Dynamic Arbitrage from Price-Based Risk Constraints

Francesco Nicolai*
BI Norwegian Business School

Simona Risteska†
Warwick Business School

March 2026 – Latest Version
January 2025 – First Version

Abstract

Under classic no-manipulation conditions on market impact, price-based risk constraints (margins, haircuts, leverage limits, volatility targets, mandates) can still generate dynamic arbitrage. We develop a refined no-dynamic-arbitrage test for such environments; it requires only the constraint rule and an estimate of market impact. The test also yields an upper bound on the size of the constrained sector consistent with non-manipulability. We apply it to volatility-managed portfolios: admissible scale is well below one day of average daily volume and below 1% of the size of the underlying market, while vulnerability rises sharply once linked notional reaches roughly one to two days of daily volume, or about 1% to 2% of underlying market size. Manipulation incentives are strongest in low-volatility states, driven by feedback between measured risk and rule-induced trading.

Keywords: price manipulation, dynamic arbitrage, market impact, initial margin, collateral haircuts, VaR constraints, risk-control overlays, forced liquidation, procyclicality, volatility targeting.
JEL Classification: G14, G12, G11, G18, G28.

*Department of Finance, BI Norwegian Business School, Nydalsveien 37, 0484, Oslo, Norway, email: francesco.nicolai@bi.no.

†Department of Finance, Warwick Business School, University Of Warwick, Scarman Rd, Coventry CV4 7AL, UK, email: simona.risteska@wbs.ac.uk

1 Introduction

A price-based constraint maps recent transaction prices into future binding demand through margin requirements, leverage limits, collateral haircuts, or mandated exposure. Prices therefore clear trades today and, at the same time, determine tomorrow's feasible positions. When the constrained sector is large enough to move markets,[1] a strategic trader can profit by moving the sampled statistic that the constraint uses and then unwinding against the predictable rebalancing it induces.

How large can the constrained sector become before its own rule-induced trading makes the market manipulable? We provide an implementable test to answer this question. We first show that classical no-manipulation restrictions are incomplete once sampled prices mechanically determine future binding demand. A market can satisfy the Huberman and Stanzl (2004) condition that every admissible round trip has nonnegative expected execution cost and still admit profitable round trips, because current trades now affect future constraints and therefore future forced flow. Theorem 1 in Section 4 derives the missing no-dynamic-arbitrage restriction for this environment and turns it into a capacity test for any disclosed price-based rule. Given a market-impact estimate and the mapping from sampled transaction prices into next-period margin requirements, haircuts, leverage limits, or target exposures, the test computes the largest size of the constrained sector consistent with non-manipulability. In practice, it is an eigenvalue test on an augmented impact matrix that incorporates the endogenous trading generated by the constrained sector.

Section 7 applies the test to volatility-managed indices used in structured products and indexed annuities. We show that admissible scale is definitely small relative to market liquidity. Using only a market-impact estimate and the published volatility-control rule, we compute a size bound for the linked sector and vulnerability curves that show, for each sector size, how often the no-manipulation condition fails. We do this for a single-asset rule (Template A) and a portfolio rule (Template B), allowing for different assumptions on liquidity, target volatility, implementation lags, volatility dynamics, cross-impact, and return correlations. With a manipulation horizon of $T = 126$ days, the conservative capacity bound $W_{\max}(\mathcal{Z})$ is 0.165 days of ADV in Template A and 0.125, 0.122, and 0.230 in Template B for portfolios of $N = 2, 4, 8$ assets. At the benchmark scale of one day of volume, the no-manipulation condition fails 66.8% of the time with Template A and 88.8%, 94.0%, and 78.8% with Template B.[2] Figure 1 plots this object directly: for each value of W, it shows the fraction of

[1] Evidence that formulaic mandates and rule-triggered collateralization have grown to a material share of modern markets includes: indexation in U.S. long-term mutual funds and ETFs (Investment Company Institute, 2025) and broader discussions of passive investing (Sushko and Turner, 2018); central-bank estimates putting assets in volatility-targeting and related volatility-sensitive strategies as high as USD 2 trillion globally (European Central Bank, 2020); the expansion of indexed annuities and the associated rise of custom indices with built-in volatility-control features (Meisenzahl et al., 2025; American Academy of Actuaries, 2026a); and the magnitude and stress sensitivity of derivatives margining (BIS, 2022; IOSCO, 2022; ISDA, 2025). Public descriptions of the underlying maps from prices into target exposures appear in index-provider methodologies (MSCI, 2021) and in prospectus-level term sheets and filings for risk-controlled indices embedded in structured products (J.P. Morgan Chase, 2023; Goldman Sachs, 2025).

[2] Volatility-managed products typically allocate between a risky ETF and a low-risk asset such as cash, T-bills, or a bond index. The ADV and portfolio weight reported here refer to the risky leg. Suppose the risky leg is implemented through a liquid ETF with \$5 billion ADV and \$600 billion in net assets. If the risky allocation averages 0.80 of AUM, then $W_{\max}(\mathcal{Z}) = 0.165$ implies a risky notional of $0.165 \times (\$5 \text{ billion}) = \$0.825$ billion and total linked AUM of $\$0.825/0.80 = \$1.03$ billion across the volatility-managed complex. This corresponds to only $\$0.825/\$600 = 0.14\%$ of the underlying ETF on the risky leg, and $\$1.03/\$600 = 0.17\%$ in total linked AUM relative to the ETF's size. At $W = 1$, which is about 0.83% of the size of the underlying ETF, the risky notional equals one ADV (\$5 billion) and total linked AUM is $\$5/0.80 = \$6.25$ billion; at that scale, the test fails in 66.8% of stress states, so the market is not manipulation-proof. These

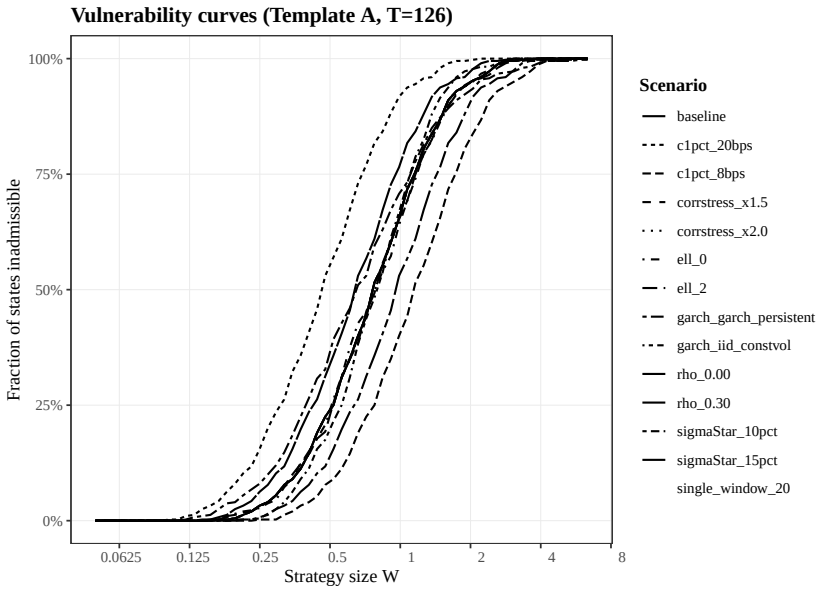

Figure 1 Vulnerability curve for a volatility-controlled rule over horizon $T = 126$ days.
Each stress state z is a joint draw of the rule inputs (recent returns that determine the realized-volatility estimate and regime) and a local liquidity environment (impact operator $\mathcal{I}$). For each z we compute the statewise capacity $W_{\max}(z) = \sup\{W \geq 0 : \lambda_{\min}^{\mathcal{R}_T}(\widehat{H}(z; W)) \geq 0\}$, the largest sector scale W that passes the restricted-eigenvalue screen over horizon T. The figure plots the vulnerability curve $\Pr_z[W > W_{\max}(z)]$ against W, where W is the sector's linked notional expressed in ADV units (so $W = 1$ means linked notional equals one day of average trading volume; a unit change in exposure trades one ADV). This is the $T = 126$ panel of Figure 2.

states in which the test fails.

We characterize the optimal attack in Section 5. The attacker trades to move the sampled risk statistic that governs future constraints and then profits by unwinding against the predictable rebalancing induced by the tighter constraint. The strength of this motive is governed by two objects: how sensitive the sampled statistic is to current prices, and how strongly constrained demand responds to that statistic. For strategies based on convex risk measures, including volatility-control rules, the trigger is often inventory-light because the attacker can raise measured risk with offsetting trades inside the sampling window rather than by carrying a large directional position. Because the risk measure is computed over a rolling window, the induced tightening and its reversal are mechanically timed, which creates a predictable opportunity to enter and exit. This mechanism differs from corners, squeezes, and standard price-level manipulation. The attacker does not profit by impairing the net worth of the constrained sector, but by moving a disclosed statistic that mechanically reallocates demand over time. We also study the multi-asset case in Section 6. Optimal exploitation can be reduced to the two-dimensional span of the portfolio that is the most efficient at moving risk and the portfolio that lies in the direction of liquidation.

figures are consistent with Khomyn et al. (2024) who find that the median ETF's secondary-market dollar ADV is about 0.9% of AUM per day. In this setting, profitable round trips can be arbitrarily small. The attack is also inventory-light: the trigger can be implemented through a short-lived spike-and-revert strategy that raises measured volatility while largely undoing the price-level change, allowing the attacker to unwind quickly without warehousing a large directional position.

Notably, the timing of mechanical trading rules matters. Lagged updating yields a well-posed induced response, whereas instantaneous updating can be singular or nearly singular, generating extreme amplification. For volatility-managed portfolios, this distinction corresponds to whether the exposure applied at t is computed from a volatility estimate that excludes the price changes created by the t-rebalance (lagged updating) or instead responds to prices within the same sampling window (instantaneous updating). Without a lag, the rule can approach a fixed-point singularity: arbitrarily small price changes can imply arbitrarily large mandated rebalancing. Lags and smoothing therefore act as market stability conditions and as devices to reduce procyclicality.

Another implication of mandated strategies based on convex risk measures is that fragility need not be greatest in high-volatility states. It can be greatest in calm markets, because the risk input is then easiest to move. For target-volatility rules, manipulability is highest when measured volatility is low: inverse-volatility scaling makes exposure most sensitive to the volatility estimate, and a given increase in realized variance has the largest percentage effect on measured volatility when the baseline is low. Calm markets are therefore precisely the states in which a small spike-and-revert can move measured volatility enough to induce material and predictable rebalancing. This is a rule-based version of the paradox of financial stability (Borio and Drehmann, 2009). It is related to, but distinct from, leverage-cycle mechanisms (Brunnermeier and Pedersen, 2009; Geanakoplos, 2010; Adrian and Shin, 2010). In leverage-cycle models, amplification operates through balance-sheet capacity: requirements tighten because collateral values fall or funding conditions worsen. Here amplification arises because requirements are a deterministic function of a sampled risk statistic computed from recent prices, and that statistic can be moved by trading even without any change in fundamentals, collateral values, or intermediary capital.

1.1 Contributions to the Literature

Our first contribution is to extend the classical no-dynamic-arbitrage benchmark to environments in which transaction prices mechanically determine future binding demand. In the benchmark of Huberman and Stanzl (2004) and the subsequent no-dynamic-arbitrage literature (Gatheral, 2010; Schneider and Lillo, 2019), trades move execution prices but do not alter tomorrow's constraints or tomorrow's order flow. With price-based constraints, that exogeneity breaks down: current transaction prices enter next-period requirements and therefore generate predictable rebalancing by the constrained sector. We show that classical impact restrictions are no longer sufficient for market viability in this environment. Theorem 1 in Section 4 derives the missing admissibility condition by augmenting the impact-cost matrix to internalize rule-induced flow. Dynamic arbitrage is present exactly when the resulting matrix fails to be positive semidefinite on round trips, which yields a tractable eigenvalue test for manipulation-proofness.

Our second contribution is to turn this admissibility restriction into an implementable capacity diagnostic. Given a disclosed price-based rule and a market-impact calibration, we compute a conservative bound on the scale of the constrained sector that is consistent with non-manipulability, together with vulnerability curves that map sector size into the probability that the market fails the test. In that sense, the paper is close in spirit to work that transforms observable inputs into supervision-relevant objects such as systemic-risk measures, capital shortfalls, and vulnerability di-

agnostics (Acharya et al., 2012; Adrian and Brunnermeier, 2016; Acharya et al., 2017; Brownlees and Engle, 2017; Greenwood et al., 2015). The diagnostic requires only the rules followed by the constrained sector and an estimate of impact, and therefore provides a practical screen for how large a mechanically constrained sector can become before its own induced trading creates exploitable round trips.

Our third contribution is to characterize optimal strategic trading when forced flow is both predictable and manipulable. Unlike predatory-trading and order-anticipation models, where the relevant imbalance is typically taken as exogenous from the attacker's perspective (Brunnermeier and Pedersen, 2005; Attari et al., 2005; Carlin et al., 2007; Rostek and Weretka, 2015; Sannikov and Skrzypacz, 2016; Lou et al., 2013; Fardeau, 2021), the imbalance here is endogenously generated by a disclosed rule that maps sampled prices into future requirements. Mechanism-wise, the paper is closest to contract-based and benchmark manipulation, where trading moves a reference statistic that then governs payoffs or behavior (Kumar and Seppi, 1992; Hillion and Suominen, 2004; Dutt and Harris, 2005; Onur and Reiffen, 2018; Zhang, 2022). Our setting replaces settlement or marking rules with a risk-based constraint that deterministically generates future forced flow, and it delivers a characterization of the optimal trigger-and-reverse round trip, including the inventory-light nature of attacks under convex risk measures and the two-portfolio reduction in the multi-asset case.

Finally, this paper is closely related to Nicolai (2026), which establishes an impossibility result for price-based risk constraints and studies the mechanism design of such rules. That paper shows that no price-based, risk-sensitive rule can simultaneously guarantee liquidity continuity and round-trip manipulation-proofness, and it characterizes the resulting design responses. We take the institutional environment and the mechanical trading rule as given and ask a different question: for a given rule and a given impact calibration, when does the induced feedback admit dynamic arbitrage, and how large can the constrained sector become before that happens? This yields a computable admissibility test, conservative capacity bounds, and vulnerability curves.

Roadmap. Section 2 sets up the baseline impact environment and reviews the classical no-manipulation condition, in both the single- and multi-asset cases. Section 3 introduces deterministic price-based constraints and closes the feedback loop, and Section 4 derives the corresponding augmented no-arbitrage condition (Theorem 1). Sections 5 and 6 solve the optimal attack in the single- and multi-asset cases and establish the two-portfolio reduction in the multi-asset setting (Theorem 2). Section 7 implements the stress test for volatility-controlled indices, and Section 8 concludes.

2 Model Setup

2.1 Single Asset

Fix a finite horizon T with trading dates $t = 0, 1, \dots, T-1$. Random variables are defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq -1}, \mathbb{P})$. The sigma-field $\mathcal{F}_t$ represents public information available immediately after outcomes at date t are realized, with $\mathcal{F}_{-1}$ denoting initial information before the first trade. A time-t trade u_t is therefore chosen using $\mathcal{F}_{t-1}$.

Definition 1 (Inventory dynamics and admissible strategies)**.** *A single-asset strategy is a predictable*

process $u = (u_t)_{t=0}^{T-1}$ *with* $u_t \in \mathbb{R}$*, meaning* u_t *is* $\mathcal{F}_{t-1}$*-measurable for each* $t = 0, \dots, T-1$*. We interpret* $u_t > 0$ *as a purchase and* $u_t < 0$ *as a sale. Inventory evolves according to*

$$x_{t+1} = x_t + u_t, \qquad x_0 = 0. \tag{1}$$

A strategy is admissible *if it is predictable and uniformly bounded: there exists* $\bar{u} > 0$ *such that* $|u_t| \le \bar{u}$ *for all* t *and all* $\omega \in \Omega$*.*

Definition 2 (Round trips)**.** *An admissible strategy* u *is a round trip if it starts and ends flat,*

$$x_T = 0, \tag{2}$$

Equivalently, $\sum_{t=0}^{T-1} u_t = 0$

2.1.1 Prices

We distinguish the unaffected price S_t, which evolves absent the strategy's trades, from the execution price P_t, at which the strategy transacts and which incorporates market impact. Let $(S_t)_{t=0}^{T-1}$ be an exogenous process adapted to $(\mathcal{F}_t)_{t \ge -1}$.

Assumption 1 (Unaffected price)**.** $(S_t)_{t=0}^{T-1}$ *is an* $\mathcal{F}_t$*-martingale and* $\mathbb{E}[|S_t|] < \infty$ *for all* t*.*

Assumption 1 rules out predictable drift in S_t. Any expected profits must therefore come from market impact and from the deterministic institutional feedback introduced below, not from forecasting S_t. Execution prices typically move against the trader: buys execute above S_t and sells below S_t. We capture this wedge with a standard transient linear impact specification.

Assumption 2 (Transient linear impact)**.** *Fix coefficients* $G_0, G_1, \dots, G_{T-1} \in \mathbb{R}$ *and an impact slope* $\eta > 0$*. Given a strategy* u*, the execution price at time* t *is*

$$P_t(u) = S_t + \eta \sum_{s=0}^{t} G_{t-s} u_s. \tag{3}$$

The term $\eta G_{t-s} u_s$ is the contribution of the time-s trade to the time-t execution price. Previous flows are allowed to influence current prices. The sequence (G_ℓ) controls persistence: small G_ℓ at large lags means impact decays quickly. A useful benchmark is purely temporary impact, $G_0 = 1$ and $G_\ell = 0$ for all $\ell \ge 1$, in which case $P_t(u) = S_t + \eta u_t$.

2.1.2 Execution cost and finite-time dynamic arbitrage

Trading at execution prices generates the realized cash outflow

$$\sum_{t=0}^{T-1} u_t P_t(u).$$

For a round trip this is the entire payoff, since the strategy starts and ends flat. Define realized trading profit by

$$\Pi(u) = -\sum_{t=0}^{T-1} u_t P_t(u).$$

Under the martingale benchmark for the unaffected price, a round trip cannot generate expected profits from S_t alone.

Lemma 1. *Under Assumption 1, for any admissible round trip u,*

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t S_t\right] = 0. \tag{4}$$

Proof. Let x_t denote inventory. Since $u_t = x_{t+1} - x_t$,

$$\begin{aligned}\sum_{t=0}^{T-1} u_t S_t &= \sum_{t=0}^{T-1} (x_{t+1} - x_t) S_t \\ &= x_T S_{T-1} - x_0 S_0 + \sum_{t=1}^{T-1} x_t (S_{t-1} - S_t).\end{aligned} \tag{5}$$

For a round trip, $x_0 = 0$ and $x_T = 0$ a.s., so the boundary terms vanish. Predictability implies x_t is $\mathcal{F}_{t-1}$-measurable. Taking expectations and using $\mathbb{E}[S_t \mid \mathcal{F}_{t-1}] = S_{t-1}$ yields, for each $t \geq 1$,

$$\mathbb{E}[x_t (S_{t-1} - S_t)] = \mathbb{E}[x_t \, \mathbb{E}[S_{t-1} - S_t \mid \mathcal{F}_{t-1}]] = 0.$$

Summing over t gives the claim. □

We start from the classical finite-horizon notion of price manipulation in Huberman and Stanzl (2004) (and its transient-impact refinements in Gatheral (2010)). Throughout we impose this no-manipulation benchmark, so any profitable round trip identified below cannot be attributed to the classic mechanism in which a trader moves prices through impact and then unwinds. The gains we study instead run through the institutional feedback that maps transacted prices into binding requirements and induced forced flow.[3]

Definition 3 (Finite-time dynamic arbitrage)**.** *Under Assumptions 1–2, an admissible round trip u is a price manipulation strategy if its expected execution cost is negative:*

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t(u)\right] < 0. \tag{6}$$

Equivalently, $\mathbb{E}[\Pi(u)] > 0$ for $\Pi(u) = -\sum_{t=0}^{T-1} u_t P_t(u)$. The impact model is manipulation free (no finite-time dynamic arbitrage) if no admissible round trip is a price manipulation strategy.

[3] Attention is restricted to admissible strategies (predictable and uniformly bounded) as in Definition 1. This excludes doubling-type schemes based on unbounded positions or arbitrarily large interim losses.

