the relaxed dual diagnostic decays more slowly, so the exact large-book exponent remains open (Appendix E.4). The diagnostics span n_{eff} from 1 to 99; rates beyond that are extrapolation.

What the fine-tuning result covers. The small gain from separate cost numbers is a robustness result over the tested sweeps, not a universal bound. The local theorem explains it near $p = 2$, and the stress test shows it survives unequal volatility, liquidity and signals, neutrality and cross-impact in one deliberately difficult construction. It concerns the best linear rule; whether the optimal nonlinear policy wants direction-dependent structure is open. Zero correlation is degenerate: uncorrelated assets without cross-impact separate into single-asset problems, the analogy Brokmann et al. (2025) draw with Liu (2004).

The economy. The evidence comes from synthetic economies designed to isolate the mechanism. The benchmark uses exchangeable assets, equal liquidity and symmetric AR(1) signals; the factor book relaxes those within a constructed 12-asset economy. The exchangeable structure also does real work in Section 6: a general covariance may not admit a single family indexed by n_{eff} . There is no bid–ask spread. Adding one gives a no-trade region, the classical proportional-cost problem and the $p = 1$ limit of Section 4, rather than the power cost studied here.

Static impact. Impact is temporary and static, with no decay. That is what keeps concave multi-asset impact well behaved: with decay, nonlinear impact can admit price manipulation (Gatheral, 2010; Gatheral et al., 2011; Schneider and Lillo, 2019), and Hey et al. (2024) show the concavity must then act along factor directions. Our total cost is a sum of nonnegative per-period terms, so no round trip can profit from its own impact, for any p and either cost model. That lets us compare the two forms on equal terms, including per-asset concavity acting name by name. The price is generality: a transient per-asset cost with cross-impact would have to be rebuilt in factor form.

Linear policies. Restricting to linear rules and certifying the restriction with a dual bound follows Moallemi and Sağlam (2017). The fine-tuning question is defined only for linear rules, which have a cost matrix Q to tune.

12 Future directions

The limits above define six concrete next steps.

1. *Identify the pooling structure in data.* The slippage regression of Section 10, run separately for algorithmic and risk-transfer execution, would estimate where each desk lies between the two regimes.
2. *Study the intermediate, factor-pooled cost.* The concave cross-impact model of Hey et al. (2024, 2025) pools trades within factors and sits between our two limits; where it falls on the diversification curve is open.
3. *Add transient impact.* This forces the per-asset cost into factor form (Section 11) and would test whether the dichotomy survives impact decay.
4. *Certify the large-book results.* The ten-name dichotomy has only achieved lower bounds, and the large-book decay rate is not pinned (Section 11). Richer dual penalties are admissible (Appendix E.3), but a certificate for the $1/n_{\text{eff}}$ rate needs a qualitatively new one.
5. *Look beyond linear rules.* The next policy question is whether the optimal nonlinear rule wants direction-dependent structure that the best linear rule does not.
6. *Calibrate to a real book.* Estimated covariances, cross-impact and a bid–ask component would replace the synthetic economy.

13 Conclusion

Six findings carry the paper. Each is a measured result with a stated mechanism unless marked as proved.

1. *The broad name-by-name book does not shrink the demonstrated nonlinear advantage.* Under package pricing, our nonlinear rule’s achieved gain falls from 2.131% at two names to 0.609% at ten ($p = 1.3$). Under name-by-name pricing it stays near 3%: 3.060% at two names and 2.932% at ten (Section 7).
2. *The demonstrated model risk is asymmetric.* At ten names our nonlinear rule recovers about five times more over the best linear rule under name-by-name pricing than under package pricing. Assuming the milder model may therefore miss the larger observed gain; the opposite choice mainly risks excess complexity (Section 10).
3. *A long-short pair is not diversified.* A pair that always buys one stock and sells the other in fixed proportion is really trading one thing, the spread. For trading costs it behaves like a single stock, where the demonstrated nonlinear gain is largest (Section 6.1).
4. *The cheaper big trades become, the worse a linear rule does.* At the square-root law the two-asset cost of linearity is 0.662% to 2.386%; at $p = 1.3$ it is 2.144% to 7.234%, about three times larger (Section 4). A local formula predicts how nonlinear the rule should be (Section 5).
5. *One well-fitted cost number captures nearly all directional-calibration value; the cost model dominates.* Separate cost numbers add at most 0.33% in the main runs and 0.45% anywhere measured. Choosing package versus name-by-name pricing moves the achieved gain by 2.3 points, at least five times more (Sections 8 and 9).
6. *At aggregate-zero points, per-asset impact generally adds a leverage condition* (proved locally). In our neutral test book the extra directional detail is worth nothing under package pricing and 0.106% under name-by-name pricing (Appendix A.1).

The practical ordering is therefore clear. First determine whether execution prices the portfolio as one package or works each asset through a separate liquidity pool. Then decide whether the curvature justifies a nonlinear trade-sizing rule. Only after those choices is it worth asking whether one effective cost coefficient should become several. The first question can be estimated from a desk’s own execution records (Section 10).

References

- Eduardo Abi Jaber, Eyal Neuman, and Sturmius Tuschmann. Optimal portfolio choice with cross-impact propagators. *Mathematical Finance*, 2026. doi: 10.1111/mafi.70025. Forthcoming. arXiv:2403.10273.
- Aurélien Alfonsi, Florian Klöck, and Alexander Schied. Multivariate transient price impact and matrix-valued positive definite functions. *Mathematics of Operations Research*, 41(3): 914–934, 2016. doi: 10.1287/moor.2015.0761.

- Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *The Journal of Risk*, 3(2):5–39, 2001. doi: 10.21314/JOR.2001.041.
- Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. Direct estimation of equity market impact. *Risk*, 18(7):57–62, 2005.
- Emmanuel Bacry, Adrian Iuga, Matthieu Lasnier, and Charles-Albert Lehalle. Market impacts and the life cycle of investors orders. *Market Microstructure and Liquidity*, 1(2):1550009, 2015. doi: 10.1142/S2382626615500094.
- Erhan Bayraktar, Thomas Cayé, and Ibrahim Ekren. Asymptotics for small nonlinear price impact: A PDE approach to the multidimensional case. *Mathematical Finance*, 31(1):36–108, 2021. doi: 10.1111/mafi.12283.
- Joseph J. Beaman. Non-linear quadratic Gaussian control. *International Journal of Control*, 39(2):343–361, 1984. doi: 10.1080/00207178408933170.
- R. J. Bell and P. Dean. Atomic vibrations in vitreous silica. *Discussions of the Faraday Society*, 50:55–61, 1970. doi: 10.1039/DF9705000055.
- Michael Benzaquen, Iacopo Mastromatteo, Zoltán Eisler, and Jean-Philippe Bouchaud. Dissecting cross-impact on stock markets: An empirical analysis. *Journal of Statistical Mechanics: Theory and Experiment*, 2017:023406, 2017. doi: 10.1088/1742-5468/aa53f7.
- J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research. Springer, 2000. doi: 10.1007/978-1-4612-1394-9.
- Richard C. Booton. Nonlinear control systems with random inputs. *IRE Transactions on Circuit Theory*, 1(1):9–18, 1954. doi: 10.1109/TCT.1954.6373354.
- Xavier Brokman, David Itkin, Johannes Muhle-Karbe, and Peter Schmidt. Tackling nonlinear price impact with linear strategies. *Mathematical Finance*, 35(2):422–440, 2025. doi: 10.1111/mafi.12449.
- David B. Brown and James E. Smith. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science*, 57(10):1752–1770, 2011. doi: 10.1287/mnsc.1110.1377.
- David B. Brown, James E. Smith, and Peng Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4):785–801, 2010. doi: 10.1287/opre.1090.0796.
- Frederic Bucci, Michael Benzaquen, Fabrizio Lillo, and Jean-Philippe Bouchaud. Crossover from linear to square-root market impact. *Physical Review Letters*, 122:108302, 2019. doi: 10.1103/PhysRevLett.122.108302.
- Frederic Bucci, Iacopo Mastromatteo, Zoltán Eisler, Fabrizio Lillo, Jean-Philippe Bouchaud, and Charles-Albert Lehalle. Co-impact: Crowding effects in institutional trading activity. *Quantitative Finance*, 20(2):193–205, 2020. doi: 10.1080/14697688.2019.1660398.

- Bruno Durin. Signal-wise performance attribution for constrained portfolio optimisation. Working paper, Capital Fund Management, 2014. URL <https://arxiv.org/abs/1404.4798>. Revised August 2014.
- James B. Farison and Moufid S. Elhazzouri. Control of non-linear stochastic sampled-data systems with non-quadratic criteria by equivalent LQ modelling. *International Journal of Control*, 45(6):2035–2044, 1987. doi: 10.1080/00207178708933865.
- Nicolae Gârleanu and Lasse Heje Pedersen. Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, 68(6):2309–2340, 2013. doi: 10.1111/jofi.12080.
- Nicolae Gârleanu and Lasse Heje Pedersen. Dynamic portfolio choice with frictions. *Journal of Economic Theory*, 165:487–516, 2016. doi: 10.1016/j.jet.2016.06.001.
- Jim Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759, 2010. doi: 10.1080/14697680903373692.
- Jim Gatheral, Alexander Schied, and Alla Slynko. Exponential resilience and decay of market impact. In Frédéric Abergel, Bikas K. Chakrabarti, Anirban Chakraborti, and Manipushpak Mitra, editors, *Econophysics of Order-Driven Markets*, pages 225–236. Springer Milan, 2011. doi: 10.1007/978-88-470-1766-5_15.
- Paolo Guasoni and Marko H. Weber. Nonlinear price impact and portfolio choice. *Mathematical Finance*, 30(2):341–376, 2020. doi: 10.1111/mafi.12234.
- Martin B. Haugh and Chun Wang. Dynamic portfolio execution and information relaxations. *SIAM Journal on Financial Mathematics*, 5(1):316–359, 2014. doi: 10.1137/120896761.
- G. Heess. Application of state space statistical linearization to optimal stochastic control of non-linear systems. *International Journal of Control*, 11(4):697–701, 1970. doi: 10.1080/00207177008905950.
- Natascha Hey, Iacopo Mastromatteo, and Johannes Muhle-Karbe. Concave cross impact. *SSRN Working Paper*, 2024. doi: 10.2139/ssrn.5046242.
- Natascha Hey, Eyal Neuman, and Sturmius Tuschmann. Nonparametric estimation of self- and cross-impact. *arXiv preprint*, 2025. arXiv:2510.06879.
- David C. Hoaglin and Roy E. Welsch. The hat matrix in regression and ANOVA. *The American Statistician*, 32(1):17–22, 1978. doi: 10.1080/00031305.1978.10479237.
- Kenneth A. Kavajecz and Donald B. Keim. Packaging liquidity: Blind auctions and transaction efficiencies. *Journal of Financial and Quantitative Analysis*, 40(3):465–492, 2005. doi: 10.1017/S0022109000001836.
- I. E. Kazakov. Approximate probabilistic analysis of the accuracy of operation of essentially nonlinear systems. *Avtomatika i Telemekhanika*, 17(5):385–409, 1956. In Russian. English translation: *Automation and Remote Control* 17, 423–450, 1956.
- Hong Liu. Optimal consumption and investment with transaction costs and multiple risky assets. *The Journal of Finance*, 59(1):289–338, 2004. doi: 10.1111/j.1540-6261.2004.00634.x.