2.1.3 A discrete-time no-manipulation condition

Stack trades into $u = (u_0, \dots, u_{T-1})^\top \in \mathbb{R}^T$. Define the lower-triangular impact matrix $K \in \mathbb{R}^{T\times T}$ with elements

$$K_{t,s} = G_{t-s}\,\mathbf{1}_{\{t\geq s\}}, \qquad t, s \in \{0, 1, \dots, T-1\}, \tag{7}$$

so that, for each t,

$$(Ku)_t = \sum_{s=0}^{t} G_{t-s} u_s.$$

Under Assumption 2, the execution price is $P_t(u) = S_t + \eta(Ku)_t$, and total execution cost is

$$\sum_{t=0}^{T-1} u_t P_t(u) = \sum_{t=0}^{T-1} u_t S_t + \eta\, u^\top K u.$$

For a round trip, Lemma 1 implies $\mathbb{E}[\sum_t u_t S_t] = 0$, so expected cost is determined by the impact term. It is convenient to define the following symmetric matrix

$$H = \eta\,(K + K^\top). \tag{8}$$

Since $u^\top K u$ is a scalar, $u^\top K u = u^\top K^\top u$, hence

$$\eta\, u^\top K u = \frac{1}{2}\, u^\top H u.$$

Finally, consider the subspace of deterministic round trips. Our notion of dynamic arbitrage restricts the attention to these trades:

$$\mathcal{R}_T = \left\{ u \in \mathbb{R}^T : \sum_{t=0}^{T-1} u_t = 0 \right\}. \tag{9}$$

We now derive the classical no-manipulation criterion (Huberman and Stanzl, 2004) on price impact, which states that H should be positive semidefinite to prevent dynamic arbitrage.

Proposition 1. *Suppose Assumptions 1–2 hold. Then, for any admissible round trip u,*

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t(u)\right] = \frac{1}{2}\,\mathbb{E}\left[u^\top H u\right]. \tag{10}$$

Consequently, the following are equivalent:

(i) The impact model is manipulation free in the sense of Definition 3.

(ii) The symmetric matrix H is positive semidefinite on $\mathcal{R}_T$, that is, $u^\top H u \geq 0$ for all $u \in \mathcal{R}_T$.

Proof. Fix an admissible round trip u. By Lemma 1,

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t(u)\right] = \mathbb{E}\left[\sum_{t=0}^{T-1} u_t \big(P_t(u) - S_t\big)\right].$$

Assumption 2 gives $P_t(u) - S_t = \eta(Ku)_t$, so

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t(u)\right] = \eta\,\mathbb{E}[u^\top K u] = \frac{1}{2}\,\mathbb{E}\Big[u^\top \eta(K + K^\top)u\Big] = \frac{1}{2}\,\mathbb{E}[u^\top H u],$$

which proves (10).

(ii)⇒(i): If $u^\top H u \geq 0$ for all $u \in \mathcal{R}_T$, then for any admissible round trip the realized trade vector satisfies $u(\omega) \in \mathcal{R}_T$ a.s., hence $u(\omega)^\top H u(\omega) \geq 0$ a.s. Taking expectations and using (10) yields $\mathbb{E}[\sum_t u_t P_t(u)] \geq 0$, so no admissible round trip is a manipulation strategy.

(i)⇒(ii): If H is not positive semidefinite on $\mathcal{R}_T$, then there exists $\tilde{u} \in \mathcal{R}_T$ with $\tilde{u}^\top H \tilde{u} < 0$. Consider the deterministic strategy $u = \tilde{u}$. It is a round trip by construction, and it is admissible for any bound $\bar{u} \geq \max_t |\tilde{u}_t|$. Then (10) gives

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t(u)\right] = \frac{1}{2}\,\tilde{u}^\top H \tilde{u} < 0,$$

so u is a manipulation strategy, contradicting (i). □

Proposition 1 is the discrete-time Huberman and Stanzl (2004) benchmark: on the round-trip set $\mathcal{R}_T$, impact behaves like a true trading cost if and only if H is positive semidefinite. We impose this baseline below. The mechanism studied in this paper therefore does not rely on classic impact-based manipulation; it relies on the additional feedback that maps transaction prices into binding requirements and induced forced flow.

2.1.4 Example: temporary linear impact

A convenient benchmark is purely temporary impact: set $G_0 = 1$ and $G_\ell = 0$ for all $\ell \geq 1$. Then only the contemporaneous trade affects the execution price. In this case $K = I_T$, where I_T is the $T \times T$ identity matrix, so

$$H = \eta(K + K^\top) = 2\eta\, I_T.$$

Hence $u^\top H u = 2\eta\,\|u\|^2 \geq 0$ for all $u \in \mathbb{R}^T$, and in particular for all round trips $u \in \mathcal{R}_T$. By Proposition 1, the model is manipulation free: every admissible round trip has nonnegative expected execution cost.

2.2 More than one asset

Fix N assets. At each trading date $t = 0, 1, \ldots, T-1$, the trader chooses a trade vector $u_t \in \mathbb{R}^N$, where $(u_t)_n$ is the trade in asset n. Inventory is a vector process $x_t \in \mathbb{R}^N$ evolving as

$$x_{t+1} = x_t + u_t, \qquad x_0 = 0.$$

A vector round trip starts and ends flat,

$$x_T = 0 \text{ a.s.} \qquad \text{equivalently} \qquad \sum_{t=0}^{T-1} u_t = 0$$

with equalities understood componentwise. As before, $S_t \in \mathbb{R}^N$ denotes the unaffected price vector and $P_t(u) \in \mathbb{R}^N$ the execution price vector induced by the strategy.

Assumption 3 (Multi-asset unaffected prices). *$(S_t)_{t=0}^{T-1}$ is an $(\mathcal{F}_t)$-martingale in $\mathbb{R}^N$ and $\mathbb{E}[\|S_t\|] < \infty$ for all t.*

Lemma 2 (Multi-asset martingale benchmark). *Under Assumption 3, let $u = (u_t)_{t=0}^{T-1}$ be a predictable $\mathbb{R}^N$-valued process and uniformly bounded: there exists $\bar{u} > 0$ such that $\|u_t\| \le \bar{u}$ for all t and all ω. If $\sum_{t=0}^{T-1} u_t = 0$, then*

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t^\top S_t\right] = 0. \tag{11}$$

Proof. Let $x_t = \sum_{s=0}^{t-1} u_s$ be inventory, so $u_t = x_{t+1} - x_t$. We have

$$\sum_{t=0}^{T-1} u_t^\top S_t = x_T^\top S_{T-1} - x_0^\top S_0 + \sum_{t=1}^{T-1} x_t^\top (S_{t-1} - S_t).$$

On a round trip $x_0 = x_T = 0$ a.s. Predictability implies x_t is $\mathcal{F}_{t-1}$-measurable, and the martingale property yields $\mathbb{E}[S_t \mid \mathcal{F}_{t-1}] = S_{t-1}$, hence for each $t \ge 1$,

$$\mathbb{E}\Big[x_t^\top (S_{t-1} - S_t)\Big] = \mathbb{E}\Big[x_t^\top \, \mathbb{E}[S_{t-1} - S_t \mid \mathcal{F}_{t-1}]\Big] = 0.$$

Summing over t gives the claim. □

Assumption 4 (Transient linear cross-impact). *Fix matrices $G_0, G_1, \ldots, G_{T-1} \in \mathbb{R}^{N \times N}$. Given a strategy u, the execution price vector at time t is*

$$P_t(u) = S_t + \sum_{s=0}^{t} G_{t-s}\, u_s. \tag{12}$$

The matrix G_{t-s} maps the trade vector at time s into its contribution to the execution price vector at time t. Diagonal entries capture own-asset impact; off-diagonal entries capture cross-impact. Lag dependence allows impact to persist and decay over time. A useful benchmark is purely temporary cross-impact: take $G_0 = A$ for some $A \in \mathbb{R}^{N \times N}$ and $G_\ell = 0$ for all $\ell \ge 1$. Then $P_t(u) = S_t + Au_t$, so only contemporaneous trades affect execution prices.

2.2.1 Execution cost and manipulation

The realized execution cost of a multi-asset strategy is $\sum_{t=0}^{T-1} u_t^\top P_t(u)$. On a round trip this equals minus realized profit. Stack trades over time into $U \in \mathbb{R}^{NT}$, $U = (u_0^\top, \ldots, u_{T-1}^\top)^\top$. Define the block

lower-triangular matrix $\mathbf{K} \in \mathbb{R}^{NT \times NT}$ by its $N \times N$ blocks

$$\mathbf{K}_{t,s} = G_{t-s}\,\mathbf{1}_{\{t \geq s\}}, \qquad t, s \in \{0, 1, \ldots, T-1\}. \tag{13}$$

Then the t-th block of $\mathbf{K}U$ is

$$(\mathbf{K}U)_t = \sum_{s=0}^{t} G_{t-s} u_s,$$

so Assumption 4 implies $P_t(u) = S_t + (\mathbf{K}U)_t$. As in the single asset case, it is convenient to define the symmetric matrix

$$\mathbf{H} = \mathbf{K} + \mathbf{K}^\top.$$

Collect vector round trips in

$$\mathbf{R}_T = \left\{ U \in \mathbb{R}^{NT} : \sum_{t=0}^{T-1} u_t = 0 \right\}, \tag{14}$$

where the constraint is a vector equality in $\mathbb{R}^N$.

Proposition 2. *Suppose Assumptions 3 and 4 hold. Then, for any admissible round trip u,*

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t^\top P_t(u) \right] = \frac{1}{2}\,\mathbb{E}\left[U^\top \mathbf{H} U \right]. \tag{15}$$

Consequently, the model is manipulation free if and only if $\mathbf{H}$ is positive semidefinite on $\mathbf{R}_T$.

Proof. Fix an admissible round trip u. Lemma 2 gives $\mathbb{E}[\sum_{t=0}^{T-1} u_t^\top S_t] = 0$, hence

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t^\top P_t(u) \right] = \mathbb{E}\left[\sum_{t=0}^{T-1} u_t^\top \big(P_t(u) - S_t \big) \right].$$

By Assumption 4, $P_t(u) - S_t = \sum_{s=0}^{t} G_{t-s} u_s$, so

$$\sum_{t=0}^{T-1} u_t^\top \big(P_t(u) - S_t \big) = \sum_{t=0}^{T-1} \sum_{s=0}^{t} u_t^\top G_{t-s} u_s = U^\top \mathbf{K} U.$$

Since $U^\top \mathbf{K} U$ is a scalar, $U^\top \mathbf{K} U = U^\top \mathbf{K}^\top U$, hence

$$U^\top \mathbf{K} U = \frac{1}{2}\, U^\top (\mathbf{K} + \mathbf{K}^\top) U = \frac{1}{2}\, U^\top \mathbf{H} U.$$

Taking expectations yields (15). The model admits a manipulation strategy if and only if there exists $U \in \mathbf{R}_T$ with $U^\top \mathbf{H} U < 0$, which is equivalent to failure of positive semidefiniteness of $\mathbf{H}$ on $\mathbf{R}_T$. □

With multiple assets, round-trip profits can arise from cross-impact even when each asset in isolation is manipulation free. If $\mathbf{H}$ fails to be positive semidefinite on $\mathbf{R}_T$, then some vector round trip has negative expected execution cost: trading in one asset moves execution prices in another, altering the cost of subsequent legs and allowing the strategy to close out flat at a gain. Proposition 2 gives the finite-horizon condition that rules out such cross-impact loops within the transient linear class.

2.2.2 Example: two-asset arbitrage

A concrete setting to illustrate the cross-asset no-arbitrage condition is a pair of tightly linked instruments with asymmetric hedging pressure, for example a very liquid index future (asset 1) and a less liquid cash instrument that tracks the same index (asset 2), such as an ETF or a basket proxy. When a trader buys the future aggressively, liquidity providers often hedge by buying the cash instrument. This transmits order flow from the future into the cash execution price. The reverse channel can be weaker if trading in the cash instrument is more fragmented or less informative. Empirically, lead-lag and price-discovery patterns of this type are well documented in futures versus spot settings (Kawaller et al., 1987; Chan, 1992). Model this with purely temporary cross-impact, $P_t(u) = S_t + Au_t$, and

$$A = \begin{pmatrix} \lambda_F & \epsilon \\ \Lambda & \lambda_C \end{pmatrix}, \qquad \lambda_F > 0,\ \lambda_C > 0,\ \Lambda > \epsilon \geq 0.$$

The diagonal terms λ_F, λ_C are own-impact. The off-diagonal terms are cross-impact: trading asset 1 shifts the execution price of asset 2 with slope Λ, while trading asset 2 shifts the execution price of asset 1 with a small slope ϵ. Expected execution cost depends on

$$H = \frac{A + A^\top}{2} = \begin{pmatrix} \lambda_F & (\epsilon + \Lambda)/2 \\ (\epsilon + \Lambda)/2 & \lambda_C \end{pmatrix}.$$

A profitable round trip exists if and only if H is not positive semidefinite, equivalently if there exists $q \neq 0$ with $q^\top A_s q < 0$. In two dimensions this is equivalent to

$$\det(H) = \lambda_F \lambda_C - \left(\frac{\epsilon + \Lambda}{2}\right)^2 < 0.$$

Cross-impact can dominate own-impact, creating a direction in which quadratic impact cost is negative. The simplest direction of attack is the spread $v = (1, -1)^\top$ (buy asset 1, sell asset 2). Along v,

$$v^\top A v = v^\top A_s v = \lambda_F + \lambda_C - (\epsilon + \Lambda).$$

If $\epsilon + \Lambda > \lambda_F + \lambda_C$, then trading the spread has negative quadratic cost. With $T = 2$, consider the round trip $u_0 = v$ and $u_1 = -v$, which opens the spread at $t = 0$ and fully reverses it at $t = 1$. The strategy starts flat and ends flat, and

$$\mathbb{E}\left[\sum_{t=0}^{1} u_t^\top P_t(u)\right] = \mathbb{E}\left[\sum_{t=0}^{1} u_t^\top A u_t\right] = 2\, v^\top A v.$$

For instance, if $\lambda_F = \lambda_C = 0.50$, $\epsilon = 0.05$, and $\Lambda = 1.50$, then $v^\top A v = 0.50 + 0.50 - (0.05 + 1.50) = -0.55$, so expected execution cost equals $2(-0.55) = -1.10 < 0$ and expected profit is strictly positive. The opening trade in the liquid instrument pushes the execution price of the other leg through hedging-induced cross-impact, while the reversal inherits a favorable wedge that more than offsets the own-impact paid on both legs. Proposition 2 rules this out by requiring $\mathbf{H}$ is p.s.d. on $\mathbf{R}_T$, so that no vector round trip can have negative expected quadratic impact cost.

3 Price-based constraints and forced deleveraging

We now introduce a constrained sector that follows a mechanical rule and show that this restores manipulation opportunities even when the standard no-manipulation condition holds. Including such a sector is economically natural because it is pervasive in modern markets. Many trading mandates and market institutions update binding constraints mechanically from recent transaction prices. Examples include risk-managed mandates (volatility targeting, risk parity, drawdown or VaR limits, and multi-strategy hedge funds with tight risk constraints) and risk regulation that rescales exposure when measured risk rises. Broker-dealers and prime brokers update haircuts, credit limits, and internal risk limits from the same inputs; central counterparties reset initial margins from recent prints; and stress-test triggers and risk-control overlays often take the same form. The same structure is explicit in crypto: derivatives venues update margins and liquidation thresholds from mark prices, and lending protocols revalue collateral from oracle marks and liquidate mechanically when collateral ratios breach preset thresholds. Even passive trading rules, such as mechanical momentum or reversal strategies or equal-weighted portfolios, adjust holdings based on realized returns.[4] The common structure is that past transaction prices determine tomorrow's feasible position. With price impact, this creates a feedback loop: prices affect tomorrow's constraint; a binding constraint forces trades; forced trades move prices.

Fix trading dates $t = 0, \dots, T-1$, and let P_t denote transaction prices (scalar in the single-asset case, vector in the multi-asset case). Fix a window length $m \geq 0$. At each date t, the rule maps the most recent $m+1$ transaction prices into a scalar risk input,

$$\Gamma_t = \Gamma(P_{t-m}, \dots, P_t), \tag{16}$$

and sets the next-date requirement

$$M_{t+1} = g(\Gamma_t), \tag{17}$$

where g is nondecreasing. A constrained sector holds positions X_t and must satisfy a requirement-based feasibility constraint. A representative example is

$$|X_t(i)|\, M_t \leq E(i) \quad \text{for all constrained agents } i, \tag{18}$$

where $E(i)$ denotes equity. When the constraint binds, an increase in M_t forces a reduction in $|X_t(i)|$. Aggregating across constrained agents yields an induced sector demand rule $X_t = X(M_t)$. The mechanism follows from the mechanical sensitivity of constrained demand to past prices, i.e. the fact that $dX_t/dP_{t-s} \neq 0$. As long as this sensitivity exists, the mechanism follows. Fix a binding requirement M_0 and define the local sensitivity $B = -X'(M_0) > 0$. Rebalancing generates forced order flow $v_t = X_t - X_{t-1}$. Let u_t denote the strategic trader's order, and define total order flow

$$q_t = u_t + v_t.$$

[4] Market-cap-weighted index funds, however, are immune to this type of rebalancing, except for rebalancing due to changes in index composition (Sammon and Shim, 2026).

Appendix A microfounds the mapping $X(M)$, defines $B = -X'(M_0)$ at a binding state, and records the fully constrained benchmark $X = W/M$ (hence $B = W/M_0^2$).

Assumption 5 (Temporary linear impact with total order flow)**.** *Transaction prices satisfy*

$$P_t = S_t + \mathcal{A}q_t, \tag{19}$$

where $\mathcal{A} = \eta > 0$ in the single-asset case, and $\mathcal{A} = A \in \mathbb{R}^{N\times N}$ in the multi-asset case, with A symmetric and $z^\top Az \geq 0$ for all $z \in \mathbb{R}^N$.

Assumption 5 imposes the classical no-manipulation benchmark: absent the price-based constraint, the impact block cannot generate positive expected profits from any admissible round trip (Huberman and Stanzl, 2004; Gatheral, 2010). It follows that any profitable round trip in the rest of the paper is not impact-only manipulation. It is created by a distinct channel: transaction prices enter the rule (16)–(17), the rule moves a binding requirement, and the requirement induces forced flow v_t that feeds back into prices. The mechanism is pinned down by four primitives: the statistic Γ_t, the schedule g, the constrained-sector mapping $X(M)$ (equivalently, the mapping from M_t into v_t), and the impact map from total flow q_t into transaction prices P_t. The remainder of the paper takes these primitives as given, characterizes when their composition admits profitable round trips, and derives stress tests for a given rule.

3.1 Profit decomposition and the mechanical source of arbitrage

Consider the single-asset case. For an order sequence $u = (u_t)_{t=0}^{T-1}$, realized trading profit is

$$\Pi(u) = -\sum_{t=0}^{T-1} u_t P_t. \tag{20}$$

For a round trip, Lemma 1 implies

$$\mathbb{E}[\Pi(u)] = -\mathbb{E}\left[\sum_{t=0}^{T-1} u_t \left(P_t - S_t\right)\right].$$

Under Assumption 5, $P_t - S_t = \eta q_t = \eta(u_t + v_t)$. Substituting gives

$$\mathbb{E}[\Pi(u)] = -\eta\,\mathbb{E}\left[\sum_{t=0}^{T-1} u_t^2\right] - \eta\,\mathbb{E}\left[\sum_{t=0}^{T-1} u_t v_t\right]. \tag{21}$$

The term $-\eta\sum_t u_t^2$ is the direct impact cost. The term $-\eta\sum_t u_t v_t$ is the only channel through which expected profit can be positive: it is positive exactly when u_t tends to have the opposite sign of forced flow v_t. For instance, during forced liquidation $v_t < 0$, a buy order $u_t > 0$ makes $u_t v_t < 0$ and increases expected profit. The source of arbitrage is that v_t is endogenous and rule-driven. Forced flow is the mechanical rebalancing of a constrained sector,

$$v_t = X_t - X_{t-1}, \qquad X_t = X(M_t),$$

where the requirement $M_{t+1} = g(\Gamma_t)$ is computed from past transaction prices $\Gamma_t = \Gamma(P_{t-m}, \ldots, P_t)$. Because transaction prices satisfy $P_t = S_t + \eta(u_t + v_t)$, the strategic order u_t affects the realized prices sampled by the rule, hence future requirements M_{t+1} and therefore future forced flow $v_{t+1}, v_{t+2}, \ldots$.

4 An augmented no-arbitrage condition for price-based constraints

We now derive the correct no-arbitrage condition for this economy. We maintain the standard finite-horizon benchmark for market impact: absent any constraint feedback, every admissible round trip has nonnegative expected execution cost (Huberman and Stanzl, 2004; Gatheral, 2010; Schneider and Lillo, 2019) (Propositions 1–2). This section asks whether a deterministic price-based constraint can nevertheless generate profitable round trips. The manipulator's profits here do not come from an arbitrage embedded in the impact operator: by assumption, the impact block is manipulation-free. The channel is also distinct from standard predatory trading, where the counterparty's liquidation path is taken as exogenous to the attacker (Brunnermeier and Pedersen, 2005; Carlin et al., 2007). In our setting, the liquidation schedule is mechanically generated by a publicly specified constraint and shifts, at the margin, with trades that move the sampled statistic.We study this feedback locally by linearizing around a configuration in which the requirement is binding. Theorem 1 provides a test for this linearized feedback loop: manipulation is absent if and only if the symmetric cost matrix that endogenizes this feedback loop is positive semidefinite on the round-trip subspace. The test depends on the impact map $\mathcal{I}$, the sensitivity of the risk statistic J, the margin-schedule slope $s = g'(\Gamma_0)$, the requirement-to-position sensitivity B, and the timing operators L and D that govern implementation lags of the strategy. Although the condition is derived from a linearization, Appendix C shows that failure of the test implies a truly profitable, sufficiently small round trip in the exact nonlinear model (Proposition 9).