- Iacopo Mastromatteo, Michael Benzaquen, Zoltán Eisler, and Jean-Philippe Bouchaud. Trading lightly: Cross-impact and optimal portfolio execution. *Risk*, 2017. arXiv:1702.03838.
- Xiaoling Mei, Victor DeMiguel, and Francisco J. Nogales. Multiperiod portfolio optimization with multiple risky assets and general transaction costs. *Journal of Banking & Finance*, 69: 108–120, 2016. doi: 10.1016/j.jbankfin.2016.04.002.
- Jeffrey Meli and Zornitsa Todorova. Portfolio trading in corporate bond markets. *AFA 2024 Annual Meeting Paper*, 2024.
- Ciamac C. Moallemi and Mehmet Sağlam. Dynamic portfolio choice with linear rebalancing rules. *Journal of Financial and Quantitative Analysis*, 52(3):1247–1278, 2017. doi: 10.1017/S0022109017000345.
- Ludovic Moreau, Johannes Muhle-Karbe, and H. Mete Soner. Trading with small price impact. *Mathematical Finance*, 27(2):350–400, 2017. doi: 10.1111/mafi.12098.
- Johannes Muhle-Karbe and Chase Tracy. Stochastic liquidity as a proxy for concave multi-asset propagator models. *SSRN Working Paper*, 2024. doi: 10.2139/ssrn.4885733.
- Robert M. Price. A useful theorem for nonlinear devices having Gaussian inputs. *IRE Transactions on Information Theory*, 4(2):69–72, 1958. doi: 10.1109/TIT.1958.1057444.
- Alexander Schied, Torsten Schöneborn, and Michael Tehranchi. Optimal basket liquidation for CARA investors is deterministic. *Applied Mathematical Finance*, 17(6):471–489, 2010. doi: 10.1080/13504860903565050.
- Michael Schneider and Fabrizio Lillo. Cross-impact and no-dynamic-arbitrage. *Quantitative Finance*, 19(1):137–154, 2019. doi: 10.1080/14697688.2018.1467033.
- Michael Tomas, Iacopo Mastromatteo, and Michael Benzaquen. How to build a cross-impact model from first principles: Theoretical requirements and empirical results. *Quantitative Finance*, 22(6):1017–1036, 2022. doi: 10.1080/14697688.2021.2020328.
- Bence Tóth, Yves Lempérière, Cyril Deremble, Joachim de Lataillade, Julien Kockelkoren, and Jean-Philippe Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1:021006, 2011. doi: 10.1103/PhysRevX.1.021006.
- Kazuo Yoshida. A method of optimal control of non-linear stochastic systems with non-quadratic criteria. *International Journal of Control*, 39(2):279–291, 1984. doi: 10.1080/00207178408933165.
- Elia Zarinelli, Michele Treccani, J. Doyne Farmer, and Fabrizio Lillo. Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity*, 1(2):1550004, 2015. doi: 10.1142/S2382626615500042.

A Directional calibration: derivations and zero conditions

This appendix records the local expansion and zero-condition split summarized in Section 8. In the measured symmetric sweep the directional-over-scalar gain peaks at 0.312% of net value at $p = 1.3$ and 0.109% at the square-root law $p = 1.5$; Section 11 states the scope of those numbers.

Let $\varepsilon = 2 - p$, and let $\eta \in \mathbb{R}^m$ be a smooth local chart for the rich effective-matrix family around the quadratic optimum $Q_0 = \kappa\Lambda$, represented by η_0 . Write $F(\eta, \varepsilon)$ for the true objective of the linear policy constructed from η . The scalar family $q\Lambda$ is a curve through η_0 ; denote its nonzero tangent by $s \in \mathbb{R}^m$. For the two-mode diagonal chart used in the symmetric benchmark,

$$\eta = \theta = (\log q_1, \log q_2), \quad \eta_0 = (\log(\kappa\lambda), \log(\kappa\lambda)), \quad s = e = (1, 1)^\top.$$

For a full 2×2 symmetric matrix, $m = 3$ and s is the tangent of the curve representing $q\Lambda$ in the chosen chart.

Lemma 2 (Local regularity). *Suppose the stabilizing Riccati solution exists at $Q_0 = \kappa\Lambda$, its closed loop is strictly stable, and its control curvature is positive definite. Then the Riccati solution and its feedback gain are smooth functions of Q near Q_0 . The stationary covariance is also smooth there. If the quadratic-case trade covariance and Λ are positive definite, then the stationary objective is smooth from the side $p < 2$ near $(p, Q) = (2, Q_0)$, and $F(\eta, \varepsilon)$ is smooth near $(\eta_0, 0)$.*

The lemma is a standard implicit-function argument, but stating it matters: the second-order conclusion below does not follow from equality at $p = 2$ alone. Appendix B gives the argument.

Proposition 1 (Local scalar-directional gap). *Under Lemma 2, suppose η_0 is an isolated interior optimum and*

$$H = -F_{\eta\eta}(\eta_0, 0)$$

is positive definite. Let F_{dir} and F_{scalar} denote the values on the local maximizing branches that converge to η_0 . Set $b = F_{\eta\varepsilon}(\eta_0, 0)$. Then

$$F_{\text{dir}}(\varepsilon) - F_{\text{scalar}}(\varepsilon) = C\varepsilon^2 + o(\varepsilon^2),$$

where

$$C = \frac{1}{2} \left[b^\top H^{-1} b - \frac{(s^\top b)^2}{s^\top H s} \right] \geq 0. \quad (9)$$

Moreover, $C = 0$ exactly when the first-order optimal adjustment $H^{-1}b$ lies on the scalar tangent $\text{span}(s)$; otherwise $C > 0$.

Equation (9) is a projection formula. The directional family can move in every coordinate of its local chart; the scalar family can move only along s . The coefficient is half the squared distance, measured in the value-curvature metric H , between the first-order optimum and its projection onto that line. This is standard perturbation geometry (Bonnans and Shapiro, 2000). The portfolio-specific content is identifying the direction being projected.

Let $U := \Sigma_u$ be the quadratic-case covariance of trades and $\chi = u^\top \Lambda u$. Gaussian covariance differentiation (Price, 1958) gives the self-consistent effective matrix below. The reason it is

self-consistent is the classical equivalent-LQ argument: at a local optimum over linear policies, replace the expected nonlinear cost by its tangent plane as a function of U . The first-order conditions do not change, and that tangent plane is the expected cost of a quadratic with $Q = Q_{\text{eff}}(p, U)$. Conversely, a smooth solution of $Q = Q_{\text{eff}}(p, U(Q))$ is stationary for the true objective because the nonlinear cost and surrogate quadratic have identical first variations. Near $p = 2$, the implicit-function theorem and $H \succ 0$ identify this stationary branch with the local maximizer.

$$Q_{\text{eff}}(p, U) = \kappa \mathbb{E} \left[\chi^{p/2-1} \Lambda + (p-2) \chi^{p/2-2} \Lambda u u^\top \Lambda \right].$$

At $p = 2$ this is $\kappa \Lambda$ for every U . Therefore changes in U do not feed back at first order, and

$$\left. \frac{\partial Q_{\text{eff}}}{\partial p} \right|_{p=2} = \kappa \mathbb{E} \left[\frac{1}{2} \log(\chi) \Lambda + \frac{\Lambda u u^\top \Lambda}{\chi} \right]. \quad (10)$$

The first term in (10) is a common rescaling, which the scalar family can represent. The second is the source of directional structure.

Proposition 2 (Invariant zero and positive conditions). *In a two-asset model with arbitrary positive-definite risk and impact matrices, let*

$$\Gamma = \Lambda^{1/2} U \Lambda^{1/2}$$

be the quadratic-case trade covariance in impact units, and let the rich family contain all symmetric perturbations of Q_0 . Under the local identifiability and positive-Hessian conditions of Proposition 1, the coefficient C is zero if and only if Γ is proportional to the identity. If Γ has at least two unequal active eigenvalues, then $C > 0$.

This proposition removes equal volatility from the criterion. Raw covariance is not the invariant object: changing the units of an asset changes it. Impact-whitened covariance Γ compares trade variation with the liquidity metric and is unchanged by that bookkeeping. The proof uses the fact that

$$\mathbb{E} \left[\frac{v v^\top}{v^\top v} \right]$$

is proportional to the identity exactly when the covariance of the nondegenerate Gaussian vector v is proportional to the identity.

For a structured directional family rather than a full matrix, the same conclusion holds when its tangent space contains the anisotropic part of (10). Simultaneous diagonalization of Λ , Γ and the family basis is a sufficient condition and holds in our symmetric benchmark. Without that alignment, the general projection formula (9), rather than isotropy alone, decides whether the structured family improves on the scalar one.

For two assets, suppose the impact-whitened mode variances are $s_1, s_2 > 0$. An elementary angular integral gives

$$\mathbb{E} \left[\frac{v v^\top}{v^\top v} \right] = \text{diag} \left(\frac{\sqrt{s_1}}{\sqrt{s_1} + \sqrt{s_2}}, \frac{\sqrt{s_2}}{\sqrt{s_1} + \sqrt{s_2}} \right).$$

Consequently,

$$\left. \frac{d(\theta_1 - \theta_2)}{d\varepsilon} \right|_{\varepsilon=0} = - \frac{\sqrt{s_1} - \sqrt{s_2}}{\sqrt{s_1} + \sqrt{s_2}}. \quad (11)$$

If $H = \text{diag}(h_1, h_2)$ in the decoupled risk basis, then

$$C = \frac{1}{2} \frac{h_1 h_2}{h_1 + h_2} \left(\frac{\sqrt{s_1} - \sqrt{s_2}}{\sqrt{s_1} + \sqrt{s_2}} \right)^2. \quad (12)$$

Thus both ingredients are necessary: unequal impact-whitened trade scales and positive value curvature in both active modes.

Proposition 3 (Two small parameters). *Suppose the symmetric two-mode model and its local maximizing branches are jointly four-times continuously differentiable in (ρ, ε) near $(0, 0)$, with the nondegeneracy conditions above holding uniformly there. Then*

$$F_{\text{dir}} - F_{\text{scalar}} = \rho^2(2 - p)^2(c_0 + o(1)), \quad (13)$$

where $c_0 \geq 0$, and $c_0 > 0$ if the first directional adjustment changes nontrivially with ρ at zero.

Indeed, the model at $-\rho$ is the model at $+\rho$ with its market and spread labels exchanged, so the value gap is even in ρ . It vanishes for all p at $\rho = 0$ by symmetry, and it vanishes to second order in ε at $p = 2$ by Proposition 1. Joint smoothness therefore factors both ρ^2 and ε^2 from the gap. This is why the gain is especially small near the common benchmark: correlation creates a direction-specific parameter change at first order, but value is stationary at the shared quadratic optimum, so the value gain is its square.

The calculation uses no fitted coefficient from $p < 2$. At $\rho = 0.85$, it predicts the derivative in (11) as -0.616756 ; fitting the full objective at $2 - p = 0.005$ gives -0.617302 . The ratio of the actual to predicted raw value gap is 1.039, 1.015, 1.008 and 1.004 as $2 - p$ falls from 0.05 to 0.005. Near $\rho = 0$, $C(\rho)/\rho^2$ is stable at about 0.00487. The centered finite difference used to estimate the mixed derivative evaluates the smooth Gaussian objective at a nearby $p > 2$ as well as $p < 2$. The economic model is stated for $p \leq 2$; the extension is used only as a numerical derivative check.

The invariant statement was checked separately without equal volatility. In a two-asset market with volatilities 0.7 and 1.4, unequal non-diagonal liquidity, signal persistence 0.84 and 0.94, and non-diagonal signal loading, (10) predicts a full-matrix change whose finite quotient at $2 - p = 0.005$ agrees to 0.56% in relative Frobenius norm. In a second market with raw volatilities 0.5 and 2.0 but $\Lambda^{1/2}U\Lambda^{1/2} \propto I$, the full-matrix gain at $p = 1.7$ is 1.8×10^{-14} percent. Unequal raw volatility therefore does not by itself create directional value; impact-whitened anisotropy does.

That second market is not an isolated example, which matters because it also happens to have diagonal risk, diagonal impact and a symmetric signal. Relabeling holdings by an invertible D sends (Σ, Λ, B) to $(D^\top \Sigma D, D^\top \Lambda D, D^\top B)$ and carries Γ to itself, so applying random D to the uncorrelated benchmark generates markets whose risk, impact and signal loading are all non-diagonal and heterogeneous while Γ stays proportional to the identity. Across ten such markets, with raw volatility ratios up to 3.1 and $\text{cond}(\Lambda)$ up to 77, the largest full-matrix gain is 2.3×10^{-13} percent. The ε^2 rate also survives away from symmetry: holding a random heterogeneous market fixed and varying only p , so that H and Γ cannot move, the gain falls with log-log slopes of 2.09, 2.14, 2.18 and 2.24 against the predicted 2.