4.1 A linearized model

Fix a horizon T. We study a small trade-induced deviation, holding the unaffected price path S fixed. For any object Y, let δY denote its deviation from the reference value induced by trading; in particular, the fundamental price is left unaffected, so $\delta S = 0$. Let $u \in \mathbb{R}^T$ be the strategic trader's deviation from the current strategy and let $v \in \mathbb{R}^T$ be the resulting deviation in forced flow from the constrained sector. Total flow is

$$q = u + v.$$

P is the transaction price path. Total flow creates a wedge between transaction prices and unaffected prices. In deviations,

$$\delta P = \delta S + \mathcal{I} q, \tag{22}$$

where $\mathcal{I}$ is the linear map from flow paths to the induced transaction-price wedge. It is convenient to denote this wedge by

$$\delta \Delta P = \delta P - \delta S = \mathcal{I} q.$$

Let $\delta\Gamma \in \mathbb{R}^T$ collect the trade-induced deviations in the statistic over the horizon, and let $\delta\Delta P \in \mathbb{R}^T$ collect the corresponding deviations in wedge prices.

Assumption 6 (Local Regularity for Linearization)**.** *Fix the reference configuration used for the linearization in this section. There exists an open neighborhood of the reference transaction-price path (equivalently, the wedge-price path since $\delta S = 0$) such that:*

(i) The statistic map $\Gamma(P)$ is differentiable at the reference path, so the Jacobian J used in the linearization exists; the schedule g is differentiable at the relevant statistic value, so $g'(\Gamma)$ is well defined.

(ii) If the constraint includes nondifferentiable components such as caps or floors, turnover buffers, max or min operations across windows, or other piecewise constructions, the linearizations in the main text are taken at interior points where the active branch is locally constant and the defining inequalities are slack. At kink points one can instead work with one-sided derivatives or a subgradient selection; replacing J and $g'(\Gamma)$ by such a selection yields a conservative sufficient version of the linearized admissibility screen.

The statistic is a deterministic function of the sampled transaction-price path, so its first-order approximation is

$$\delta\Gamma = J\,\delta\Delta P, \tag{23}$$

where J is the Jacobian of $\Gamma(P)$ evaluated at the reference configuration, holding S fixed. To keep requirements aligned with the trading dates $t = 0, \ldots, T-1$, let $\delta M \in \mathbb{R}^T$ denote deviations in posted requirements that apply at those dates, so $(\delta M)_t$ corresponds to M_t. Under the timing $M_{t+1} = g(\Gamma_t)$, within-horizon deviations satisfy $(\delta M)_0 = 0$ and, for $t = 1, \ldots, T-1$,

$$(\delta M)_t = s\,(\delta\Gamma)_{t-1},$$

where $s = g'(\Gamma_0)$ is the local slope of the schedule at the reference statistic value. Equivalently,

$$\delta M = s\,L\,\delta\Gamma,$$

where $L \in \mathbb{R}^{T\times T}$ is the one-step lag operator.

4.1.1 From margins to forced flow

Let B denote the local sensitivity of aggregate constrained demand $X = X(M)$ at the binding point. Under the benchmark $X = W/M$, this sensitivity equals $B = W/M_0^2$; more generally, when X is differentiable at M_0 one can write $B = -X'(M_0) > 0$ (see Appendix A.1). Linearizing aggregate demand around the binding point gives

$$\delta X = -B\,\delta M.$$

Forced flow is the change in constrained demand, $v_t = X_t - X_{t-1}$, so in deviations

$$v = -B\,D\,\delta M, \tag{24}$$

where $D \in \mathbb{R}^{T\times T}$ is the first-difference operator.

4.1.2 Closing the loop

Combining the linearizations yields the induced forced-flow response to a wedge-price deviation,

$$v = -B\, s\, D\, L\, J\, \delta\Delta P. \tag{25}$$

Using $\delta\Delta P = \mathcal{I}q$ and $q = u + v$, the closed loop satisfies

$$q = u - BsDLJ\mathcal{I}q. \tag{26}$$

Define

$$\mathcal{K} = BsDLJ\mathcal{I}. \tag{27}$$

Then $(I + \mathcal{K})q = u$. $I + \mathcal{K}$ is invertible[5] and the fixed point selects a unique total-flow path for each strategic deviation u:

$$q = (I + \mathcal{K})^{-1}u, \qquad \delta\Delta P = \mathcal{I}(I + \mathcal{K})^{-1}u. \tag{28}$$

4.2 The augmented no-arbitrage condition

We derive an augmented no-arbitrage screen that internalizes deterministic constraint feedback. Fix an admissible round trip $u \in \mathcal{R}_T$. Under the martingale benchmark,

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t S_t\right] = 0$$

so expected profits depend only on the execution wedge $P - S$. In the linearized loop, (28) implies

$$P - S = \delta\Delta P = \mathcal{I}(I + \mathcal{K})^{-1}u.$$

Define the effective impact

$$\widetilde{\mathcal{I}} = \mathcal{I}(I + \mathcal{K})^{-1}.$$

This impact matrix maps the proposed trade path into the entire wedge-price path, incorporating both direct impact and the indirect component coming from induced forced flow. Expected execution cost is therefore

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t P_t\right] = \mathbb{E}\left[u^\top (P - S)\right] = \mathbb{E}\left[u^\top \widetilde{\mathcal{I}}\, u\right]. \tag{29}$$

As usual, define the augmented symmetric cost matrix

$$\widehat{H} = \widetilde{\mathcal{I}} + \widetilde{\mathcal{I}}^\top = \mathcal{I}(I + \mathcal{K})^{-1} + (I + \mathcal{K})^{-\top}\mathcal{I}^\top. \tag{30}$$

[5] Lemma 11 formalizes this argument; Appendix B proves it and contrasts lagged updating with instantaneous updating $M_t = g(\Gamma_t)$, where invertibility can fail. Near singularity, small strategic trades can be strongly amplified into total flow, producing the quasi-arbitrage pathology highlighted there.

With $\Pi(u) = -\sum_{t=0}^{T-1} u_t P_t$, it follows that

$$\mathbb{E}[\Pi(u)] = -\mathbb{E}\Big[u^\top \widetilde{\mathcal{I}}\, u\Big] = -\frac{1}{2}\,\mathbb{E}\Big[u^\top \widehat{H}\, u\Big]. \tag{31}$$

A profitable round trip exists exactly when $\widehat{H}$ has a negative direction on $\mathcal{R}_T$, meaning that the rule-induced response tilts the induced flow so that the trader can unwind on favorable terms often enough to offset direct impact costs.

Theorem 1 (Augmented no-arbitrage condition)**.** *Consider the linearized feedback system* (22)–(28) *around a binding state, and assume $I + \mathcal{K}$ is invertible.*

(i) Single asset. The combined system is margin-feedback manipulation free in the linearized class if and only if $\widehat{H}$ is positive semidefinite on $\mathcal{R}_T$.

(ii) Multi asset. Let $\mathcal{I} \in \mathbb{R}^{NT\times NT}$ map stacked total flow $q \in \mathbb{R}^{NT}$ into stacked wedge-price deviations $\delta\Delta P \in \mathbb{R}^{NT}$ via $\delta\Delta P = \mathcal{I}q$, and let $J \in \mathbb{R}^{T\times NT}$ map wedge-price paths into statistic deviations via $\delta\Gamma = J\,\delta\Delta P$. Under $M_{t+1} = g(\Gamma_t)$, write $\delta M = s\,L\,\delta\Gamma$ with $s = g'(\Gamma_0)$ and L the one-step lag operator. Suppose the constrained sector rebalances along a fixed liquidation direction $b \in \mathbb{R}^N$ with local sensitivity $B = -X'(M_0) > 0$, so the induced forced-flow deviations satisfy

$$v = -Bs\,(D \otimes I_N)(I_T \otimes b)\,L\,J\,\delta\Delta P.$$

Define

$$\mathcal{K} = Bs\,(D \otimes I_N)(I_T \otimes b)\,L\,J\,\mathcal{I} \in \mathbb{R}^{NT\times NT},$$

and

$$\widehat{\mathbf{H}} = \mathcal{I}(I + \mathcal{K})^{-1} + (I + \mathcal{K})^{-\top}\mathcal{I}^\top.$$

Then the combined system is margin-feedback manipulation free if and only if $\widehat{\mathbf{H}}$ is positive semidefinite on round trips $u \in \mathbf{R}_T$.

Proof. Part (ii) follows from the same argument after stacking time-and-asset vectors, so it suffices to prove part (i). Fix an admissible round trip u. By Assumption 1 and the round-trip condition,

$$\mathbb{E}\left[\sum_{t=0}^{T-1} u_t S_t\right] = 0.$$

Using (28), we have $P - S = \widetilde{\mathcal{I}}u$, hence (31) gives

$$\mathbb{E}[\Pi(u)] = -\frac{1}{2}\,\mathbb{E}\Big[u^\top \widehat{H}\, u\Big].$$

If $\widehat{H}$ is positive semidefinite on $\mathcal{R}_T$, then $u(\omega) \in \mathcal{R}_T$ a.s. implies $u(\omega)^\top \widehat{H}\, u(\omega) \geq 0$ a.s., so $\mathbb{E}[\Pi(u)] \leq 0$ for every admissible round trip. Conversely, if $\widehat{H}$ is not positive semidefinite on $\mathcal{R}_T$, choose $\tilde{u} \in \mathcal{R}_T$ with $\tilde{u}^\top \widehat{H}\, \tilde{u} < 0$ and take the deterministic admissible round trip $u \equiv \tilde{u}$. Then $\mathbb{E}[\Pi(u)] = -\frac{1}{2}\,\tilde{u}^\top \widehat{H}\, \tilde{u} > 0$, so margin-feedback manipulation exists. □

Appendix B, Theorem 4 rewrites the condition as a restricted eigenvalue test, which is convenient for computation.

The mechanism behind Theorem 1 is distinct from standard models of predatory trading and order anticipation (Brunnermeier and Pedersen, 2005; Carlin et al., 2007; Attari et al., 2005). In that literature, the strategic trader profits from an order imbalance generated by a counterparty's balance-sheet stress, funding constraints, or liquidation needs. Conditional on the relevant state variables, liquidation pressure is taken as given from the predator's perspective; any feedback from the predator's own trades operates primarily through the price level, which worsens the counterparty's financing terms. Here, by contrast, the imbalance is produced by a disclosed, deterministic mapping from sampled transaction prices into next-period requirements. When requirements bind, the induced rebalancing is mechanical and pinned down by observable derivatives of the rule. The strategic trader can therefore shift the magnitude and timing of future forced flow by trading to move the risk statistic, so the profit opportunity takes the form of rule-induced trigger-and-reverse trading rather than intermediation against an exogenously distressed liquidator.

4.3 Classical no-manipulation is not sufficient under deterministic margin rules

We now provide an example to show that an impact matrix that is manipulation free in the classical sense (Huberman and Stanzl, 2004) can still admit profitable round trips once a deterministic price-based feedback rule is added. Consider two trading dates $t = 0, 1$ and temporary linear impact under total flow,

$$P_t = S_t + \eta q_t, \qquad q_t = u_t + v_t, \qquad \eta > 0,$$

and take the unaffected price to be a martingale (for simplicity, $S_t = S$). The round-trip subspace is $\mathcal{R}_2 = \{u \in \mathbb{R}^2 : u_0 + u_1 = 0\}$. Without feedback, $v = 0$ and hence $q = u$, so for any $u \in \mathcal{R}_2$,

$$\Pi(u) = -\sum_{t=0}^{1} u_t P_t = -\eta(u_0^2 + u_1^2) \leq 0,$$

which is the classical no-manipulation benchmark. Now introduce a reduced-form, deterministic feedback from the date-0 trade to forced flow at date 1: set $v_0 = 0$ and $v_1 = ku_0$ with $k > 0$. Then $q_0 = u_0$ and $q_1 = u_1 + ku_0$. For the round trip $u = (q, -q)$,

$$\Pi(u) = -\eta\big(u_0 q_0 + u_1 q_1\big) = \eta(k-2)q^2,$$

so the round trip is profitable whenever $k > 2$. In this two-period setting, a deviation is profitable if it induces more than twice the flow that it initially requires. This reduced-form rule is consistent with the linearized primitives in Section 4. In deviations,

$$v = -BsDLJ\,\delta\Delta P, \qquad \delta\Delta P = \mathcal{I}q.$$

With $T = 2$ and $\mathcal{I} = \eta I_2$, choose J so that $J_{0,0} = -k/(Bs\eta)$ and all other entries are zero. Then

$$\mathcal{K} = BsDLJ\mathcal{I} = \begin{pmatrix} 0 & 0 \\ -k & 0 \end{pmatrix},$$

and the fixed point $(I_2 + \mathcal{K})q = u$ reproduces $v_0 = 0$ and $v_1 = ku_0$. Since

$$(I_2 + \mathcal{K})^{-1} = \begin{pmatrix} 1 & 0 \\ k & 1 \end{pmatrix},$$

the effective impact matrix is $\widetilde{\mathcal{I}} = \mathcal{I}(I_2 + \mathcal{K})^{-1}$ and the augmented symmetric cost matrix becomes

$$\widehat{H} = \widetilde{\mathcal{I}} + \widetilde{\mathcal{I}}^\top = \eta \begin{pmatrix} 2 & k \\ k & 2 \end{pmatrix}.$$

For $u = (-q, q) \in \mathcal{R}_2$,

$$u^\top \widehat{H} u = -2\eta(k-2)q^2 < 0 \quad \text{when } k > 2,$$

which matches the positive-profit region above.

5 Single-asset optimal attack under price-based margining

Section 4 provides a local diagnostic for when deterministic, price-based requirements overturn the classical no-dynamic-arbitrage benchmark in temporary-impact models (Huberman and Stanzl, 2004; Gatheral, 2010; Schneider and Lillo, 2019). This section takes the rule and impact as given and solves the strategic trader's finite-horizon optimal round trip in the single-asset environment. The solution separates two motives. First, the attacker earns profits by supplying liquidity against predictable, constraint-driven orders, as in models of order anticipation and predatory trading (Brunnermeier and Pedersen, 2005). Second, and in sharp contrast to that literature, the attacker can trade within the sampling window to move the rule's input and thereby reshape future forced flow. The dynamic program makes these incentives explicit.

5.1 The dynamic program

Throughout impose Assumption 5 and the fully constrained benchmark $X_t = W/M_t$ from Lemma 7. Fix a finite horizon T with trading dates $t = 0, 1, \ldots, T-1$. Initial conditions are $(P_{-1}, M_0, X_{-1}, X_0)$, where P_{-1} is the last pre-horizon transaction price. For the optimal-attack dynamic programs in Sections 5–6, we work conditional on a deterministic unaffected-price path $(S_t)_{t=-1}^{T-1}$, so the state transition is deterministic given (s_t, u_t) and we omit conditional expectations. At the start of date t, before choosing u_t, the state is

$$s_t = (x_t,\ y_t,\ P_{t-1},\ M_t,\ X_t,\ X_{t-1}). \tag{32}$$

The attacker controls inventory x_t through trades. The pair (y_t, P_{t-1}) summarizes the information used by the margin rule to map the next transaction price into the next requirement. The remaining components (M_t, X_t, X_{t-1}) pin down the constrained sector's current and lagged holdings, which determine its mechanically induced net trade on date t. The rule is Markov in (y_t, P_t):

$$y_{t+1} = F(y_t, P_t), \qquad \Gamma_t = \Gamma(y_t, P_{t-1}, P_t), \qquad M_{t+1} = g(\Gamma_t), \tag{33}$$

and under the benchmark $X_{t+1} = W/M_{t+1}$. The constrained sector's date-t forced order flow is therefore predetermined at the start of date t:

$$v_t = X_t - X_{t-1}. \tag{34}$$

The attacker observes v_t in the state and chooses u_t subject to a per-date bound $|u_t| \le \bar{u}$. To ensure feasibility of the terminal round-trip constraint $x_T = 0$, restrict actions to those that leave enough remaining capacity to unwind:

$$\mathcal{U}_t(x_t) = \{u \in \mathbb{R} : \ |u| \le \bar{u}, \ |x_t + u| \le (T - 1 - t)\bar{u}\}. \tag{35}$$

Given u_t, total order flow is $q_t = u_t + v_t$ and transaction prices follow temporary impact:

$$P_t = S_t + \eta q_t = S_t + \eta(u_t + v_t). \tag{36}$$

The attacker's realized profit is $\Pi(u) = -\sum_{t=0}^{T-1} u_t P_t$. Because we impose a terminal round trip, $x_T = 0$, and treat the unaffected price as a martingale benchmark, the unaffected-price component has zero expected contribution to profits. Conditional on the realized path (S_t), the only component of $\Pi(u)$ that depends on the choice of u_t is therefore the execution wedge $P_t - S_t$. Accordingly, we can write the per-period payoff as

$$\pi_t(s_t, u_t) = -u_t(P_t - S_t) = -\eta u_t(u_t + v_t), \tag{37}$$

which is the objective optimized by the dynamic program. The direct impact term $-\eta u_t^2$ is always negative, while the interaction term $-\eta u_t v_t$ is positive precisely when the attacker trades against the mechanically induced flow. Dynamic incentives arise because u_t changes both inventory, $x_{t+1} = x_t + u_t$, and the transaction price P_t, which enters (33) and can therefore move future requirements and future forced flow. Given s_t and u_t, the next state is

$$s_{t+1} = \mathcal{T}(s_t, u_t) = (x_t + u_t, \ F(y_t, P_t), \ P_t, \ g(\Gamma_t), \ W/g(\Gamma_t), \ X_t), \tag{38}$$

where P_t is given by (36) and Γ_t by (33). Define the value function

$$V_t(s) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \sum_{\tau=t}^{T-1} \pi_\tau(s_\tau, u_\tau), \tag{39}$$

subject to $u_\tau \in \mathcal{U}_\tau(x_\tau)$, the transition (38), and the terminal constraint $x_T = 0$. The associated Bellman

recursion is

$$V_T(s) = 0 \text{ if } x_T = 0, \qquad V_T(s) = -\infty \text{ otherwise,} \tag{40}$$

$$V_t(s) = \max_{u \in \mathcal{U}_t(x)} \{\pi_t(s,u) + V_{t+1}(\mathcal{T}(s,u))\}, \qquad t = T-1, \dots, 0. \tag{41}$$

5.2 Trade decomposition: harvesting, inventory management, and triggering

The dynamic program highlights two distinct roles of a trade. First, it determines the current payoff through (37). Second, by moving P_t, it affects the rule and can therefore reshape future forced flow. A useful benchmark holds future incentives fixed and treats v_t as predetermined, isolating the purely mechanical intermediation problem. This is the single-period analogue of supplying liquidity against a given liquidation pressure in the predatory trading literature (Brunnermeier and Pedersen, 2005; Carlin et al., 2007).

Lemma 3 (Myopic harvesting)**.** *Fix a date t and a state s_t with predetermined forced flow $v_t = X_t - X_{t-1}$. Under (37), the one-period payoff satisfies*

$$\pi_t(s_t, u) = -\eta\Big(u + \tfrac{v_t}{2}\Big)^2 + \frac{\eta}{4} v_t^2. \tag{42}$$

Consequently, holding the continuation value fixed, $\pi_t(s_t, u)$ is uniquely maximized at

$$u_t^{\text{myopic}} = -\frac{v_t}{2}, \qquad \max_{u \in \mathbb{R}} \pi_t(s_t, u) = \frac{\eta}{4} v_t^2. \tag{43}$$

If $v_t = 0$, then $\pi_t(s_t, u) \le 0$ for all u, with equality only at $u = 0$.

Proof. Completing the square in (37) gives (42). The square term is minimized at $u = -v_t/2$, yielding the unique maximizer and the maximal value in (43). If $v_t = 0$, then $\pi_t(s_t, u) = -\eta u^2 \le 0$, with equality only at $u = 0$. □

Lemma 3 provides a pure harvesting benchmark. When v_t is predetermined, gains arise only from trading against the forced order, and the myopic optimum absorbs exactly half of it. The quadratic form (42) shows that the attacker is compensated for providing immediacy against a predictable imbalance, while paying the usual temporary-impact cost on own orders.

Optimal attacks depart from the previous myopic harvesting benchmark because the trader must satisfy the terminal round-trip constraint and because u_t affects future requirements and thus future forced flow. The next result characterizes the interior optimum by rewriting the first-order condition in terms of two shadow values: the marginal value of inventory carried into $t+1$ and the marginal value of shifting the current transaction price, which affects the rule state and thus future forced flow.