The theorem explains the local regime. At $p = 1.3$, $\varepsilon = 0.7$ is already far from that regime, so the measured gains need not follow the expansion quantitatively. Against full optimizations at $\rho = 0.85$ the ratio of actual to predicted gap is 1.04 at $\varepsilon = 0.05$, then 1.27, 1.50 and 1.78 at

$p = 1.7, 1.5$ and 1.3 . The expansion under-predicts throughout, so using it to argue that the directional gain is small is conservative: the measured gaps are larger than the local theory alone would suggest, and still small. Nor does it cover inactive modes or a singular Hessian. Proposition 2 allows heterogeneous volatility and liquidity, but only through the local full-matrix comparison; it does not say that every restricted directional parameterization captures the required anisotropic correction.

A.1 Per-asset leverage at aggregate-zero points

Sections 7 and 9 report that per-asset impact rewards a directional metric several times more than aggregate impact does. That is a measured dichotomy. The local theorem turns it into a structural one, because the two costs behave differently under an exact neutrality reduction $x_{\text{phys}} = Lx$.

Aggregate impact is *closed* under the reduction: $(u_{\text{phys}}^\top \Lambda_{\text{phys}} u_{\text{phys}})^{p/2}$ equals $(u^\top \Lambda_F u)^{p/2}$ with $\Lambda_F = L^\top \Lambda_{\text{phys}} L$, so everything above applies verbatim with $(\Sigma, \Lambda, B, U, n)$ replaced by $(\Sigma_F, \Lambda_F, B_F, U_F, d)$. The criterion is also independent of the nullspace basis: under $L \mapsto LR$ the whitened covariance Γ_F maps to $O\Gamma_F O^\top$ for an orthogonal O , so its eigenvalues do not move.

Per-asset impact is not closed. A constraint mixes physical names and a p -norm is not rotation invariant, so the cost remains $\sum_i \lambda_i^{p/2} |a_i^\top u|^p$ with $a_i^\top$ the rows of L .

Proposition 4 (Per-asset leverage condition). *Let $s_i = a_i^\top U_F a_i$ and $g_i = \lambda_i^{1/2} \Lambda_F^{-1/2} a_i$, so that $\sum_i g_i g_i^\top = I$ and $\lambda_i s_i = g_i^\top \Gamma_F g_i$. Define the leverage scores*

$$\ell_i = \|g_i\|^2 = \lambda_i a_i^\top \Lambda_F^{-1} a_i, \quad \sum_i \ell_i = d.$$

Then the first directional correction of the per-asset effective matrix is, modulo a multiple of Λ_F , the whitened matrix $\frac{\kappa}{2} \sum_i \log(\lambda_i s_i) g_i g_i^\top$. If the aggregate zero condition $\Gamma_F = cI$ holds, this reduces to $\sum_i \log(c \ell_i) g_i g_i^\top$, which is proportional to the identity when the ℓ_i are equal and generically not otherwise.

The ℓ_i are statistical leverage scores, the hat-matrix diagonals of the weighted design whose rows are $\lambda_i^{1/2} a_i^\top$, with $\sum_i \ell_i = d$ the usual trace identity (Hoaglin and Welsch, 1978). Our contribution is to use them as a per-name diagnostic for whether an additional per-asset correction appears at an aggregate-zero point. The closest construction we know of in the execution literature is Mastromatteo et al. (2017), who also reduce a multi-asset impact cost to a weighted sum over projections of the trade; their basis is a spectral rotation of the impact matrix rather than the nullspace of a neutrality constraint, and they derive no per-row condition.

Thus, at an aggregate-zero point $\Gamma_F \propto I$, per-asset impact generally imposes an additional condition: the leverage-weighted log-score matrix $\sum_i \log(\ell_i) g_i g_i^\top$ must be proportional to the identity for its first directional correction to vanish. Equal leverage scores are a simple sufficient condition, not a necessary one; unequal scores can cancel in specially arranged redundant frames. Take the impact-whitened isotropic two-mode economy, so $\Gamma_F \propto I$ exactly, and realize it with three unit-depth physical assets through $L = W \Lambda_F^{1/2}$ with W an orthonormal nullspace basis. Then $L^\top L = \Lambda_F$ for *any* constraint, so every arm is the same aggregate market, while $\ell_i = \|w_i\|^2$ is set by which constraint was imposed.

arm	leverage ℓ_i	aggregate gain	per-asset gain, $p = 1.3$
beta neutral, unequal betas	0.858, 0.710, 0.432	0.000000%	0.105669%
dollar neutral, symmetric	0.667, 0.667, 0.667	0.000000%	0.000000%

Both arms sit on the aggregate zero condition, with Γ eigenvalues equal to ten digits, and the aggregate gain is zero in both at $p = 1.7, 1.5$ and 1.3 . Of these two test arms, the per-asset gain is zero only in the one with equal leverage scores. At $p = 2$ the two costs give identical values to machine precision.

The leverage statement is exact *at* the aggregate zero condition, which is what makes it a clean separation of the two criteria. Away from that point both sources of anisotropy are active and interact, and leverage spread alone does not order the measured gaps: in the factor book of Section 9 the three-direction case has the widest leverage spread and the smallest per-asset gain. The proposition identifies a specific term rather than ranking all markets away from the zero condition.

B Proofs for the directional-calibration results

Proof of Lemma 2

Write the augmented linear dynamics as $z_{t+1} = \mathcal{A}z_t + \mathcal{B}u_t + w_{t+1}$. The Riccati equation is a smooth finite-dimensional equation in its solution P and in Q . At the stabilizing solution, its derivative with respect to P is the discrete Lyapunov operator

$$\mathcal{L}_{\Phi_0}(X) = X - \Phi_0^\top X \Phi_0, \quad \Phi_0 = \mathcal{A} + \mathcal{B}K_0.$$

Strict stability makes this operator invertible, with inverse

$$\mathcal{L}_{\Phi_0}^{-1}(Y) = \sum_{j=0}^{\infty} (\Phi_0^\top)^j Y \Phi_0^j.$$

The implicit function theorem therefore makes $P(Q)$, and hence $K(Q)$, smooth near Q_0 . This argument also covers the cross term between state and control in our LQ cost: the derivative of the minimizing control drops out by the envelope theorem.

For a stable K , the stationary covariance solves

$$\Psi - \Phi_K \Psi \Phi_K^\top = W.$$

The same Lyapunov argument makes $\Psi(K)$ smooth. Finally, if $U = K\Psi K^\top$ and Λ are positive definite, then $\chi = u^\top \Lambda u > 0$ almost surely. Powers of χ times any fixed power of $|\log \chi|$ are integrable for p in a neighborhood to the left of two. Dominated differentiation under the Gaussian integral makes the expected power cost, and therefore F , smooth there. $\square$

Why the Hessian condition is checkable. Let $\mathcal{R}_0$ denote the positive-definite control curvature $R + \mathcal{B}^\top P_0 \mathcal{B}$ of the quadratic problem. The standard LQ performance-difference identity gives, near K_0 ,

$$F(\eta_0, 0) - F(\eta, 0) = \text{tr} \left(\mathcal{R}_0 [K(\eta) - K_0] \Psi_{K(\eta)} [K(\eta) - K_0]^\top \right).$$

Consequently, for a parameter direction a and $\dot{K} = D_\eta K(\eta_0)[a]$,

$$a^\top H a = 2 \operatorname{tr}(\mathcal{R}_0 \dot{K} \Psi_0 \dot{K}^\top).$$

Thus H is positive definite whenever Ψ_0 is positive definite and the policy Jacobian is injective on the cost directions being compared. These are the local identifiability conditions stated in the propositions.

In the symmetric benchmark the risk modes decouple at $p = 2$. For a mode with return variance $s_i > 0$ and effective cost $q_i > 0$, its inventory block of the Riccati matrix is the positive solution of

$$2P_i^2 + \gamma s_i P_i - \frac{1}{2} q_i \gamma s_i = 0, \quad P_i = \frac{-\gamma s_i + \sqrt{\gamma^2 s_i^2 + 4q_i \gamma s_i}}{4}.$$

The coefficient mapping incoming inventory into the trade is

$$K_{x,i} = -\frac{P_i + \gamma s_i/2}{P_i + \gamma s_i/2 + q_i/2}.$$

Direct differentiation shows $\partial K_{x,i}/\partial q_i \neq 0$ for finite positive q_i, s_i . Therefore each log-cost coordinate changes its own policy block, and the policy Jacobian is injective. Positive signal innovation variance and nonzero signal loading make the stationary state covariance positive definite in each mode. Hence $H \succ 0$ in the stated symmetric benchmark rather than merely being observed numerically.

Proof of Proposition 1

Put $\delta = \eta - \eta_0$. Under the stated smoothness assumptions, Taylor expansion around $(\eta_0, 0)$ gives

$$F(\eta_0 + \delta, \varepsilon) = F_0 + \text{terms independent of } \delta + \varepsilon b^\top \delta - \frac{1}{2} \delta^\top H \delta + o(\varepsilon^2 + \|\delta\|^2).$$

Positive definiteness of H and the implicit function theorem give

$$\delta_{\text{dir}} = \varepsilon H^{-1} b + O(\varepsilon^2).$$

On the scalar tangent $\delta = as$, maximizing the same quadratic gives

$$\delta_{\text{scalar}} = \varepsilon s \frac{s^\top b}{s^\top H s} + O(\varepsilon^2).$$

Substitution and cancellation of the terms common to both families yields (9). Cauchy–Schwarz in the H inner product makes the coefficient nonnegative. Equality holds exactly when $H^{-1}b$ belongs to $\operatorname{span}(s)$.

To identify b , first note what the equivalent matrix establishes. For a stationary linear policy, the nonlinear cost depends on the policy through its trade covariance U . Replacing that cost by its tangent plane at a local optimum preserves its first-order variation with respect to the policy. The tangent plane is $\frac{1}{2} \operatorname{tr}(Q_{\text{eff}} U)$ plus a constant. Therefore a smooth solution of $Q = Q_{\text{eff}}(p, U(Q))$ is stationary for the true objective. Local uniqueness and $H \succ 0$ identify this stationary branch with the local maximizer near $p = 2$. This establishes matching first-order conditions; it is not a claim of global optimality.

Now apply Price's theorem to $f(u) = \chi^{p/2}$ with $\chi = u^\top \Lambda u$. Since

$$\nabla_u^2 f = p \chi^{p/2-1} \Lambda + p(p-2) \chi^{p/2-2} \Lambda u u^\top \Lambda,$$

and the cost is $\kappa \mathbb{E}[f]/p$, twice its covariance gradient is the displayed $Q_{\text{eff}}(p, U)$. At $p = 2$ this matrix is independent of U , so the derivative of the self-consistent trade covariance contributes nothing at first order. Differentiating in p gives (10).

To obtain (11), diagonalise the impact-whitened trade covariance and write a standard two-dimensional Gaussian in polar coordinates. Its radius cancels from $vv^\top/(v^\top v)$, leaving

$$\frac{1}{2\pi} \int_0^{2\pi} \frac{s_1 \cos^2 t}{s_1 \cos^2 t + s_2 \sin^2 t} dt = \frac{\sqrt{s_1}}{\sqrt{s_1} + \sqrt{s_2}},$$

and the second diagonal entry follows because the trace is one. Equation (12) is the two-dimensional H -projection in closed form. $\square$

Proof of Proposition 2

Set $v = \Lambda^{1/2}u$ and diagonalise $\Gamma = \text{Cov}(v)$, with eigenvalues $g_1, \dots, g_n > 0$. Sign symmetry makes

$$D(\Gamma) := \mathbb{E} \left[\frac{vv^\top}{v^\top v} \right]$$

diagonal in the same basis. Its i th diagonal entry is

$$d_i = \mathbb{E} \left[\frac{g_i Z_i^2}{\sum_k g_k Z_k^2} \right].$$

If $g_i > g_j$, condition on all squared Gaussian coordinates except Z_i^2 and Z_j^2 , and call their contribution $R \geq 0$. Pair each realization $(Z_i^2, Z_j^2) = (a, b)$ with (b, a) . After putting the two fractions over a common positive denominator, the numerator of their contribution to $d_i - d_j$ is

$$(g_i - g_j) [R(a + b) + 2(g_i + g_j)ab] > 0$$

almost surely. Hence $g_i > g_j$ implies $d_i > d_j$. It follows that $D(\Gamma)$ is proportional to the identity if and only if Γ is proportional to the identity.