Proposition 3. *Fix $t \in \{0, \dots, T-1\}$ and a state s_t. Suppose the optimizer $u_t^\star$ in (41) is interior. Assume that V_{t+1} is differentiable at $s_{t+1}^\star = \mathcal{T}(s_t, u_t^\star)$ and that $\mathcal{T}(s_t, u)$ is differentiable at $u_t^\star$. Define*

$$\mu_{t+1} = \partial_x V_{t+1}(s_{t+1}^\star), \qquad \nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1}^\star), \qquad P_t = S_t + \eta(u_t^\star + v_t).$$

Then $u_t^\star$ satisfies the Euler equation

$$\eta(2u_t^\star + v_t) = \mu_{t+1} + \eta\,\nu_{t+1}, \tag{44}$$

or equivalently

$$u_t^\star = -\frac{v_t}{2} + \frac{1}{2\eta}\mu_{t+1} + \frac{1}{2}\nu_{t+1}. \tag{45}$$

Moreover, the resulting current-period payoff satisfies

$$\pi_t(s_t, u_t^\star) = \frac{\eta}{4}v_t^2 - \eta\Big(u_t^\star + \tfrac{v_t}{2}\Big)^2 = \frac{\eta}{4}v_t^2 - \frac{1}{4\eta}\big(\mu_{t+1} + \eta\nu_{t+1}\big)^2. \tag{46}$$

Proof. See Appendix D.2. □

The decomposition (45) isolates three forces. The term $-v_t/2$ is the myopic harvesting benchmark from Lemma 3. The inventory shadow value μ_{t+1} captures the intertemporal motive familiar from optimal execution with linear impact (Almgren and Chriss, 2001) and from transient-impact models (Obizhaeva and Wang, 2013): even if future forced flow were held fixed, the attacker would generally tilt away from myopic harvesting to manage the remaining unwind implied by $x_T = 0$. The price shadow value ν_{t+1} is specific to the feedback environment because P_t is an input to the requirement update; when a marginal change in the current price shifts future requirements and thus future forced flow, ν_{t+1} is the continuation-value benefit of moving P_t. The term $\frac{\eta}{4}v_t^2$ in equation (46) is the mechanical rent available from intermediating the predetermined imbalance, while the second term is the wedge cost of departing from the myopic liquidity-provision benchmark to manage inventory and to move the rule's input. A pure trigger trade corresponds to a choice of $u_t^\star$ that is deliberately off the harvesting benchmark, often making $\pi_t(s_t, u_t^\star)$ negative on impact, because the continuation-value gains embodied in μ_{t+1} and/or ν_{t+1} dominate.

Corollary 1 (Trigger dates versus harvest dates)**.** *Fix a date t with $v_t = 0$. Lemma 3 implies $\pi_t(s_t, u) \le 0$ for all u. Assume $0 \in \mathcal{U}_t(x_t)$ (equivalently, $|x_t| \le (T-1-t)\bar{u}$), so that $u = 0$ is feasible. Then any optimal action with $u_t^\star \neq 0$ must be chosen for its effect on continuation value in (41), since it cannot raise the current-period wedge payoff. If, in addition, the conditions of Proposition 3 hold at $(s_t, u_t^\star)$, then the decomposition (45) reduces to*

$$u_t^\star = \frac{1}{2\eta}\mu_{t+1} + \frac{1}{2}\nu_{t+1},$$

so any nonzero trade is driven by the inventory and/or rule-manipulation motives. Conversely, at any interior optimum with $\mu_{t+1} = 0$ and $\nu_{t+1} = 0$, the optimal trade reduces to pure harvesting, $u_t^\star = -v_t/2$.

Proof. See Appendix D.3. □

Dates with nonzero v_t are natural harvest dates because profits are available immediately by leaning against the forced order. When $v_t = 0$, any nonzero trade gives up wedge profits on impact and can be optimal only because it increases continuation value, either by repositioning inventory ahead of later harvest opportunities or by moving the rule's input to reshape future forced flow. In this sense, feedback makes the liquidation path endogenous: the same strategic trades that supply liquidity against future imbalances can also shift the rule that generates those imbalances, unlike standard

predatory trading and order-anticipation models in which the relevant order imbalance is taken as given (Brunnermeier and Pedersen, 2005; Carlin et al., 2007; Rostek and Weretka, 2015; Sannikov and Skrzypacz, 2016; Fardeau, 2021).

The decomposition also clarifies when the attacker may rationally trade in the same direction as forced flow. Such procyclical trading is never selected by the static harvesting benchmark, since it worsens the current payoff. It can be optimal only when continuation value dominates the harvesting motive, either because inventory considerations shift trading toward later dates or because moving the current price changes future requirements and enlarges future forced flow enough to compensate for the extra impact cost today.

Corollary 2. *Fix $t \in \{0, \dots, T-1\}$ and suppose the conditions of Proposition 3 hold at $(s_t, u_t^\star)$. Define*

$$\chi_{t+1} = \mu_{t+1} + \eta \nu_{t+1}.$$

(i) If $v_t \neq 0$, then $u_t^\star v_t > 0$ if and only if $\operatorname{sign}(\chi_{t+1}) = \operatorname{sign}(v_t)$ and $|\chi_{t+1}| > \eta |v_t|$.

(ii) The one-period payoff at the optimal action satisfies $\pi_t(s_t, u_t^\star) < 0$ if and only if $|\chi_{t+1}| > \eta |v_t|$.

Proof. See Appendix D.4. □

Trading with the forced order requires χ_{t+1} to be large enough to overturn harvesting and flip the sign of the optimal trade, which is equivalent to $|\chi_{t+1}| > \eta |v_t|$ with $\operatorname{sign}(\chi_{t+1}) = \operatorname{sign}(v_t)$. In that region the attacker necessarily accepts a negative current payoff, using the date-t trade as an investment in continuation value rather than a harvesting trade.

5.3 How the margin rule shapes the manipulation motive

Proposition 3 identifies ν_{t+1} as the continuation-value gain from moving the current transaction price through the rule update. The next result expresses ν_{t+1} directly in terms of the rule primitives (F, Γ, g) and derivatives of the value function, making transparent which parts of the rule generate a manipulation incentive and how they enter the Euler equation.

Proposition 4 (Decomposition of ν_{t+1})**.** *Fix $t \in \{0, \dots, T-1\}$ and consider any state-action pair (s_t, u_t). Let P_t be given by (36) and let $s_{t+1} = \mathcal{T}(s_t, u_t)$. Assume V_{t+1} is differentiable at s_{t+1}, F is differentiable in its price argument at (y_t, P_t), and g and Γ are differentiable at $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$ and in P_t at (y_t, P_{t-1}, P_t), respectively. Write $B_{t+1} = W/M_{t+1}^2$ and define*

$$\Lambda_{t+1} = \partial_M V_{t+1}(s_{t+1}) - B_{t+1}\, \partial_X V_{t+1}(s_{t+1}), \tag{47}$$

the shadow value of tightening the requirement when the constrained sector remains on $X_{t+1} = W/M_{t+1}$. Then $\nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1})$ satisfies

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + \big(\nabla_y V_{t+1}(s_{t+1})\big) \cdot \partial_P F(y_t, P_t) + \Lambda_{t+1}\, g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t). \tag{48}$$

Proof. See Appendix D.5. □

The first two terms in (48) reflect the fact that P_t enters the next state directly: it is stored as P_t and it updates the internal memory variable $y_{t+1} = F(y_t, P_t)$. The third term isolates the economically relevant channel. Along $X_{t+1} = W/M_{t+1}$, a marginal change dM_{t+1} induces $dX_{t+1} = -B_{t+1}dM_{t+1}$, so the change in continuation value induced by tightening is $dV_{t+1} = \Lambda_{t+1}dM_{t+1}$. The rule-induced incentive to move P_t therefore factors into the statistic sensitivity $\partial_P\Gamma(y_t, P_{t-1}, P_t)$, the local pass-through $g'(\Gamma_t)$, and the shadow value Λ_{t+1} of tightening next-period feasibility. In particular, the margin-update channel is active only when $g'(\Gamma_t)\,\partial_P\Gamma \neq 0$, and its direction is governed by $\operatorname{sign}\big(\Lambda_{t+1}g'(\Gamma_t)\partial_P\Gamma\big)$.

Corollary 3. *Under the conditions of Proposition 4, the margin-update component of ν_{t+1} in (48) is*

$$\nu^M_{t+1} = \Lambda_{t+1}\, g'(\Gamma_t)\, \partial_P\Gamma(y_t, P_{t-1}, P_t).$$

Accordingly, the margin-update channel is locally inactive whenever $g'(\Gamma_t) = 0$, $\partial_P\Gamma(y_t, P_{t-1}, P_t) = 0$, or $\Lambda_{t+1} = 0$. When it is active, its sign is

$$\operatorname{sign}(\nu^M_{t+1}) = \operatorname{sign}\big(\Lambda_{t+1}g'(\Gamma_t)\partial_P\Gamma(y_t, P_{t-1}, P_t)\big).$$

Proof. See Appendix D.6. □

Manipulation through the margin update requires both links in the feedback chain to be operative: the current price must move the statistic, $\partial_P\Gamma \neq 0$, and the statistic must move the requirement, $g'(\Gamma_t) \neq 0$. If either link is flat, marginal changes in P_t do not change M_{t+1} and thus do not change X_{t+1} under $X_{t+1} = W/M_{t+1}$; the attacker then chooses trades only to harvest predetermined flow and to manage inventory given the implied path of (v_τ). When both links are active, the induced sensitivity of next-period constrained demand to the current price is

$$\frac{dX_{t+1}}{dP_t} = -B_{t+1}\, g'(\Gamma_t)\, \partial_P\Gamma(y_t, P_{t-1}, P_t),$$

and Λ_{t+1} converts the resulting marginal tightening into continuation value along the benchmark constraint set. The margin-update component of ν_{t+1} therefore scales with $|g'(\Gamma_t)\partial_P\Gamma|$ and points in the direction determined by $\operatorname{sign}\big(\Lambda_{t+1}g'(\Gamma_t)\partial_P\Gamma\big)$.

5.4 Endgame structure, implementation lags, and closed-form harvesting

The terminal round-trip constraint forces an unwind at the end of the horizon. This endgame discipline rules out indefinitely carrying an induced position and makes late-horizon behavior largely mechanical. Implementation lags strengthen this logic by limiting the set of dates on which a trade can affect within-horizon requirements. At the terminal trade date $t = T - 1$ the unique admissible trade is $u^\star_{T-1} = -x_{T-1}$. Therefore $V_{T-1}(s_{T-1}) = \pi_{T-1}(s_{T-1}, -x_{T-1})$. With an implementation lag, the last dates are left for harvesting and inventory management.

Lemma 4. *Suppose requirements are applied with an implementation lag $\ell \geq 0$, so that $M_{t+1} = g(\Gamma_{t-\ell})$. Then for any date $t \geq T - 1 - \ell$, the action u_t cannot affect any requirements that arrive before the final trade*

date, namely $(M_{t+1}, \dots, M_{T-1})$*, and hence cannot affect any within-horizon forced flows* $(v_{t+1}, \dots, v_{T-1})$*. In particular, from such dates onward the continuation value depends on the state only through inventory, so* $\nu_{t+1} = 0$ *whenever the derivative in Proposition 3 exists.*

Therefore, from date t *onward the continuation problem reduces to harvesting against the predetermined forced-flow sequence and unwinding inventory by* T*, solved in closed form by Proposition 5 when trade bounds are slack (and by Proposition 6 under binding bounds). Equivalently, the last date at which a trade can trigger additional within-horizon forced flow is* $t_{\text{last}} = T - 2 - \ell$*.*

Proof. See Appendix D.7. □

Lemma 4 captures a practical implication of implementation lags. Late trades can still move the measured statistic, but any induced change in posted requirements arrives too late to affect constrained positions within the remaining horizon. From the last trigger date onward, the attacker faces a fixed sequence of predictable imbalance and an obligatory unwind, linking the endgame to optimal execution with exogenous flow under convex impact costs (Almgren and Chriss, 2001; Alfonsi et al., 2010) and to the broader literature on predictable trades and anticipated shocks (Lou et al., 2013; Bessembinder et al., 2016; Fardeau, 2021). More generally, once the remaining forced-flow path $(v_\tau)_{\tau=t}^{T-1}$ is fixed, the continuation problem reduces to intermediation and inventory management. The manipulation motive drops out because within-horizon requirements, and thus within-horizon forced flows, no longer depend on the current print.

Proposition 5 (Closed-form continuation when future forced flow is predetermined)**.** *Fix a date* t *and suppose that, conditional on the state* s_t*, the remaining forced flows* $(v_\tau)_{\tau=t}^{T-1}$ *are predetermined at time* t *and do not depend on the continuation actions* $(u_\tau)_{\tau=t}^{T-1}$*. Assume the per-date bounds are slack at the optimum. Let* $n = T - t$ *denote the number of remaining trading dates and define the average forced flow over the remainder of the horizon,*

$$\bar{v}_t = \frac{1}{n} \sum_{\tau=t}^{T-1} v_\tau.$$

Then the unique maximizer of the continuation problem satisfies, for $\tau = t, \dots, T-1$*,*

$$u_\tau^\star = -\frac{1}{2}(v_\tau - \bar{v}_t) - \frac{x_t}{n}. \tag{49}$$

Moreover, the associated maximal continuation profit is

$$V_t^{\text{harv}} = \frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2 + \eta x_t \bar{v}_t - \frac{\eta}{n} x_t^2. \tag{50}$$

Proof. See Appendix D.8. □

Proposition 5 shows that once the flow of the constrained sector is exogenous to the attacker, she is optimally intermediating a known sequence of imbalances while respecting the terminal unwind. The average level of forced flow does not by itself generate a timing opportunity; the timing rent comes from predictable variation around the average. This is captured by $\frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2$, which vanishes when v_τ is constant over the remaining dates, leaving only mechanical liquidation. The

remaining terms, $\eta x_t \bar{v}_t - \frac{\eta}{n} x_t^2$, summarize the inventory tradeoff: carrying inventory aligned with the average imbalance reduces the net intermediation burden, but outstanding inventory becomes increasingly costly as n shrinks because it must be unwound over fewer remaining dates under quadratic impact.

Corollary 4 (Only dispersion generates timing rents)**.** *In the setting of Proposition 5,*

$$V_t^{\text{harv}} = \frac{\eta}{4} \sum_{\tau=t}^{T-1} \left(v_\tau - \bar{v}_t\right)^2 + \eta x_t \bar{v}_t - \frac{\eta}{n} x_t^2.$$

The only component of V_t^{harv} that depends on the temporal pattern of forced flow beyond its average is the dispersion term $\frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2$. In particular, if $(v_\tau)_{\tau=t}^{T-1}$ is constant over time, say $v_\tau = c$, then

$$u_\tau^\star = -\frac{x_t}{n} \quad \textit{for all } \tau, \qquad V_t^{\text{harv}} = \eta x_t c - \frac{\eta}{n} x_t^2.$$

If, in addition, $x_t = 0$, then $u_\tau^\star = 0$ and $V_t^{\text{harv}} = 0$.

Proof. See Appendix D.9. □

Corollary 5 (A single future liquidation event)**.** *In the setting of Proposition 5, let $n = T - t$ and suppose $x_t = 0$ and forced flow is concentrated at the terminal trade date: $v_\tau = 0$ for $\tau = t, \dots, T-2$ and $v_{T-1} = v < 0$. Then $\bar{v}_t = v/n$ and the optimal trades are*

$$u_\tau^\star = \frac{v}{2n} < 0 \textit{ for } \tau = t, \dots, T-2, \qquad u_{T-1}^\star = -\frac{n-1}{2n}\, v > 0.$$

Moreover, the maximal continuation profit is

$$V_t^{\text{harv}} = \frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2 = \frac{\eta}{4} \frac{n-1}{n} v^2.$$

Proof. See Appendix D.10. □

Holding the forced-flow path fixed, continuation value increases with its temporal dispersion. Design features that bunch rebalancing into a narrow window (e.g., closing auctions, benchmark setting dates etc) therefore raise dispersion and increase the rents available from subsequent intermediation. Corollary 5 illustrates the optimal response to a bunched liquidation: the attacker accumulates inventory ahead of the event and then supplies liquidity on the liquidation date by trading against the forced order. In the full feedback model, early trades can change the forced-flow path by shifting future requirements; once a particular (v_τ) is induced, the remaining problem is exactly the harvesting-and-unwind continuation characterized above.

Finally, per-date speed limits matter for implementation. When $|u_\tau| \le \bar{u}$ binds, the unconstrained intermediation rule (49) cannot be implemented date by date. The optimum then clips the unconstrained target at $\pm\bar{u}$ on the most extreme dates and adjusts the remaining trades to satisfy the terminal inventory constraint.

Proposition 6 (Harvesting with per-date trade bounds). *Fix a date t and suppose that, conditional on s_t, the remaining forced flows $(v_\tau)_{\tau=t}^{T-1}$ are predetermined at time t and do not depend on the continuation trades $(u_\tau)_{\tau=t}^{T-1}$. Consider the continuation problem under temporary impact with per-date bounds $|u_\tau| \le \bar{u}$:*

$$\max_{(u_\tau)_{\tau=t}^{T-1}} \; -\eta \sum_{\tau=t}^{T-1} \left(u_\tau^2 + u_\tau v_\tau\right) \qquad s.t. \qquad \sum_{\tau=t}^{T-1} u_\tau = -x_t, \qquad |u_\tau| \le \bar{u} \;\; \forall \tau.$$

Assume feasibility, $|x_t| \le (T-t)\bar{u}$. Then the problem has a unique maximizer $(u_\tau^\star)_{\tau=t}^{T-1}$. Moreover, there exists a scalar $c_t \in \mathbb{R}$ such that, for $\tau = t, \dots, T-1$,

$$u_\tau^\star = \mathrm{clip}_{[-\bar{u},\bar{u}]}\left(-\frac{v_\tau}{2} + c_t\right), \tag{51}$$

where $\mathrm{clip}_{[-\bar{u},\bar{u}]}(z) = \min\{\bar{u}, \max\{-\bar{u}, z\}\}$ and c_t can be chosen so that $\sum_{\tau=t}^{T-1} u_\tau^\star = -x_t$. If at least one bound is slack at the optimizer (so $|u_\tau^\star| < \bar{u}$ for some τ), then c_t is unique. If the bounds are slack for all τ, then (51) reduces to (49).

Proof. See Appendix D.11. □

6 Multi-asset optimal attacks

Section 5 characterizes the optimal single-asset round trip and shows how trading combines harvesting against predetermined forced flow with state manipulation through the margin rule. This section extends the analysis to N assets with temporary cross-impact under total flow. Cross-impact changes both sides of the problem: it alters the rents from intermediating forced liquidation, and it alters the marginal cost of moving the transaction-price features that feed into the statistic. When the constrained sector rebalances along a fixed portfolio direction b while the risk is a scalar statistic from prices, the attacker can generally use one portfolio to move the statistic and a different portfolio to intermediate the liquidation that the rule induces. This separation has no analogue in the single-asset case, where the same impact coefficient governs both triggering and harvesting. The multi-asset problem therefore links the feedback mechanism to multivariate linear-impact execution with quadratic costs (Alfonsi et al., 2016), to cross-impact restrictions implied by no-dynamic-arbitrage (Schneider and Lillo, 2019), and to order anticipation against forced liquidation (Brunnermeier and Pedersen, 2005).

6.1 The dynamic program

Throughout impose Assumption 5 with N assets and assume A is symmetric and positive semidefinite. Fix a finite horizon T with trading dates $t = 0, 1, \dots, T-1$. Initial conditions are $(P_{-1}, M_0, X_{-1}, X_0)$, where $P_{-1} \in \mathbb{R}^N$ is the last pre-horizon transaction-price vector. As in Section 5.1, we work conditional on a deterministic unaffected-price path $(S_t)_{t=-1}^{T-1}$. At the start of date t, before choosing u_t, the state is

$$s_t = (x_t, y_t, P_{t-1}, M_t, X_t, X_{t-1}), \tag{52}$$

where $x_t \in \mathbb{R}^N$ is the attacker's inventory, (y_t, P_{t-1}) is the information the rule uses to update the next requirement, and (M_t, X_t, X_{t-1}) pin down the constrained sector's holdings and its date-t forced flow. Given (y_t, P_{t-1}) and the current transaction-price vector P_t, the rule updates via

$$y_{t+1} = F(y_t, P_t), \qquad \Gamma_t = \Gamma(y_t, P_{t-1}, P_t), \qquad M_{t+1} = g(\Gamma_t). \tag{53}$$

Under the fully constrained benchmark, the sector holds a fixed portfolio direction $b \in \mathbb{R}^N$ scaled by W/M_t:

$$X_t = \frac{W}{M_t} b. \tag{54}$$

The associated local sensitivity of holdings to the requirement is

$$-\frac{\partial X_t}{\partial M_t} = \frac{W}{M_t^2} b. \tag{55}$$

Thus the constrained sector's date-t net trade is predetermined at the start of date t as

$$v_t = X_t - X_{t-1}. \tag{56}$$

The attacker chooses $u_t \in \mathbb{R}^N$ subject to the per-date speed limit

$$\|u_t\|_1 \le \bar{u}. \tag{57}$$

To ensure feasibility of the terminal round-trip constraint $x_T = 0$, restrict attention to actions that leave enough remaining capacity to unwind:

$$\mathcal{U}_t(x_t) = \left\{ u \in \mathbb{R}^N : \ \|u\|_1 \le \bar{u}, \ \|x_t + u\|_1 \le (T - 1 - t)\bar{u} \right\}. \tag{58}$$

Given u_t, total order flow is $q_t = u_t + v_t$ and transaction prices follow temporary cross-impact under total flow,

$$P_t = S_t + A q_t = S_t + A(u_t + v_t). \tag{59}$$

Using (59), the date-t per-period objective is

$$\pi_t(s_t, u_t) = -u_t^\top (P_t - S_t) = -u_t^\top A(u_t + v_t). \tag{60}$$

The direct impact term $-u_t^\top A u_t$ is weakly negative. The interaction term $-u_t^\top A v_t$ is positive precisely when the attacker trades against the mechanically induced flow, with the relevant notion of against determined by the cross-impact matrix A. Define the transition map $s_{t+1} = \mathcal{T}(s_t, u_t)$

$$\mathcal{T}(s_t, u_t) = \Big(x_t + u_t, \ F(y_t, P_t), \ P_t, \ g(\Gamma_t), \ X_{t+1}, \ X_t \Big), \tag{61}$$

where P_t is given by (59), $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$, and

$$X_{t+1} = \frac{W}{M_{t+1}} b = \frac{W}{g(\Gamma_t)} b.$$

The value function is

$$V_t(s) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \sum_{\tau=t}^{T-1} \pi_\tau(s_\tau, u_\tau), \tag{62}$$

subject to $u_\tau \in \mathcal{U}_\tau(x_\tau)$, the transition map (61), and $x_T = 0$. The associated Bellman recursion is

$$V_T(s) = 0 \text{ if } x_T = 0, \qquad V_T(s) = -\infty \text{ otherwise}, \tag{63}$$

$$V_t(s) = \max_{u \in \mathcal{U}_t(x)} \left\{ \pi_t(s,u) + V_{t+1}(\mathcal{T}(s,u)) \right\}, \qquad t = T-1, \dots, 0. \tag{64}$$

Relative to (41), the control and inventory are vector valued and the impact matrix A determines the trading-cost geometry for both harvesting against forced flow and moving the rule's price inputs.