Equation (10) can be written

$$\partial_p Q_{\text{eff}}|_{p=2} = \kappa \Lambda^{1/2} \left[\frac{1}{2} \mathbb{E} \log(v^\top v) I + D(\Gamma) \right] \Lambda^{1/2}.$$

Therefore its first directional component is absent exactly when $\Gamma \propto I$. For the full symmetric-matrix family, the projection criterion in Proposition 1 then gives $C = 0$ under isotropy and $C > 0$ otherwise. For a smaller directional family, the same argument applies when its tangent contains the anisotropic part of (10); simultaneous diagonalization is sufficient. $\square$

Proof of Proposition 3

Let $\Delta(\rho, \varepsilon)$ denote the difference between the two local optimized values. Exchanging ρ and $-\rho$ swaps the market and spread labels, so Δ is even in ρ . At $\rho = 0$, exchangeability and local

uniqueness make the directional optimum scalar for every small ε ; hence $\Delta(0, \varepsilon) = 0$. Joint smoothness and evenness therefore give

$$\Delta(\rho, \varepsilon) = \rho^2 G(\rho, \varepsilon)$$

for a smooth G .

At $\varepsilon = 0$, both families contain the exact quadratic optimum, and Proposition 1 also gives $\partial_\varepsilon \Delta(\rho, 0) = 0$. Thus $G(\rho, 0) = \partial_\varepsilon G(\rho, 0) = 0$, including at $\rho = 0$ by continuity. A second application of the same factorization gives

$$G(\rho, \varepsilon) = \varepsilon^2 \tilde{G}(\rho, \varepsilon).$$

Taking $c_0 = \tilde{G}(0, 0)$ proves (13). Nonnegativity follows from nesting. It is strict when the first directional adjustment has a nonzero derivative in ρ and the Hessian is positive definite. $\square$

Proof of Proposition 4

Let $a_i^\top$ denote row i of the feasible map L , and put $s_i = a_i^\top U_F a_i > 0$. For a centered Gaussian trade,

$$\mathbb{E}|a_i^\top u|^p = m_p s_i^{p/2}, \quad m_p = 2^{p/2} \frac{\Gamma((p+1)/2)}{\sqrt{\pi}}.$$

The expected per-asset cost is therefore

$$\frac{\kappa m_p}{p} \sum_i \lambda_i^{p/2} s_i^{p/2}.$$

Taking its covariance gradient and matching it to the gradient of $\frac{1}{2} \text{tr}(Q_{\text{eff,sep}} U_F)$ gives

$$Q_{\text{eff,sep}}(p, U_F) = \kappa m_p \sum_i \lambda_i^{p/2} s_i^{p/2-1} a_i a_i^\top.$$

At $p = 2$, $m_2 = 1$ and

$$Q_{\text{eff,sep}}(2, U_F) = \kappa \sum_i \lambda_i a_i a_i^\top = \kappa \Lambda_F,$$

independently of U_F . Thus, as in (10), the derivative of U_F does not feed back at first order. Differentiating with respect to p at two gives a common scalar multiple of Λ_F from m_p , plus the directional term

$$\frac{\kappa}{2} \sum_i \lambda_i \log(\lambda_i s_i) a_i a_i^\top.$$

Whitening by $\Lambda_F^{-1/2}$ and setting $g_i = \lambda_i^{1/2} \Lambda_F^{-1/2} a_i$ turns this term into

$$\frac{\kappa}{2} \sum_i \log(\lambda_i s_i) g_i g_i^\top.$$

Because $\Lambda_F = \sum_i \lambda_i a_i a_i^\top$, $\sum_i g_i g_i^\top = I$. Moreover,

$$g_i^\top \Gamma_F g_i = \lambda_i a_i^\top U_F a_i = \lambda_i s_i.$$

Under the aggregate zero condition $\Gamma_F = cI$, this becomes $\lambda_i s_i = c \|g_i\|^2 = c \ell_i$, proving the stated form. If all ℓ_i are equal, the correction is $\log(c\ell) \sum_i g_i g_i^\top = \log(c\ell) I$ and therefore has no directional part. Unequal scores generically leave a non-scalar weighted sum; the qualifier is necessary because specially arranged rows can cancel. $\square$

C Proof of Lemma 1

For $u \sim N(0, \Sigma_u)$ and symmetric Λ , the standard Gaussian trace identities give

$$\mathbb{E}[u^\top \Lambda u] = \text{tr}(\Sigma_u \Lambda), \quad \text{Var}(u^\top \Lambda u) = 2 \text{tr}((\Sigma_u \Lambda)^2).$$

Dividing and substituting (8) gives $\text{Var}(\chi)/\mathbb{E}[\chi]^2 = 2/n_{\text{eff}}$. $\square$

Checked numerically against Monte Carlo at $n = 2, 5, 20$: 1.7341 versus 1.7326, 1.0411 versus 1.0398, 0.2932 versus 0.2929.

D Evaluating linear policies without simulation

Linear policies are scored without Monte Carlo noise, using deterministic quadrature, which is what makes the comparisons in Table 1 precise. Under $u = Kz$ the state is a Gaussian AR(1), so its stationary covariance solves a discrete Lyapunov equation and the profit and risk terms are traces. The cost term needs $\mathbb{E}[(u^\top \Lambda u)^{p/2}]$ for Gaussian u , which we compute two ways as a cross-check.

For up to three assets, factor the expectation into a radial part (a chi-square moment) and an angular integral over the sphere, evaluated by quadrature. For any number of assets, use the identity $x^s = \frac{s}{\Gamma(1-s)} \int_0^\infty t^{-s-1} (1 - e^{-tx}) dt$ for $s \in (0, 1)$. Taking expectations turns the hard object into the Laplace transform of a weighted chi-square, which is a product:

$$\mathbb{E}[q^s] = \frac{s}{\Gamma(1-s)} \int_0^\infty t^{-s-1} \left[1 - \prod_i (1 + 2t\mu_i)^{-1/2} \right] dt, \quad s = p/2,$$

with μ_i the eigenvalues of $\Sigma_u \Lambda$. This is a single integral in any dimension. Substituting $t = e^v$ makes the tail behavior explicit. We evaluate the middle interval by deterministic quadrature and add analytically integrated leading asymptotic corrections to both omitted tails. Near $p = 2$ the lower tail decays like $e^{(1-p/2)v}$, and simply truncating it loses over ten percent of the integral while appearing to converge.

The test suite compares the two deterministic routes through three dimensions and checks the general-dimensional route independently against Monte Carlo in ten dimensions.

E Numerical details for the bounds

This appendix collects the numerical evidence behind Section 3 and the large-book rate quoted in Section 11.

E.1 Grid dynamic programming

We discretize the signal with a Rouwenhorst chain, put positions on a box, restrict y to the same grid as x so the position dynamics need no interpolation, and run relative value iteration. Against the exact $p = 2$ optimum the position error falls as the square of the grid spacing:

$$9.71\%, 4.33\%, 2.41\%, 1.08\%, 0.58\%, 0.39\%, 0.27\%, 0.15\%, 0.096\%$$

at 21 through 201 points per axis. A global fit gives error $\propto h^{2.008}$ across a tenfold refinement, and a 50% wider box at the same spacing matches the value to eight digits, so spacing rather

than truncation is what remains. Reaching 0.05% from the position grid alone takes about 280 points per axis. Signal discretization adds a separate error: at $p = 2$, moving from seven to nine or more signal states shifts the value by about 0.03 percentage points, and convergence is slower below $p = 2$.

The binding obstacle is where the curvature sits. When trading cost is a sum over axes, the Bellman maximization factors into one-dimensional passes and costs $O(n_x^3)$ per sweep rather than $O(n_x^4)$. The per-asset cost (3) is such a sum at every p ; the aggregate cost (2) is one only at $p = 2$, because below it the outer power couples the position axes. The choice of cost model is therefore also a choice of computational complexity, and adaptive grids or parametric value functions would still have to confront that coupling.

E.2 A per-asset robustness check on the lower-bound policy

For the per-asset cost, we check whether Equation (6) leaves an easy gain behind by building a second policy without it. Per-asset separability makes a fine grid affordable: we solve the discretized dynamic program on a ladder to 161 points per axis, take one off-lattice policy-improvement step, and simulate the result in the continuous model.

It is weaker in all nine cells below $p = 2$. At $\rho = 0.6$ the grid policy earns 0.321%, 1.021% and 2.254% over the best linear rule at $p = 1.7, 1.5, 1.3$, against 0.393%, 1.252% and 2.777% for Equation (6). The ordering is the same at $\rho = 0$ and $\rho = 0.85$, with margins of 0.05 to 0.71 percentage points, well outside the Monte Carlo standard errors. Moving the grid policy's actions off the lattice adds only 0.04–0.10 points and reverses no margin. At $p = 2$ the same construction returns -0.006% to -0.010% . The check supports the two-parameter family as the stronger implemented per-asset policy; it neither checks the aggregate bracket nor shows that the family is globally optimal.

E.3 Certifying the dual inner problem

An under-solved inner maximum is not an upper bound (Haugh and Wang, 2014), so we make the numerical step certifying. If $\hat{y}$ is the optimizer's candidate and μ is the smallest eigenvalue of $\gamma\Sigma$, μ -strong concavity gives, for every y ,

$$f(y) \leq f(\hat{y}) + \nabla f(\hat{y})^\top (y - \hat{y}) - \frac{\mu}{2} \|y - \hat{y}\|^2, \quad \text{hence} \quad \max_y f(y) \leq f(\hat{y}) + \frac{\|\nabla f(\hat{y})\|^2}{2\mu}.$$

We add the correction $\|\nabla f(\hat{y})\|^2/(2\mu)$ on every simulated path, so a line-search termination cannot invalidate the bound. The correction is in the units of the per-period objective g . Across the nine cells of Table 1 (200 paths of 400 periods each), its largest per-period value is 1.2×10^{-6} , at $\rho = 0.6$, $p = 1.3$, where the best linear policy earns 0.851 per period: 0.00014% of the value, below 0.0002% and too small to change any displayed upper endpoint.

Richer penalties are admissible. The penalty's dependence on the control is $h(y, f') - \mathbb{E}[h(y, Af + \varepsilon)]$, so concavity requires only that what survives be affine in y : $\pi_t = y_t^\top \Phi(f_t, \varepsilon_t)$ with $\mathbb{E}[\Phi \mid \mathcal{F}_t] = 0$. A quadratic value function forces Φ linear in ε . A cubic extension improves the relaxed value modestly but does not change its rate in the number of assets.

E.4 The large-book rate

Under aggregate impact we pushed the book to a hundred assets. If the achieved gain of our nonlinear policy decays exactly as $1/n_{\text{eff}}$, multiplying it by n_{eff} gives a flat line:

n	2	3	5	10	20	35	50	75	100
n_{eff}	1.71	2.56	4.39	9.20	19.09	34.04	49.02	74.01	99.01
achieved $\times n_{\text{eff}}$	3.61	—	4.94	5.65	5.88	—	6.04	—	—
relaxed diagnostic $\times n_{\text{eff}}$	5.9	7.0	7.8	9.3	10.6	11.8	12.5	14.4	15.9

The achieved product approaches about 6.1 over the largest reported dimensions, consistent with $6.1/n_{\text{eff}}$. The relaxed product does not plateau; a power fit gives about $17/n_{\text{eff}}^{0.78}$. This older run predates the residual correction of Appendix E.3, so we use it as a rate diagnostic, not as a certified upper bound.