6.2 Myopic harvesting and Euler equations

The multi-asset problem has the same structure as the single-asset case: gains come from trading against forced flow, while dynamic incentives arise because current trades move both inventory and the transaction-price vector that enters the rule update. As a benchmark, treat v_t as predetermined and hold the continuation value fixed.

Lemma 5 (Myopic harvesting in multiple assets)**.** *Fix a date t and a state s_t with predetermined forced flow $v_t = X_t - X_{t-1}$. Under (60), the one-period payoff satisfies*

$$\pi_t(s_t, u) = -\left(u + \tfrac{v_t}{2}\right)^\top A \left(u + \tfrac{v_t}{2}\right) + \frac{1}{4} v_t^\top A v_t. \tag{65}$$

Consequently, holding the continuation value fixed, $\pi_t(s_t, u)$ is concave in u. The set of unconstrained maximizers is

$$u_t^{\text{myopic}} = -\frac{v_t}{2} + n, \qquad An = 0, \tag{66}$$

and the maximal rent equals $\max_{u \in \mathbb{R}^N} \pi_t(s_t, u) = \frac{1}{4} v_t^\top A v_t$. If A is positive definite, the maximizer is unique and equals $u_t^{\text{myopic}} = -v_t/2$.

Proof. Using (60) and symmetry of A,

$$\pi_t(s_t, u) = -u^\top A(u + v_t) = -u^\top A u - u^\top A v_t.$$

Expanding the quadratic form gives

$$\left(u + \tfrac{v_t}{2}\right)^\top A \left(u + \tfrac{v_t}{2}\right) = u^\top A u + u^\top A v_t + \tfrac{1}{4} v_t^\top A v_t,$$

which rearranges to (65). Since A is positive semidefinite, $z^\top A z$ is convex in z, so $-\left(u + \frac{v_t}{2}\right)^\top A \left(u + \frac{v_t}{2}\right)$ is concave in u. The square term is minimized when $A(u + v_t/2) = 0$, equivalently $u = -v_t/2 + n$ with $An = 0$, and substituting into (65) yields the maximal value. If A is positive definite, then $An = 0$ implies $n = 0$, so the maximizer is unique. □

Lemma 5 is the vector analogue of Lemma 3. When v_t is predetermined, gains arise only from intermediating that flow, and the instantaneous harvesting rent is $\frac{1}{4} v_t^\top A v_t$. Cross-impact affects this rent

through both the impact intensities and the off-diagonal entries of A.

An optimal attack typically deviates from myopic harvesting because the attacker must unwind by T and because current trades move the transaction-price vector that enters the requirement update. The next result extends Proposition 3 by decomposing the interior first-order condition into harvesting, inventory management, and rule manipulation.

Proposition 7. *Fix $t \in \{0, \dots, T-1\}$ and a state s_t. Suppose the optimizer $u_t^\star$ in (64) is interior, $u_t^\star \in \operatorname{int} \mathcal{U}_t(x_t)$, and assume that V_{t+1} is differentiable at the continuation state $s_{t+1}^\star = \mathcal{T}(s_t, u_t^\star)$. Define*

$$\mu_{t+1} = \nabla_x V_{t+1}(s_{t+1}^\star) \in \mathbb{R}^N, \qquad \nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1}^\star) \in \mathbb{R}^N, \qquad P_t = S_t + A(u_t^\star + v_t),$$

where $\frac{d}{dP_t}$ denotes the total derivative through all components of s_{t+1} that depend on P_t under the transition map. Then $u_t^\star$ satisfies the Euler equation

$$A(2u_t^\star + v_t) = \mu_{t+1} + A\,\nu_{t+1}. \tag{67}$$

If A is positive definite, (67) uniquely determines the interior optimizer:

$$2u_t^\star = -v_t + A^{-1}\mu_{t+1} + \nu_{t+1}.$$

If A is only positive semidefinite, the objective and the rule respond only to the price-moving component of trades. Trading in an impact-neutral direction does not move the price under (59) and does not generate cashflows, so such directions are economically irrelevant for harvesting and for triggering. Let $A^\dagger$ denote the Moore–Penrose pseudoinverse and write $\Pi_A = A^\dagger A$ for the operator that extracts the price-moving component of a trade. Then (67) pins down $\Pi_A u_t^\star$, and the canonical choice that sets the impact-neutral component to zero satisfies

$$\Pi_A u_t^\star = -\frac{1}{2}\,\Pi_A v_t + \frac{1}{2} A^\dagger \mu_{t+1} + \frac{1}{2}\,\Pi_A \nu_{t+1}. \tag{68}$$

If A is positive definite, then $A^\dagger = A^{-1}$ and $\Pi_A = I$, so (68) reduces to the previous expression. Moreover, the resulting current-period wedge payoff can be written as

$$\pi_t(s_t, u_t^\star) = \frac{1}{4} v_t^\top A v_t - \Big(u_t^\star + \tfrac{v_t}{2}\Big)^\top A \Big(u_t^\star + \tfrac{v_t}{2}\Big) = \frac{1}{4} v_t^\top A v_t - \frac{1}{4}\big(\mu_{t+1} + A\nu_{t+1}\big)^\top A^\dagger \big(\mu_{t+1} + A\nu_{t+1}\big). \tag{69}$$

Proof. See Appendix E.2. □

The decomposition (68) has the same logic as (45), but cross-impact changes what it means to lean against forced flow and what it means to move prices. On the harvesting side, the attacker is paid only for absorbing the part of the forced imbalance that actually moves execution prices. If a component of v_t does not move transaction prices under the impact matrix A, it cannot generate rents and it does not create a harvesting opportunity. On the inventory side, cross-impact changes which portfolios are cheap to trade. The same marginal value of carrying inventory, μ_{t+1}, translates into smaller position adjustments in directions where execution costs are higher, exactly as in multivariate optimal execution with quadratic costs (Alfonsi et al., 2016). Finally, the manipulation motive is tied to how trades move the transaction-price vector that the rule observes. The continuation-value

sensitivity ν_{t+1} matters only insofar as the attacker can move the relevant prints through impact; directions that do not move prices cannot be used to shift the statistic or tomorrow's requirement.

6.3 Chain-rule decomposition of the rule-manipulation motive

Proposition 7 identifies ν_{t+1} as the continuation-value gain from moving the current transaction-price vector through the rule update. As in the single-asset case, ν_{t+1} can be expressed in terms of the rule primitives (F, Γ, g) and marginal values in the continuation problem.

Proposition 8 (Chain-rule decomposition of ν_{t+1})**.** *Fix $t \in \{0, \dots, T-1\}$ and consider any state-action pair (s_t, u_t). Let P_t be given by (59) and let $s_{t+1} = \mathcal{T}(s_t, u_t)$. Assume V_{t+1} is differentiable at s_{t+1}, F is differentiable in its price argument at (y_t, P_t), and the maps g and Γ are differentiable at $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$ and in P_t at (y_t, P_{t-1}, P_t), respectively. Define*

$$\Lambda_{t+1} = \partial_M V_{t+1}(s_{t+1}) - B_{t+1}^{\top} \nabla_X V_{t+1}(s_{t+1}), \tag{70}$$

the shadow value of tightening the requirement along $X_{t+1} = (W/M_{t+1})b$. Let

$$\nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1}) \in \mathbb{R}^N$$

denote the total derivative of the continuation value with respect to the current print P_t through the transition map. Then

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + \big(\partial_P F(y_t, P_t)\big)^{\top} \nabla_y V_{t+1}(s_{t+1}) + \Lambda_{t+1}\, g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t). \tag{71}$$

Proof. See Appendix E.3. □

Relative to the single-asset case, the key new feature is that the rule may be sensitive to particular price combinations, and cross-impact affects how cheaply the attacker can move those combinations because the print responds to trades through $P_t = S_t + A(u_t + v_t)$.

6.4 The geometry of multi-asset attacks

In multiple assets, the attacker faces two economically distinct objects. First, forced flow arrives in the constrained sector's portfolio direction b, so harvesting rents are earned by trading in ways that load on that liquidation pressure. Second, the rule reacts to a scalar statistic, so the ability to trigger future requirements is governed by the price direction to which the statistic is sensitive, summarized by the gradient $\partial_P \Gamma$. Cross-impact A connects these objects by mapping trades into prices and by determining execution costs. The results below formalize two practical implications: the cheapest way to move the statistic is to trade a particular portfolio, and once the rule marks a fixed portfolio, the optimal attack can be implemented using only two portfolios.

Lemma 6 (Minimal-cost portfolios)**.** *Fix a date t and suppose the statistic is differentiable in the current transaction-price vector with gradient $j_t = \partial_P \Gamma(y_t, P_{t-1}, P_t) \in \mathbb{R}^N$. Under (59), an infinitesimal trade*

perturbation $\delta u \in \mathbb{R}^N$ induces $\delta P_t = A\,\delta u$ and therefore the first-order statistic change

$$\delta\Gamma_t = j_t^\top \delta P_t = j_t^\top A\,\delta u.$$

Assume $j_t^\top A j_t > 0$. Among all perturbations that deliver a given marginal statistic move $\delta\Gamma_t = \delta$, the one with minimal instantaneous impact cost $\delta u^\top A\,\delta u$ is

$$\delta u^\star = \frac{\delta}{j_t^\top A j_t}\, j_t, \qquad \min_{\{\delta u:\ j_t^\top A\,\delta u = \delta\}} \delta u^\top A\,\delta u = \frac{\delta^2}{j_t^\top A j_t}.$$

If $j_t^\top A j_t = 0$, then $A j_t = 0$ and hence $j_t^\top A\,\delta u = 0$ for all δu; the statistic is locally insensitive to trading at date t.

Proof. A marginal trade perturbation δu moves the price by $\delta P_t = A\,\delta u$, so the associated marginal statistic move is the exposure $j_t^\top \delta P_t = j_t^\top A\,\delta u$. Fix a target move $\delta\Gamma_t = \delta$ and consider the least-cost way to achieve it. Under temporary impact, the instantaneous cost of δu is the quadratic form $\delta u^\top A\,\delta u$, so we minimize this cost subject to delivering the required exposure:

$$\min_{\delta u}\ \delta u^\top A\,\delta u \qquad \text{s.t.} \qquad j_t^\top A\,\delta u = \delta.$$

The cheapest way to hit a one-dimensional exposure constraint is to load on the exposure direction itself. Formally, the first-order condition for the constrained minimization implies that the cost gradient $2A\,\delta u$ is proportional to the constraint gradient $A j_t$, so any cost-minimizing perturbation satisfies $A\,\delta u \propto A j_t$. This pins down the price-moving component of δu to be proportional to j_t, and the exposure constraint then fixes the scale. When $j_t^\top A j_t > 0$, the unique price-moving minimizer is

$$\delta u^\star = \frac{\delta}{j_t^\top A j_t}\, j_t,$$

and the corresponding minimal cost is $\delta^2/(j_t^\top A j_t)$. If $j_t^\top A j_t = 0$, then $A j_t = 0$, so trading cannot move the statistic at first order: for any perturbation δu we have $j_t^\top A\,\delta u = (A j_t)^\top \delta u = 0$. □

Lemma 6 isolates the cheapest way to move the rule's input. Only the statistic gradient j_t matters: the rule reacts to the current prices through the single exposure $j_t^\top P_t$, so the lowest-cost way to change Γ_t at the margin is to trade the portfolio aligned with j_t. Cross-impact enters through the scalar $j_t^\top A j_t$, which is the execution-cost intensity of moving that exposure. There is no reason for this trigger portfolio to coincide with the liquidation portfolio b, so in multiple assets the portfolios that generate forced flow and the portfolios that most efficiently move the statistic typically differ. When the rule depends on a portfolio d and its value $d^\top P_t$, one has $j_t \propto d$, and the relevant trigger cost becomes $d^\top A d$.

Theorem 2. *Suppose the margin rule computes its scalar statistic by marking a fixed portfolio direction $d \in \mathbb{R}^N$: there exist functions $\bar{F}$ and $\bar{\Gamma}$ such that*

$$y_{t+1} = \bar{F}(y_t, d^\top P_t), \qquad \Gamma_t = \bar{\Gamma}(y_t, d^\top P_{t-1}, d^\top P_t),$$

for all relevant t. Consider the multi-asset dynamic program in Section 6.1 under temporary total-flow impact (59) and constrained demand (54). Consider the relaxed problem in which the per-date trade bound (57) (and hence the admissible set (58)) is removed but the terminal round-trip constraint $x_T = 0$ is retained. Then the supremum value of the relaxed problem is unchanged when restricting attention to attacks $(u_t)_{t=0}^{T-1}$ that trade only the two economically relevant portfolios,

$$u_t \in \operatorname{span}\{b, d\} \qquad \text{for all } t.$$

In particular, if the supremum is attained then there exists an optimal attack $(u_t^\star)_{t=0}^{T-1}$ satisfying this span restriction. Moreover, once trades are restricted to $\operatorname{span}\{b, d\}$, cross-impact matters only through three scalar cost coefficients:

$$b^\top A b, \qquad b^\top A d = d^\top A b, \qquad d^\top A d.$$

These numbers summarize, respectively, the execution-cost intensity of trading the liquidation portfolio b, the execution-cost intensity of trading the marked portfolio d, and the cross-impact interaction between the two portfolios. When $N = 1$ this reduction collapses to the single-asset model.

Proof. See Appendix E.4. □

Theorem 2 says that, when the rule marks a fixed portfolio d, the attacker never needs to trade a large menu of assets. Two portfolios span everything that matters. The constrained sector is forced to rebalance only in the liquidation portfolio b, so the only immediate profits available for harvesting come from trading in ways that absorb that b-direction order imbalance. The rule, on the other hand, reacts only to the portfolio price $d^\top P_t$, so the only way to trigger or reshape future requirements is to trade in ways that move $d^\top P_t$. Trading any portfolio component outside the span of $\{b, d\}$ does not help absorb the forced b-imbalance, and does not move the portfolio price the rule uses. Once attention is restricted to combinations of b and d, cross-impact affects outcomes only through three transparent scalars:

$$b^\top A b, \qquad d^\top A d, \qquad b^\top A d.$$

These summarize, respectively, how costly it is to intermediate liquidation in portfolio b, how costly it is to move the marked prices in portfolio d, and how much trading in one portfolio moves the other through cross-impact.

7 Volatility-controlled indices in structured products

Volatility-controlled (also called risk-control, volatility-managed, or target-volatility) indices are a large, rule-driven segment of asset markets, and they embody the design problem at the center of this paper. The underlying methodology is publicly disclosed and mechanically implemented: a short return history (often 20 days) is mapped into a risk statistic (volatility) and then into next-day exposure, often with caps, floors, smoothing, and implementation lags (Krein and Fernandez, 2012; Moreira and Muir, 2017; MSCI, 2021; S&P Dow Jones Indices, 2025). These indices are embedded in retail structured products and index-linked annuities, so the mapping from sampled transaction prices into mandated reallocations is a contractual trading rule. As a result, when realized volatility

rises and exposure needs to be reduced, replication requires predictable rebalancing in the hedge instruments (usually ETFs), and when volatility falls the rule predictably re-levers.

This setting is interesting because volatility-control indices are both large in practice and unusually transparent in their design. The methodology is disclosed and fixes the mapping from transaction prices into the risk statistic and then into next-day exposure. The remaining input is market impact in the hedge instruments used for replication; no holdings, dealer inventories, or ownership data are needed to test whether the sector is potentially manipulable. Since the market impact of the volatility-managed complex grows with its size, a candidate methodology together with an impact calibration implies an explicit capacity bound via Theorem 1, namely the largest linked notional for which the induced feedback mechanism remains manipulation-free in the linearized system.

7.1 Product types

Typically, three parties are involved in the volatility-managed complex: the index sponsor (and calculation agent) publishes the methodology and the resulting index levels and weights; the issuer or insurer sells a contract whose payoff references the index, so the liability inherits that methodology mechanically; and a hedge desk replicates the liability using liquid instruments (ETFs, futures, swaps, and cash). When the index rebalances, replication requires a mechanical rebalancing in the hedging instruments. Appendix F details the institutional setting and maps a published methodology into the rule objects used in our computations; the institutional descriptions follow American Academy of Actuaries (2026b); MSCI (2021); S&P Dow Jones Indices (2025); BlackRock (2023). Most marketed designs share the same blueprint: a risky reference portfolio is held alongside cash or a bond-like instrument, and a volatility-control rule adjusts the risky weight to keep risk near a stated target. We focus on two types that capture the feedback channel in our model:

1. *Single underlying (Template A).* The reference portfolio is a single parent index. The risky weight is a capped inverse-volatility function of a rolling volatility estimate (MSCI, 2021; S&P Dow Jones Indices, 2025).

2. *Multi-underlying, fixed basket (Template B).* The reference portfolio is a fixed multi-asset basket (often implemented with ETFs or futures). A scalar volatility-control factor scales the entire basket relative to cash (BlackRock, 2023; S&P Dow Jones Indices, 2025).

Some products add an additional layer: the reference portfolio itself changes over time based on signals (for example, trend or rotation rules), and the same volatility-control rule is then applied to the selected portfolio. This adds state variables (the signals), but the price-to-demand feedback studied here still operates.

7.2 Optimal attacks against pure target-volatility rules

To gain intuition for how volatility-managed portfolios map into the results in the rest of the paper, this section specializes the finite-horizon single-asset dynamic program in Section 5.1 to the interior region of a pure target-volatility methodology. In this region the exposure schedule is smooth, so the price-to-requirement mapping is differentiable. We therefore abstract from caps, floors, and turnover

buffers. The risky weight is the target volatility divided by a rolling realized-volatility estimate, evaluated with an implementation lag ℓ:

$$w_{t+1} = \frac{\sigma^*}{\widehat{\sigma}_{t-\ell}}, \qquad M_{t+1} = \frac{1}{w_{t+1}} = \frac{\widehat{\sigma}_{t-\ell}}{\sigma^*} = \frac{\sqrt{\Gamma_{t-\ell}}}{\sigma^*}, \qquad s_t = \frac{\partial M_{t+1}}{\partial \Gamma_{t-\ell}} = \frac{1}{2\sigma^* \widehat{\sigma}_{t-\ell}}. \tag{72}$$

Replication in the hedge instrument implies that the volatility-managed sector holds $X_{t+1} = W/M_{t+1}$ and generates forced flow $v_{t+1} = X_{t+1} - X_t$. Prices obey temporary impact, $P_t = S_t + \eta(u_t + v_t)$, and the attacker earns the payoff (37). The attacker harvests predictable rebalancing by trading against v_t, while also valuing trades that shift the current price because it enters future requirements and thereby reshapes future forced flow. At an interior optimum, Proposition 3 implies the decomposition

$$u_t^\star = -\frac{v_t}{2} + \frac{1}{2\eta}\mu_{t+1} + \frac{1}{2}\nu_{t+1}, \tag{73}$$

with (μ_{t+1}, ν_{t+1}) defined there. The first term is the myopic harvesting benchmark from Lemma 3. The second term reflects inventory management induced by the terminal round-trip constraint $x_T = 0$. The third term is the feedback motive: it values moving the current price because it shifts the rule input and therefore changes future weights and flows. Proposition 4 expresses ν_{t+1} through the rule derivatives. In a target-volatility methodology, the objects entering that chain rule are the schedule slope, given by (72), the state update F, given by the deterministic shift associated with a rolling window, and the statistic derivative with respect to the current print, which is available in closed form. As a result, the manipulation motive, and hence the optimal policy, can be computed directly from the disclosed methodology.

For Template A, the statistic is rolling realized variance computed from log returns,

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right), \qquad \Gamma_t = \frac{a_{\text{ann}}}{m}\sum_{i=0}^{m-1} r_{t-i}^2. \tag{74}$$

A minimal Markov state for the rolling window is the $(m–1)$-vector of past returns $y_t = (r_{t-1}, \ldots, r_{t-m+1})$, which updates by a deterministic shift once the new price P_t pins down r_t. A hard m-day window requires tracking which returns are about to roll out. This bookkeeping is crucial because roll-off mechanically changes $\widehat{\sigma}$ and hence next-day exposure even absent new shocks, generating predictable components of future forced flow that the attacker can anticipate. We discuss this in the next subsection.

7.2.1 A predictable flow reversal

Fix a date j and consider a perturbation that changes r_j while leaving all other returns unchanged. Because (74) is a rolling average of squared returns, the term r_j^2 enters Γ_t for every t whose m-day window contains j. Holding other returns fixed,

$$\delta\Gamma_t = \frac{a_{\text{ann}}}{m}\,\delta(r_j^2)\cdot \mathbf{1}\{t \in \{j, \ldots, j+m-1\}\}. \tag{75}$$

The derivative drops to zero at $t = j+m$ when r_j rolls out. With implementation lag ℓ, this block shift in Γ generates a block of tighter applied requirements: it affects M_{t+1} when $t-\ell \in \{j, \dots, j+m-1\}$, i.e. for dates $t+1 \in \{j+\ell+1, \dots, j+\ell+m\}$, and it disappears one day later. Since forced flow is the first difference $v_{t+1} = X_{t+1} - X_t$ with $X_{t+1} = W/M_{t+1}$, the rule induces a predictable entry-exit pattern: deleveraging when the tightened block starts to apply, followed by releveraging when it ends (shifted by the baseline one-day delay and the disclosed lag ℓ). This is a distinctive feature of hard rolling windows and a transparent instance of backward-looking procyclicality (Glasserman and Wu, 2018; BIS, 2010; ESRB, 2020). With exponentially weighted or infinite-memory estimators, the same logic applies, but the discrete roll-off is replaced by gradual decay.

7.2.2 The fragility of low-volatility states

The key object for manipulability is the marginal pass-through from a sampled transaction price into next-day forced flow. In the interior region of a target-volatility rule, $g(\Gamma) = \sqrt{\Gamma}/\sigma^*$, so the local schedule slope $s_t = g'(\Gamma_{t-\ell})$ in (72) is steepest when the measured variance $\Gamma_{t-\ell}$ is small. The linked sector holds

$$X_{t+1} = \frac{W}{M_{t+1}} = \frac{W\sigma^*}{\sqrt{\Gamma_{t-\ell}}},$$

which implies the volatility-managed demand sensitivity

$$B_{t+1} = \frac{W}{M_{t+1}^2} = \frac{W\sigma^{*2}}{\Gamma_{t-\ell}}, \qquad B_{t+1}s_t = \frac{W\sigma^*}{2\,\Gamma_{t-\ell}^{3/2}}.$$

The product $B_{t+1}s_t$ is the mechanical gain that maps a marginal change in the risk statistic into a marginal tightening of requirements and, in binding states, into marginal forced rebalancing. Its $\Gamma^{-3/2}$ scaling implies that the rule is locally most responsive in low-variance states.

The attacker does not choose Γ directly; she trades to move transaction prices, which move returns and therefore measured risk. The marginal effect of a transaction price on next-day forced flow is therefore proportional to

$$\underbrace{B_{t+1}s_t}_{\text{pass-through into demand}} \times \underbrace{\partial_P \Gamma}_{\text{sensitivity of the measured input to prices}}.$$

For rolling-window variance, $\partial_P\Gamma$ is itself largest after quiet spells: when recent dispersion is low, a given marginal perturbation to the transaction price produces a larger marginal change in the variance estimate. Hence tranquil configurations are the most manipulable. They combine large baseline exposure, so a given tightening translates into larger dollar flow, a steep schedule, so the same statistic move generates a larger change in next-day requirements, and a sensitive estimator, so the statistic is easier to shift. These are not crisis states in the standard sense; they are precisely the states in which the methodology is designed to take the most risk and to react most strongly to the measured input. In such states, a small attack that shifts the statistic can deliver high returns per unit of impact cost once the induced reversal is harvested.

This logic is related to, but distinct from, the leverage-cycle narrative. In leverage-cycle models,

low measured risk relaxes financing terms, balance sheets expand, and fragility follows from large positions (Basak and Shapiro, 2001; Brunnermeier and Pedersen, 2009; Geanakoplos, 2010; Adrian and Shin, 2010). Here the sector is indeed largest when measured risk is low, but even holding sector size fixed, fragility is amplified by two additional channels: the requirement schedule is locally steepest and the risk statistic is locally most sensitive to prices. These features are typically absent from standard balance-sheet-based amplification mechanisms. This is the rule-based "paradox of financial stability" applied to a concrete price-to-trade mapping (Borio and Drehmann, 2009; Nicolai, 2026).

For volatility-managed strategies, the point is sharper than the statement that they de-risk after spikes: they deliberately lever up after quiet spells, and those are exactly the configurations in which the first-order mapping from transaction prices to future forced flow is most sensitive. This is why, in our admissibility and capacity calculations, binding constraints are often pinned down by low-volatility, high-exposure states rather than only by conventional stress episodes. Operationally, the test is therefore a forward-looking stress test for benign states: given a disclosed methodology and standard impact inputs, one can quantify how close low-volatility configurations are to violating admissibility, even when backward-looking risk measurement labels them as safe.

7.2.3 Inventory-light volatility triggering

For any calendar date j,

$$\frac{\partial \Gamma_t}{\partial P_j} = \frac{2a_{\text{ann}}}{m\,P_j}\Big(\mathbf{1}\{t-m+1 \le j \le t\}\, r_j - \mathbf{1}\{t-m \le j \le t-1\}\, r_{j+1}\Big), \tag{76}$$

where $r_j = \log(P_j/P_{j-1})$. The key implication is that a single transaction price P_j enters two adjacent returns with opposite signs: it increases r_j and decreases r_{j+1}. Since Γ_t aggregates squared returns, this adjacent-returns channel allows an attacker to move the volatility input using short-lived variation in prices, even if the price level is quickly restored.

To see the mechanism, consider a two-day buy-then-sell disturbance that generates returns $+a$ on day j and $-a$ on day $j+1$, so the price level is largely reversed over two days. Holding the rest of the window fixed, the incremental contribution of these two returns to Γ_t is approximately

$$\delta\Gamma_t \approx \frac{a_{\text{ann}}}{m}a^2\,\mathbf{1}\{t \in \{j,\dots,j+m-1\}\} + \frac{a_{\text{ann}}}{m}a^2\,\mathbf{1}\{t \in \{j+1,\dots,j+m\}\}.$$

Both large-magnitude returns enter the rolling window, and both increase the statistic because the mapping from returns to Γ_t is convex. The attacker therefore pays execution cost to create a volatility spike, but can reverse the price level quickly, keeping net inventory small relative to the gross trading used to generate the two returns. In the dynamic program this matters because continuation value can be increased by shifting Γ without carrying a large directional position.

This channel is distinct from standard predatory-trading logic (Brunnermeier and Pedersen, 2005; Carlin et al., 2007; Attari et al., 2005). In those models, the order imbalance is driven by an external balance-sheet shock, and the attacker positions ahead of, and intermediates through, a liquidation episode whose timing and magnitude are taken as given. The limiting friction is typically inventory

risk and funding, because profits require warehousing exposure while the distressed trader unwinds. Here, by contrast, the ability to generate a volatility spike while keeping directional inventory small is exactly what makes the attack cheaper and easier to execute.

7.2.4 The effect of rule parameters: lags, horizons and target volatilities

Disclosed target-volatility methodologies often impose an implementation lag ℓ. In the dynamic program, Lemma 4 implies that if $M_{t+1} = g(\Gamma_{t-\ell})$, then for any $t \geq T-1-\ell$ a date-t trade cannot affect any requirements, and hence any forced flows, that arrive before the terminal date. In this endgame region the feedback motive disappears (formally, $\nu_{t+1} = 0$ whenever the derivative in Proposition 3 exists), and the continuation problem reduces to intermediation and inventory unwind against a predetermined forced-flow sequence, with the closed-form solution in Proposition 5. Equivalently, the last date at which a trade can still trigger additional within-horizon forced flow is $t_{\text{last}} = T-2-\ell$. Under the common $\ell = 1$ convention in target-volatility indices, the final two dates are therefore governed purely by harvesting and inventory unwind, with no incentive to move the statistic.

Window length m enters through the statistic Jacobian (76). A longer window attenuates the marginal effect of any single return on Γ through the explicit $1/m$ scaling, but it also extends the period over which a given return remains in the window, delaying roll-off and stretching the induced reversal. These forces shift optimal triggering earlier in the horizon and increase the value of having sufficient remaining time to harvest the eventual exit.

The target volatility σ^* scales requirements through $M_{t+1} = \widehat{\sigma}_{t-\ell}/\sigma^*$ and steepens the local pass-through via $s_t = 1/(2\sigma^*\widehat{\sigma}_{t-\ell})$. Holding linked notional W fixed, the resulting gain in the requirement-to-flow channel is

$$B_{t+1}s_t = \frac{W\sigma^*}{2\,\Gamma_{t-\ell}^{3/2}},$$

which increases in σ^*: higher targets mechanically raise exposure and amplify the marginal flow response to a given perturbation of the statistic.

7.3 Testing vulnerability

Fix a market state z, a linked notional W, and a disclosed methodology. These inputs pin down the augmented symmetric cost matrix $\widehat{H}(z;W)$. By Theorem 1, the economy is free of dynamic arbitrage at $(z;W)$ if and only if $\widehat{H}(z;W)$ is positive semidefinite on the round-trip subspace. This yields the statewise capacity bound

$$W_{\max}(z) = \sup\Big\{W \geq 0 :\ \widehat{H}(z;W) \text{ is positive semidefinite on round trips}\Big\}.$$

Given a stress set of states $\mathcal{Z}$, we report the conservative bound

$$W_{\max}(\mathcal{Z}) = \inf_{z\in\mathcal{Z}} W_{\max}(z),$$

and the vulnerability curve, which for each candidate scale W reports the fraction of states $z \in \mathcal{Z}$ at which admissibility fails, i.e. $\Pr_{z\in\mathcal{Z}}[W > W_{\max}(z)]$.

We construct $\mathcal{Z}$ from simulated price paths rather than from holdings or order flow. This matches the objective: an ex ante test that can be run before a product reaches scale, using public rule information and standard liquidity inputs. The disclosed methodology fixes the mapping from recent returns to next-period exposure; execution costs enter only through an externally calibrated impact model under current market conditions. Appendix F details the simulation design and the impact calibration used for the reported results.

7.3.1 Results

We implement the admissibility test for the two volatility-control templates (A and B). We consider attacker horizons of 3, 6, and 12 months, i.e. $T \in \{63, 126, 252\}$ trading days. For each parameter configuration we generate a stress set $\mathcal{Z}$ of $|\mathcal{Z}| = 400$ simulated market states (Appendix F.4). In each state $z \in \mathcal{Z}$ we compute the statewise capacity bound $W_{\max}(z)$, expressed in days of ADV of the risky leg. For a candidate sector scale W, admissibility fails in state z whenever $W > W_{\max}(z)$.

Table 1 summarizes baseline magnitudes at $T = 126$. For Template A (single-asset rule), the conservative capacity is $W_{\max}(\mathcal{Z}) = 0.165$ days of ADV, while the median across stress states is 0.760. For Template B (portfolio rule), the conservative capacity is 0.125 for $N = 2$, 0.122 for $N = 4$, and 0.230 for $N = 8$. Vulnerability rises quickly at economically relevant scales: at $W = 1$ day of ADV, the test fails in 66.8% of stress states for Template A and in 78.8% to 94.0% for Template B; at $W = 2$ days of ADV, failure rates rise to 94.5% to 99.5%.

To give these magnitudes an economic interpretation, note that volatility-managed strategies often implement a dynamic allocation between a liquid ETF and a low-risk asset such as cash, Treasury bills, or a bond-like security. Khomyn et al. (2024) report that the median U.S. ETF in their sample has an ETF turnover ratio of 2.17, where turnover is defined as annualized daily dollar volume in the ETF's secondary market divided by AUM. This implies a daily ADV equal to $2.17/252 = 0.86\%$ of net assets. Therefore, a capacity of 0.165 days of ADV implies that the volatility-managed complex becomes manipulable once its size reaches about $0.165 \times 0.86\% = 0.14\%$ of the underlying ETF's net assets. Likewise, a size of $W = 2$ days of ADV corresponds to about $2 \times 0.86\% = 1.72\%$ of the underlying ETF's net assets. At that scale, the volatility-managed complex is almost always susceptible to this form of manipulation.

Figures 2 and 3 plot vulnerability curves for Template A and for Template B with $N = 2$, for $T \in \{63, 126, 252\}$. These curves complement the worst-state bound $W_{\max}(\mathcal{Z})$ by showing how failure probability varies with sector scale. Lower target volatility, $\sigma^* = 10\%$, increases $W_{\max}(\mathcal{Z})$ (Table 2) and shifts the vulnerability curve downward at a given W relative to baseline, while weaker liquidity, $c_{1\%} = 20$ bps, shifts curves upward. In Template B, changes in σ^* can move the strategy in and out of the leverage cap in a subset of stress states, so conservative capacity need not be monotone in σ^* (Appendix H, Table 8). Appendix H reports $W_{\max}(\mathcal{Z})$ for all scenarios and horizons (Tables 4–7) and the corresponding vulnerability curves for $N \in \{4, 8\}$ (Figures 4–5).

Table 2 varies the target volatility σ^* and the impact calibration, expressed as the cost in bps of trading 1% of volume, $c_{1\%}$. Lowering the target to $\sigma^* = 10\%$ raises conservative capacity by 25% in both templates relative to baseline. Raising the target to $\sigma^* = 15\%$ reduces capacity in Template A, while leaving capacity essentially unchanged in Template B with $N = 2$ because higher target scaling

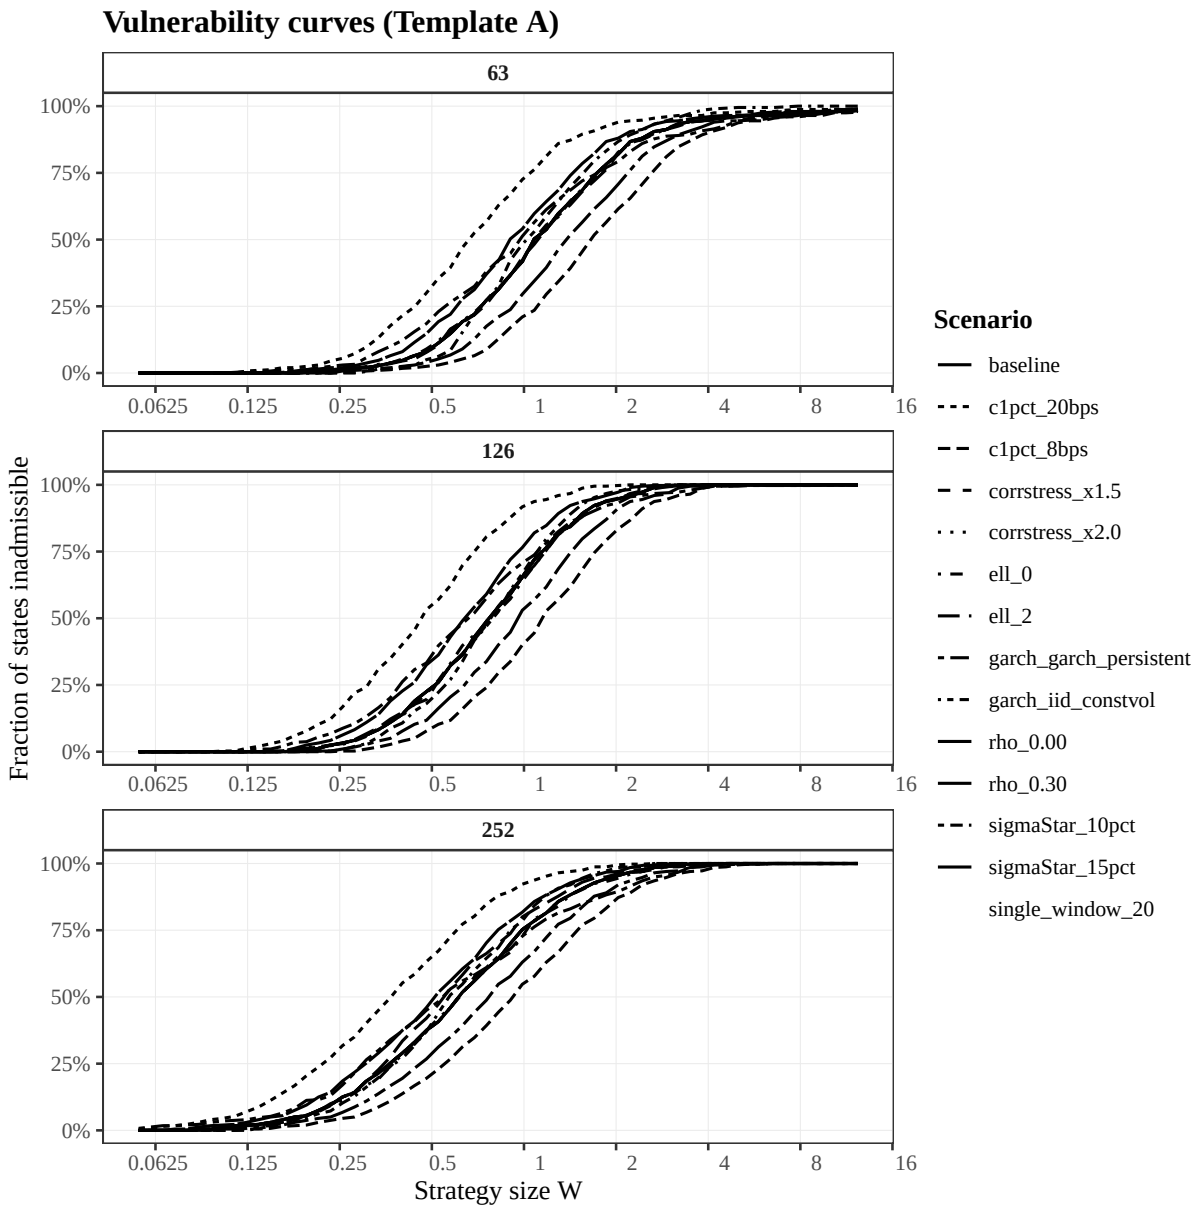

Figure 2 Vulnerability curves for Template A.

Panels correspond to $T \in \{63, 126, 252\}$. In each panel, the curve for a given scenario plots $\Pr_{z \in \mathcal{Z}}[W > W_{\max}(z)]$ across the $|\mathcal{Z}| = 400$ stress states as a function of sector scale W (in days of ADV). Line types correspond to one-at-a-time deviations from the baseline configuration: baseline ($\sigma^* = 12.5\%$, $c_{1\%} = 12$ bps, lag $\ell = 1$, windows $\{20, 60\}$ with max aggregation, and garch_moderate); target-volatility shifts sigmaStar_10pct and sigmaStar_15pct; liquidity shifts c1pct_8bps and c1pct_20bps; lag shifts ell_0 and ell_2; volatility window restriction single_window_20; and different GARCH models garch_iid_constvol and garch_persistent.

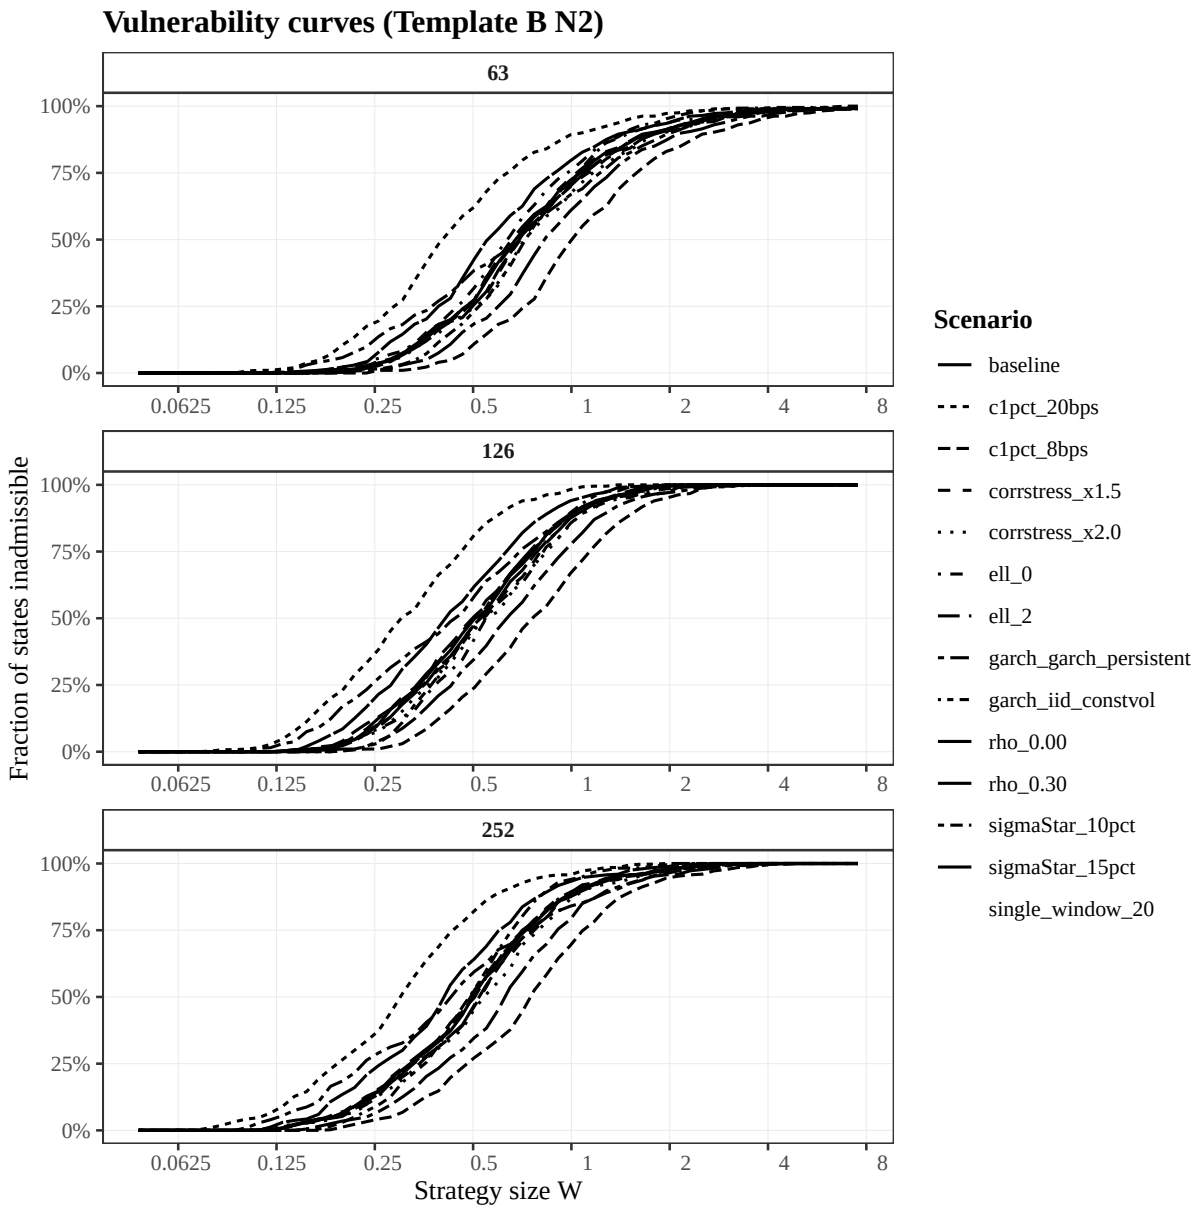

Figure 3 Vulnerability curves for Template B with $N = 2$.

Panels correspond to $T \in \{63, 126, 252\}$. In each panel, the curve for a given scenario plots $\Pr_{z \in \mathcal{Z}}[W > W_{\max}(z)]$ across the $|\mathcal{Z}| = 400$ stress states as a function of sector scale W (in days of ADV). Line types correspond to one-at-a-time deviations from the baseline configuration: baseline ($\sigma^* = 12.5\%$, $c_{1\%} = 12$ bps, lag $\ell = 1$, windows $\{20, 60\}$ with max aggregation, and garch_moderate); target-volatility shifts sigmaStar_10pct and sigmaStar_15pct; liquidity shifts c1pct_8bps and c1pct_20bps; lag shifts ell_0 and ell_2; volatility window restriction single_window_20; and different GARCH models garch_iid_constvol and garch_persistent.

Table 1 Baseline admissibility and vulnerability at horizon $T = 126$.

$W_{\max}(\mathcal{Z}) = \min_{z\in\mathcal{Z}} W_{\max}(z)$ is the conservative admissible scale (worst state in the stress set). $W_{0.05}$ is the 5th percentile of $W_{\max}(z)$ across $z \in \mathcal{Z}$, and "Median" is the median of $W_{\max}(z)$. Probabilities report the fraction of stress states with $W_{\max}(z) < W$ evaluated at $W = 1$ and $W = 2$, i.e., at one and two days of ADV.

		$W_{\max}(\mathcal{Z})$	$W_{0.05}$	Median($W_{\max}$)	$\Pr[W_{\max} < 1]$	$\Pr[W_{\max} < 2]$
Template A	$N = 1$	0.165	0.293	0.760	66.8%	94.5%
Template B	$N = 2$	0.125	0.218	0.505	88.8%	99.5%
Template B	$N = 4$	0.122	0.238	0.502	94.0%	98.5%
Template B	$N = 8$	0.230	0.408	0.716	78.8%	96.8%

Table 2 Comparative statics at horizon $T = 126$.

Comparative statics for $T = 126$: target volatility σ^* and liquidity $c_{1\%}$. Entries report $W_{\max}(\mathcal{Z})$; parentheses report the ratio relative to the baseline within template.

Scenario	Template A ($N = 1$)	Template B ($N = 2$)
baseline	0.165 (1.00x)	0.125 (1.00x)
$\sigma^* = 10\%$	0.206 (1.25x)	0.156 (1.25x)
$\sigma^* = 15\%$	0.137 (0.83x)	0.128 (1.03x)
$c_{1\%} = 8$ bps	0.247 (1.50x)	0.188 (1.50x)
$c_{1\%} = 20$ bps	0.099 (0.60x)	0.075 (0.60x)

increases the incidence of the leverage cap in a subset of stress states (Appendix H, Table 8). Liquidity shifts are close to proportional: improving impact from $c_{1\%} = 12$ bps (baseline) to 8 bps increases $W_{\max}(\mathcal{Z})$ by 50%, whereas deteriorating to $c_{1\%} = 20$ bps reduces $W_{\max}(\mathcal{Z})$ by 40%.

Table 3 highlights how capacity varies with horizon and selected feedback stresses. In Template A, conservative capacity declines with horizon in the baseline case, from 0.191 at $T = 63$ to 0.081 at $T = 252$, and more persistent volatility dynamics further tighten the long-horizon bound. In Template B, diversification and cross-impact jointly matter for capacity: for $N = 8$ at $T = 126$, increasing cross-impact from $\rho = 0$ to $\rho = 0.3$ lowers $W_{\max}(\mathcal{Z})$ from 0.287 to 0.209. Appendix H reports the full grid of horizons and scenarios and provides diagnostic plots for the cross-state distribution of $W_{\max}(z)$ and related quantities.

The magnitudes in these exercises imply that admissible scale is typically measured in fractions of ADV rather than in multiple days. Conservative bounds $W_{\max}(\mathcal{Z})$ are often around 0.1 to 0.3 days of ADV across templates and scenarios (Tables 1 and 3). Consistent with this, vulnerability rises quickly at economically large scales: for $W \in \{1, 2\}$ days of ADV, the test fails in a large share of stress states even under baseline configurations (Table 1 and Figures 2–3). A second implication is that tail states drive supervisory risk. Cross-state medians can be several times larger than worst-state bounds (Table 1), so diagnostics based on typical conditions would materially overstate safe capacity. A third implication is that liquidity is first order determinant of the likelihood of success of an attack: changes in the impact calibration $c_{1\%}$ shift capacity close to proportionally (Table 2), so credible use of the test requires updating the calibration as market conditions evolve. Finally, implementation details matter: target scaling interacts with caps and other nonlinear features so that capacity need

Table 3 Horizon and feedback-stress summary for conservative admissible scale $W_{\max}(\mathcal{Z})$. Entries are $W_{\max}(\mathcal{Z}) = \min_{z\in\mathcal{Z}} W_{\max}(z)$ in ADV-normalized units. Selected scenarios are shown; Appendix H reports $W_{\max}(\mathcal{Z})$ for all scenarios and horizons.

Scenario	$T = 63$	$T = 126$	$T = 252$
Template A: baseline	0.191	0.165	0.081
Template A: $\sigma^* = 10\%$	0.239	0.206	0.101
Template A: $c_{1\%} = 20$ bps	0.115	0.099	0.049
Template A: GARCH persistent	0.110	0.120	0.035
Template B ($N = 2$): baseline	0.158	0.125	0.115
Template B ($N = 2$): $\sigma^* = 10\%$	0.197	0.156	0.144
Template B ($N = 2$): $c_{1\%} = 20$ bps	0.095	0.075	0.069
Template B ($N = 8$): baseline	0.311	0.230	0.286
Template B ($N = 8$): cross-impact $\rho = 0$	0.395	0.287	0.360
Template B ($N = 8$): cross-impact $\rho = 0.3$	0.281	0.209	0.259

not move monotonically with σ^* in portfolio implementations (Table 2 and Appendix H, Table 8), and cross-impact can materially tighten capacity in diversified versions (Table 3).

8 Conclusion

Price-based risk constraints map sampled transaction prices into next-period requirements such as margins, haircuts, leverage limits, and mandated exposures. Once these rules bind for a sufficiently large sector, they make part of future order flow mechanically predictable. This changes the relevant no-arbitrage benchmark. Classical no-manipulation restrictions for market impact are no longer sufficient, because current trades move current prices but also change the sampled statistic used by the rule, alter tomorrow's requirement, and thereby reshape tomorrow's forced trading. A market can therefore satisfy the standard impact-side no-manipulation condition and still admit profitable round trips.

The paper's main result is to derive the correct no-manipulation condition for this environment and to express it in a form that can be implemented. Theorem 1 shows that the relevant condition must account for both the trader's direct price impact and the predictable rebalancing induced by the constrained sector's rule. This yields an ex ante stress test for any disclosed mechanical strategy. Given a rule, an impact calibration, and an assumed sector scale, the test determines whether profitable round trips exist, how often they arise across states, and how large the mechanical sector can become before those opportunities appear. Looking at one state at a time, the test gives the largest sector size that is still safe in that particular state. Looking across the full set of states, it also gives a single conservative safe size: the largest sector size that remains safe even in the most fragile states we consider. The same exercise yields a vulnerability curve showing, for any proposed sector size, the share of states in which the no-manipulation condition fails. This makes the test useful for a regulator seeking to limit the size of a sector that employs mechanical trading rules.

We apply the test to volatility-managed indices. The relevant scales are modest. Worst-case admissible scale is small relative to market liquidity, and the conservative bound is on the order of

fractions of ADV. At a collective size of 1 day of ADV, the no-manipulation condition already fails in between 66.8% and 94.0% of states, and at a size of 2 days of ADV it fails in 94.5% to 99.5% of states. Public price-based rules therefore imply a relatively small safe scale, and once that scale is exceeded, manipulability becomes a regular feature of the environment.

More broadly, the paper shows that whenever a disclosed rule maps past transaction prices into future binding demand, market safety cannot be assessed from impact restrictions alone. One must also account for the endogenous order flow generated by the rule itself. For mechanical strategies that operate at scale, admissibility is therefore a joint property of the trading rule, market liquidity, and sector size. That is the relevant object for design, supervision, and market-scale assessment. Beyond volatility-managed portfolios, these methods apply whenever markets are not perfectly elastic and a significant fraction of investors follow mechanical trading rules that respond to past prices, a common feature of modern markets.

References

Acharya, Viral, Robert Engle, and Matthew Richardson, "Capital Shortfall: A New Approach to Ranking and Regulating Systemic Risks," *American Economic Review*, 2012, *102* (3), 59–64.

Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon, and Matthew Richardson, "Measuring Systemic Risk," *The Review of Financial Studies*, 2017, *30* (1), 2–47.

Adrian, Tobias and Hyun Song Shin, "Liquidity and leverage," *Journal of Financial Intermediation*, 07 2010, *19* (3), 418–437.

_ and Markus K. Brunnermeier, "CoVaR," *American Economic Review*, 2016, *106* (7), 1705–1741.

Alfonsi, Aurelien, Alexander Schied, and Florian Klöck, "Multivariate transient price impact and matrix-valued positive definite functions," *Mathematics of Operations Research*, 2016, *41* (3), 914–934.

_ , Antje Fruth, and Alexander Schied, "Optimal execution strategies in limit order books with general shape functions," *Quantitative Finance*, 2010, *10* (2), 143–157.

Almgren, Robert and Neil Chriss, "Optimal execution of portfolio transactions," *Journal of Risk*, 2001, *3* (2), 5–39.

American Academy of Actuaries, "Fixed Indexed Annuities: Product Mechanics and Risk Management," Technical Report, American Academy of Actuaries February 2026.

_ , "Fixed Indexed Annuities: Product Mechanics and Risk Management," February 2026. Life Experience Committee report.

Attari, Mukarram, Antonio S. Mello, and Martin E. Ruckes, "Arbitraging Arbitrageurs," *The Journal of Finance*, October 2005, *60* (5), 2471–2511.

Barclays Bank, "Pricing Supplement (Form 424B2): Risk Control Underlier Disclosure," August 2021. SEC EDGAR filing.

Basak, Suleyman and Alexander Shapiro, "Value-at-Risk-Based Risk Management: Optimal Policies and Asset Prices," *The Review of Financial Studies*, 2001, *14* (2), 371–405.

Bessembinder, Hendrik, Allen Carrion, Laura Tuttle, and Kumar Venkataraman, "Liquidity, Resiliency and Market Quality Around Predictable Trades: Theory and Evidence," *Journal of Financial Economics*, 2016, *121* (1), 142–166.

BIS, "The Role of Margin Requirements and Haircuts in Procyclicality," Technical Report 36, Bank for International Settlements 03 2010.

_ , "Review of Margining Practices," Technical Report, Bank for International Settlements 09 2022.

BlackRock, "BlackRock Adaptive U.S. Equity Index Series Methodology," March 2023. Methodology document.

Bloomberg Index Services Limited, "Bloomberg Volatility Target Indices: Multi-Asset Specifications," August 2025. Index specifications.

BNP Paribas Indices, "BNP Paribas Multi-Asset Trend Index: Brochure," 2025. Brochure.

Borio, Claudio and Mathias Drehmann, "Towards an operational framework for financial stability:" fuzzy" measurement and its consequences," *Documentos de Trabajo (Banco Central de Chile)*, 2009, (544), 1.

Boyd, Stephen and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

Brownlees, Christian and Robert F. Engle, "SRISK: A Conditional Capital Shortfall Measure of Systemic Risk," *The Review of Financial Studies*, 2017, *30* (1), 48–79.

Brunnermeier, Markus K. and Lasse Heje Pedersen, "Predatory Trading," *The Journal of Finance*, 2005, *60* (4), 1825–1863.

_ and _ , "Market Liquidity and Funding Liquidity," *The Review of Financial Studies*, 2009, *22* (6), 2201–2238.

Carlin, Bruce Ian, Miguel S. Lobo, and S. Viswanathan, "Episodic Liquidity Crises: Cooperative and Predatory Trading," *The Journal of Finance*, 2007, *62* (5), 2235–2274.

Chan, K., "A Further Analysis of the Lead-Lag Relationship Between the Cash Market and Stock Index Futures Market," *Review of Financial Studies*, 1992.

Dutt, Hans R. and Lawrence E. Harris, "Position limits for cash-settled derivative contracts," *Journal of Futures Markets*, 2005, *25* (10), 945–965.

ESRB, "Mitigating the Procyclicality of Margins and Haircuts in Derivatives Markets and Securities Financing Transactions," Technical Report, European Systemic Risk Board 01 2020.

European Central Bank, "Volatility-targeting strategies and the market sell-off," Financial Stability Review (Box/Focus), European Central Bank 05 2020.

Fardeau, Vincent, "Strategic Trading around Anticipated Supply/Demand Shocks," 2021. Preprint.

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz, "Trading Costs," 2018. Working paper.

Gatheral, Jim, "No-dynamic-arbitrage and market impact," *Quantitative Finance*, 2010, *10* (7), 749–759.

Geanakoplos, John, "The Leverage Cycle," in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual 2009, Volume 24*, University of Chicago Press, 2010.

Glasserman, Paul and Qi Wu, "Persistence and Procyclicality in Margin Requirements," *Management Science*, 2018.

Goldman Sachs, "Autocallable Goldman Sachs Momentum Builder Focus ER Index-Linked Notes due 2032 (Form 424B2)," SEC EDGAR filing December 2025.

Greenwood, Robin, Augustin Landier, and David Thesmar, "Vulnerable Banks," *Journal of Financial Economics*, 2015, *115* (3), 471–485.

Hillion, Pierre and Matti Suominen, "The manipulation of closing prices," *Journal of Financial Markets*, 2004, *7* (4), 351–375.

Huberman, Gur and Werner Stanzl, "Price Manipulation and Quasi-Arbitrage," *Econometrica*, 2004, *72* (4), 1247–1275.

Investment Company Institute, "Active and Index Combined Long-Term Mutual Funds and Exchange-Traded Funds (ETFs): December 2025," ICI Research and Statistics 2025. Table of total net assets by active versus index; accessed 2026-02-15.

IOSCO, "BCBS-CPMI-IOSCO finalise analysis of margining practices during the March 2020 market turmoil," Press release (IOSCO/MR/23/2022) September 2022.

ISDA, "ISDA Margin Survey: Year-end 2024," May 2025.

J.P. Morgan, "J.P. Morgan Mozaic Fixed Income Index (USD): Terms and Methodology Summary," 2015. SEC filing (free writing prospectus).

J.P. Morgan Chase, "Pricing Supplement: Notes Linked to the S&P 500 Daily Risk Control 5% Excess Return Index due April 3, 2025," Preliminary pricing supplement / term sheet March 2023.

Kawaller, Ira G., Paul D. Koch, and Timothy W. Koch, "The Temporal Price Relationship between S&P 500 Futures and the S&P 500 Index," *The Journal of Finance*, 12 1987, *42* (5), 1309–1329.

Khomyn, Marta, Putnins, and Marius Zoican, "The value of ETF liquidity," *The Review of Financial Studies*, 2024, *37* (10), 3092–3148.

Krein, David and Jeffrey Fernandez, "Indexing Volatility Risk Control," *The Journal of Index Investing*, 2012, *3* (2), 62–75.

Kumar, Praveen and Duane J. Seppi, "Futures Manipulation with "Cash Settlement"," *Journal of Finance*, 1992, *47* (4), 1485–1502.

Kyle, Albert S. and Anna A. Obizhaeva, "Market Microstructure Invariance: Empirical Hypotheses," *Econometrica*, 2016, *84* (4), 1345–1404.

LIMRA, "2024 Retail Annuity Sales Power to a Record $432.4 Billion," January 2025. Press release.

Lou, Dong, Hongjun Yan, and Jinfan Zhang, "Anticipated and Repeated Shocks in Liquid Markets," *Review of Financial Studies*, 2013, *26* (8), 1891–1912.

Meisenzahl, Ralf R., Jackson Overpeck, and Andy Polacek, "Life Insurers' Private Credit Investments and Annuity Market Share Capture," Working Paper WP 2025-09, Federal Reserve Bank of Chicago June 2025.

Moreira, Alan and Tyler Muir, "Volatility-Managed Portfolios," *The Journal of Finance*, 2017, *72* (4), 1611–1644.

Morgan Stanley, "Pricing Supplement (Form 424B2): S&P 500 Daily Risk Control 5% USD Excess Return Index Risk Disclosure," August 2025. SEC EDGAR filing.

MSCI, "MSCI Risk Control Indexes," November 2020. Factsheet (CFS1120).

_ , "MSCI Risk Control Indexes Methodology," August 2021. Methodology document.

_ , "MSCI Risk Control Indexes Methodology," August 2024. Methodology document.

_ , "MSCI Indexes Underlying U.S. Indexed Annuities," June 2025.

Nicolai, Francesco, "An Impossibility Theorem for Price-Based Risk Constraints," 2026. Working Paper.

Obizhaeva, Anna A. and Jiang Wang, "Optimal trading strategy and supply/demand dynamics," *Journal of Financial Markets*, 2013, *16* (1), 1–32.

Onur, Esen and David Reiffen, "The effect of settlement rules on the incentive to Bang the Close," *Journal of Futures Markets*, 2018, *38* (8), 841–864.

Rostek, Marzena and Marek Weretka, "Dynamic Thin Markets," *Review of Financial Studies*, 2015, *28* (10), 2946–2992.

Sammon, Marco and John J Shim, "Index rebalancing and stock market composition: Do indexes time the market?," *Journal of Financial Economics*, 2026, *177*, 104229.

Sannikov, Yuliy and Andrzej Skrzypacz, "Dynamic Trading: Price Inertia and Front-Running," 2016. Unpublished.

Schneider, Michael and Fabrizio Lillo, "Cross-impact and no-dynamic-arbitrage," *Quantitative Finance*, 2019, *19* (1), 137–154.

Securities and Exchange Commission, "Registration for Index-Linked Annuities and Registered Market Value Adjustment Annuities; Amendments to Form N-4 for Index-Linked Annuities, Registered Market Value Adjustment Annuities, and Variable Annuities; Other Technical Amendments," July 2024. Final rule release No. 33-11294.

S&P Dow Jones Indices, "Limiting Risk Exposure with S&P Risk Control Indices," 2021. Research note.

_ , "Limiting Risk Exposure with S&P Risk Control Indices," August 2021. Research note.

_ , "S&P Dow Jones Risk Control Indices Parameters," August 2025. Parameter document.

Stachurski, John and Thomas J. Sargent, "Economic Networks: Theory and Computation," 2022. QuantEcon Book I, Theorem 6.1.14.

Sushko, Vladyslav and Grant Turner, "The implications of passive investing for securities markets," BIS Quarterly Review, Bank for International Settlements 03 2018.

UBS Investment Bank, "UBS Multi Asset Strategy Tactical Rotation (UBS-MASTR) Index: Brochure," August 2025. Brochure.

Zhang, Anthony Lee, "Competition and Manipulation in Derivative Contract Markets," *Journal of Financial Economics*, 2022, *144* (2), 396–413.

Internet Appendix

for

Dynamic Arbitrage from Price-Based Risk Constraints

A Constrained sector demand

This appendix formalizes how a posted requirement translates into forced order flow by a constrained sector. The main results use only the fact that in a binding long regime, a higher requirement reduces the sector's feasible position, with a finite marginal sensitivity. Allowing constrained investors to anticipate the update and choose targets strategically changes that sensitivity but does not overturn the direction of the effect in binding states as shown in Nicolai (2026).

Assumption 7 (Constrained sector aggregate demand)**.** *A continuum of competitive traders $i \in [0,1]$ has equity $E(i) > 0$ and faces the feasibility constraint*

$$|x_t(i)|\, M_t \le E(i). \tag{77}$$

At date t, the sector specifies a public target exposure $\bar{x}_t \in \mathbb{R}$ before the date-t requirement M_t is implemented. Each trader implements the target up to the feasibility cap:

$$x_t(i) = \operatorname{sgn}(\bar{x}_t) \min\left\{ |\bar{x}_t|, \frac{E(i)}{M_t} \right\}. \tag{78}$$

Aggregate constrained demand is $X_t = \int_0^1 x_t(i)\, di$.

The timing captures slow-moving exposure targets (index replication, mandate weights, hedges) together with mechanical compliance with a posted requirement. The truncation form is a convenient reduced form; the analysis only needs that X_t is decreasing in M_t in binding states, with a finite local sensitivity

$$B = -\frac{dX}{dM} > 0.$$

Lemma 7 (Fully constrained benchmark)**.** *Suppose $\bar{x}_t \ge 0$ and $\bar{x}_t \ge \max_{i\in[0,1]} E(i)/M_t$, so that the feasibility constraint binds for all constrained traders. Let*

$$W = \int_0^1 E(i)\, di$$

denote total equity in the constrained sector. Then

$$X_t = \frac{W}{M_t}, \qquad \frac{dX_t}{dM_t} = -\frac{W}{M_t^2}. \tag{79}$$

The constrained sector rebalances each period to its feasible aggregate position X_t. Its aggregate net trade is

$$v_t = X_t - X_{t-1}. \tag{80}$$

When the sector is long, an increase in M_t reduces feasible demand X_t, so $v_t < 0$ corresponds to forced selling.

A.1 Partial discretion in the constrained sector

The closed-form benchmark above is useful for transparency, but the mechanism does not rely on it. What matters is monotonicity of feasible demand in the requirement, and a local linear approximation of how requirement changes translate into forced flow.

Lemma 8. *Fix equity $E > 0$ and let $V : \mathbb{R} \to \mathbb{R}$ be concave. Consider the investor problem*

$$x(M) \in \arg\max_{x \in \mathbb{R}} V(x) \quad \text{s.t.} \quad |x| \le \frac{E}{M}, \qquad M > 0.$$

Let $x^u \in \arg\max_{x \in \mathbb{R}} V(x)$ be an unconstrained optimum. Then

$$x(M) = \operatorname{sgn}(x^u) \min\{|x^u|, E/M\}.$$

In particular, if $x^u \ge 0$, then $x(M) = \min\{x^u, E/M\}$ is weakly decreasing in M, and if the constraint binds ($x(M) = E/M$) then $dx/dM = -E/M^2 < 0$.

Proof. The feasible set is the interval $[-E/M, E/M]$. If x^u lies inside this interval it remains optimal. If $x^u > E/M$ (resp. $x^u < -E/M$), concavity implies the best feasible choice is the right (resp. left) endpoint. This yields the truncation formula. When $x^u \ge 0$, the cap E/M falls in M, so $x(M)$ is weakly decreasing, with derivative $-E/M^2$ wherever binding. □

Consider a continuum of long constrained investors $i \in [0, 1]$ with equity $E(i) > 0$ and unconstrained targets $x^u(i) \ge 0$. Lemma 8 implies

$$x(i, M) = \min\left\{x^u(i), \frac{E(i)}{M}\right\}, \qquad X(M) = \int_0^1 x(i, M)\, di.$$

For each i, $x(i, M)$ is weakly decreasing in M, hence $X(M)$ is weakly decreasing in M. At a reference point M_0, define the set of binding investors

$$\mathcal{B}_0 = \left\{i \in [0, 1] : x^u(i) > \frac{E(i)}{M_0}\right\}.$$

Differentiating at M_0 gives

$$X'(M_0) = -\int_{\mathcal{B}_0} \frac{E(i)}{M_0^2}\, di < 0.$$

Lemma 9. *Fix a reference requirement $M_0 > 0$. Suppose $X : (0, \infty) \to \mathbb{R}$ is decreasing and differentiable in a neighborhood of M_0, with $X'(M_0) < 0$ and X' continuous at M_0. Define forced flow by*

$$v_{t+1} = X(M_{t+1}) - X(M_t),$$

and define the local sensitivity $B = -X'(M_0) > 0$. Then, as $(M_t, M_{t+1}) \to (M_0, M_0)$,

$$v_{t+1} = -B\,(M_{t+1} - M_t) + o(|M_{t+1} - M_t|)\,.$$

In particular, local forced flow depends on $X(\cdot)$ only through B. In the benchmark $X(M) = W/M$ of Lemma 7,

one has $B = W/M_0^2$.

Proof. Write $\Delta M_{t+1} = M_{t+1} - M_t$. Differentiability at M_0 implies

$$X(M_{t+1}) - X(M_t) = X'(M_0)\Delta M_{t+1} + o(|\Delta M_{t+1}|) \qquad \text{as } (M_t, M_{t+1}) \to (M_0, M_0).$$

Since $v_{t+1} = X(M_{t+1}) - X(M_t)$ and $B = -X'(M_0)$, the expansion follows. For $X(M) = W/M$, $X'(M) = -W/M^2$, hence $B = W/M_0^2$. □

B Augmented no-arbitrage condition

In this appendix we provide two complements to Section 4. First, we characterize when the linear system $(I + \mathcal{K})q = u$ is well posed. Second, we derive the eigenvalue-based simplification used in Section B.1. For a given attack path u, invertibility of $I + \mathcal{K}$ pins down a unique induced total-flow path q and hence a unique $P - S$. If $I + \mathcal{K}$ is singular, the same u can generate multiple q and thus multiple execution-price paths; in that case the linearized profit functional is not well defined.

It turns out that timing matters. Under the baseline convention $M_{t+1} = g(\Gamma_t)$, a trade at date t affects feasibility only from $t + 1$ onward. Over a finite horizon, the resulting feedback operator $\mathcal{K}$ is strictly lower triangular, so $I+\mathcal{K}$ is automatically invertible (Lemma 11). Non-invertibility arises only with within-period updating, for example $M_t = g(\Gamma_t)$, where same-date trades can move sampled prices and relax or tighten feasibility immediately.

Definition 4 (Margin-feedback manipulation and quasi-arbitrage)**.** *Under the usual assumptions*

(i) A margin-feedback manipulation is an admissible round trip $u \in \mathcal{U}_T \cap \mathcal{R}_T$ with $\mathbb{E}[\Pi(u)] > 0$.

(ii) A margin-feedback quasi-arbitrage is a sequence $\{u^n\}_{n\geq 1} \subset \mathcal{U}_T \cap \mathcal{R}_T$ with $\mathbb{E}[\Pi(u^n)] \to \infty$ and

$$\frac{\mathbb{E}[\Pi(u^n)]}{\sqrt{\mathbb{V}\mathrm{ar}(\Pi(u^n))}} \to \infty.$$

This quasi-arbitrage notion is the one in Huberman and Stanzl (2004): if a positive-mean round trip exists, scaling position size generates a sequence with diverging expected profits and Sharpe ratios.

Lemma 10. *Consider the linearized feedback system (22)–(28). For a proposed trade vector $u \in \mathbb{R}^T$, the induced total flow $q \in \mathbb{R}^T$ and wedge-price path $P - S$ satisfy*

$$(I + \mathcal{K})q = u, \qquad P - S = \mathcal{I}q,$$

where I is the $T \times T$ identity matrix, $\mathcal{K}$ is the loop operator in (27), and $\mathcal{I}$ is the impact operator. The following statements are equivalent:

(i) $I + \mathcal{K}$ is invertible on $\mathbb{R}^T$.

(ii) For every $u \in \mathbb{R}^T$ there exists a unique $q \in \mathbb{R}^T$ solving $(I+\mathcal{K})q = u$; equivalently $P-S = \mathcal{I}(I+\mathcal{K})^{-1}u$.

(iii) For every $u \in \mathcal{R}_T$ there exists a unique $q \in \mathbb{R}^T$ solving $(I + \mathcal{K})q = u$.

Proof. See (Stachurski and Sargent, 2022, Theorem 6.1.14). □

Lemma 11 (Invertibility of $I + \mathcal{K}$ under lagged updating)**.** *Fix a horizon T and stack time-indexed paths in chronological order as column vectors $x = (x_0, \ldots, x_{T-1})^\top$. Suppose $M_{t+1} = g(\Gamma_t)$, and forced flow is the constrained sector's net trade $v_t = X_t - X_{t-1}$ with $X_t = X(M_t)$. Assume additionally that (i) the linearized impact operator $\mathcal{I}$ is such that its $T \times T$ matrix representation satisfies $\mathcal{I}_{t,s} = 0$ for all $s > t$, and (ii) Γ_t depends only on prices up to date t, so its Jacobian $J_{t,s} = 0$ for all $s > t$. Then $\mathcal{K}q$ is such that its elements*

$\mathcal{K}_{t,s} = 0$ for all $s \geq t$, i.e., it is strictly lower triangular. Then $I + \mathcal{K}$ is invertible on $\mathbb{R}^T$, and

$$(I + \mathcal{K})^{-1} = \sum_{n=0}^{T-1} (-\mathcal{K})^n.$$

Proof. The one-step lag operator L satisfies $(Lx)_0 = 0$ and $(Lx)_t = x_{t-1}$ for $t = 1, \ldots, T-1$, so L is strictly lower triangular. The first-difference operator D satisfies $(Dx)_0 = x_0$ and $(Dx)_t = x_t - x_{t-1}$ for $t = 1, \ldots, T-1$, so D is lower triangular. Hence DL is strictly lower triangular. J and $\mathcal{I}$ are lower triangular, so their product $J\mathcal{I}$ is lower triangular. Since $\mathcal{K} = BsDLJ\mathcal{I}$ and scalar factors do not affect triangularity, it follows that $\mathcal{K} = (DL)(J\mathcal{I})$ is strictly lower triangular, equivalently $\mathcal{K}_{t,s} = 0$ for all $s \geq t$. Strict lower-triangularity implies that the T-th power vanishes: $\mathcal{K}^T = 0$. Hence

$$(I + \mathcal{K}) \sum_{n=0}^{T-1} (-\mathcal{K})^n = I - (-\mathcal{K})^T = I,$$

so $I + \mathcal{K}$ is invertible with inverse given by the displayed finite sum. This is the finite-horizon case of the Neumann-series identity; see (Stachurski and Sargent, 2022, Theorem 1.2.5). □

This proves lagged margin adjustment automatically guarantees invertibility. We then proceed to show that contemporaneous margin adjustment can destroy well-posedness.

Theorem 3 (Singular feedback with instantaneous updating)**.** *Consider the linearized system (22)–(28).*

$$(I + \mathcal{K})q = u, \qquad P - S = \mathcal{I}q.$$

(i) ***Lagged updating implies automatic well-posedness.*** *Under the timing $M_{t+1} = g(\Gamma_t)$, $I + \mathcal{K}$ is invertible on $\mathbb{R}^T$ (Lemma 11).*

(ii) ***Instantaneous updating can generate singularity and non-uniqueness.*** *If within-period dependence $M_t = g(\Gamma_t)$ is allowed, let $\mathcal{K}^0$ denote the corresponding feedback matrix, so the fixed point becomes*

$$(I + \mathcal{K}^0)q = u, \qquad P - S = \mathcal{I}q.$$

Invertibility of $I + \mathcal{K}^0$ is no longer automatic. If $I + \mathcal{K}^0$ is singular, then even at $u = 0$ the induced flow path is not uniquely determined, so the wedge path $P - S$ and the induced execution-price path are not single-valued functions of u. Realized profit is undefined.

(iii) ***Near singularity implies large amplification; failure of the sign test yields quasi-arbitrage.*** *If $I + \mathcal{K}$ is invertible, then*

$$\sup_{\|u\|_2 = 1} \|(I + \mathcal{K})^{-1} u\|_2 = \|(I + \mathcal{K})^{-1}\|_2 = \frac{1}{\sigma_{\min}(I + \mathcal{K})},$$

and the same identity holds with $\mathcal{K}$ replaced by $\mathcal{K}^0$ whenever $I + \mathcal{K}^0$ is invertible. Hence if $\sigma_{\min}(I + \mathcal{K})$ converges to zero while invertibility is maintained, arbitrarily small inputs can generate arbitrarily large induced flows, and therefore arbitrarily large price responses.

Moreover when $\widehat{H}$ is not semipositive definite, there exists a margin-feedback quasi-arbitrage in the sense of Definition 4(ii).

Proof. Part (i) is Lemma 11.
For part (ii), if $I + \mathcal{K}^0$ is singular, choose $q_0 \neq 0$ such that

$$(I + \mathcal{K}^0)q_0 = 0.$$

Then at $u = 0$, both $q = 0$ and $q = q_0$ solve

$$(I + \mathcal{K}^0)q = u,$$

so the induced flow path is not unique. It remains to show that the price path is not unique either. $\mathcal{I}q_0$ cannot be zero, since if $\mathcal{I}q_0 = 0$, then we also have $\mathcal{K}^0 q_0 = 0$, hence

$$(I + \mathcal{K}^0)q_0 = q_0,$$

which contradicts $q_0 \neq 0$. Therefore $\mathcal{I}q_0 \neq 0$, so

$$P - S = \mathcal{I}q$$

is not uniquely determined even at $u = 0$. Since realized profit is computed from the induced execution-price path, it is not well defined.
For part (iii), for any invertible matrix A,

$$\sup_{\|u\|_2=1} \|A^{-1}u\|_2 = \|A^{-1}\|_2 = \frac{1}{\sigma_{\min}(A)}.$$

Applying this with $A = I + \mathcal{K}$ gives

$$\sup_{\|u\|_2=1} \|(I + \mathcal{K})^{-1}u\|_2 = \|(I + \mathcal{K})^{-1}\|_2 = \frac{1}{\sigma_{\min}(I + \mathcal{K})},$$

and the same argument applies to $A = I + \mathcal{K}^0$ whenever $I + \mathcal{K}^0$ is invertible. Hence, if $\sigma_{\min}(I + \mathcal{K}) \downarrow 0$ while invertibility is preserved, the worst-case gain of the feedback map explodes: arbitrarily small inputs can generate arbitrarily large induced flows, and therefore arbitrarily large price responses.
Finally, if $\widehat{H}$ is not positive semidefinite on $\mathcal{R}_T$, there exists $u_0 \in \mathcal{R}_T$ such that

$$u_0^\top \widehat{H} u_0 < 0.$$

Choosing the admissibility bound large enough so that $u_0 \in \mathcal{U}_T \cap \mathcal{R}_T$, the quadratic profit formula (31) yields

$$\mathbb{E}[\Pi(u_0)] = -\frac{1}{2} u_0^\top \widehat{H} u_0 > 0.$$

For $n \geq 1$, let $u^n = nu_0$. Since the induced wedge path is linear in u, expected profit scales quadrati-

cally:

$$\mathbb{E}[\Pi(u^n)] = -\frac{1}{2}\, u^{n\top} \widehat{H} u^n = n^2 \mathbb{E}[\Pi(u_0)].$$

By Lemma 1, the only random part is the unaffected-price term, so

$$\mathbb{V}\mathrm{ar}(\Pi(u^n)) = n^2 \mathbb{V}\mathrm{ar}(u_0^\top S).$$

Therefore

$$\frac{\mathbb{E}[\Pi(u^n)]}{\sqrt{\mathbb{V}\mathrm{ar}(\Pi(u^n))}} = n \frac{\mathbb{E}[\Pi(u_0)]}{\sqrt{\mathbb{V}\mathrm{ar}(u_0^\top S)}} \to \infty.$$

This is exactly the quasi-arbitrage conclusion in Definition 4(ii). □

Part (iii) uses the standard quasi-arbitrage notion of Huberman and Stanzl (2004). The result implies that once the sign condition fails, there exists a round-trip u_0 whose expected profit is positive, and scaling that same trading pattern up produces a sequence of admissible round trips $u^n = nu_0$ whose mean profit grows like n^2 while the only random component, the unaffected-price term $u^{n\top} S$, has standard deviation that grows only like n. The Sharpe ratio of the scaled round trip can be made arbitrarily large, which is exactly the sense in which Definition 4 calls the configuration vulnerable to margin-feedback quasi-arbitrage.

B.1 Eigenvalue tests

Theorem 1 says that manipulation is ruled out if the quadratic form $u^\top \widehat{H} u$ is nonnegative for every round trip $u \in \mathcal{R}_T$. In practice it is useful to compress this into a single number: the worst-case value of $u^\top \widehat{H} u$ over unit-norm round trips (Boyd and Vandenberghe, 2004, Section A.5.4).

Theorem 4. *Under the assumptions of Theorem 1 and with $\widehat{H}$ as in (30), define*

$$\lambda_{\min}^{\mathcal{R}_T}(\widehat{H}) = \inf \left\{ u^\top \widehat{H} u : \; u \in \mathcal{R}_T, \; \|u\|_2 = 1 \right\}.$$

Then $\widehat{H}$ is positive semidefinite on $\mathcal{R}_T$ if and only if $\lambda_{\min}^{\mathcal{R}_T}(\widehat{H}) \geq 0$. Equivalently, with $Q = -\frac{1}{2}\widehat{H}$,

$$\lambda_{\max}^{\mathcal{R}_T}(Q) = \sup \left\{ u^\top Q u : \; u \in \mathcal{R}_T, \; \|u\|_2 = 1 \right\} \leq 0.$$

The same statements hold in the multi-asset model after replacing $\mathcal{R}_T$ by $\mathbf{R}_T$ and $\widehat{H}$ by $\widehat{\mathbf{H}}$.

Proof. If $\widehat{H}$ is positive semidefinite on $\mathcal{R}_T$, then $u^\top \widehat{H} u \geq 0$ for every unit-norm $u \in \mathcal{R}_T$, so the infimum over that set is nonnegative, that is, $\lambda_{\min}^{\mathcal{R}_T}(\widehat{H}) \geq 0$. Conversely, if $\lambda_{\min}^{\mathcal{R}_T}(\widehat{H}) \geq 0$, then $u^\top \widehat{H} u \geq 0$ for every unit-norm $u \in \mathcal{R}_T$, hence $\widehat{H}$ is positive semidefinite on $\mathcal{R}_T$. The reformulation in terms of $Q = -\frac{1}{2}\widehat{H}$ follows directly from $u^\top Q u = -\frac{1}{2}\, u^\top \widehat{H} u$. □

C Robustness to nonlinearity

Theorem 1 is obtained by linearizing (Γ, g, X) around a binding configuration. This appendix shows that, under local smoothness and away from kinks, if the test fails then the exact nonlinear feedback system admits a genuinely profitable sufficiently small round trip.

Lemma 12. *Fix a binding reference configuration (z, W). Assume:*

(i) The map from transaction prices to the statistic Γ, and the update maps g and $X(\cdot)$, are twice continuously differentiable in a neighborhood of the reference point.

(ii) Locally, no piecewise component switches regime (caps, floors, kinks), as in Assumption 6(ii).

(iii) $I + \mathcal{K}(z; W)$ is invertible.

Then there exists $\rho > 0$ such that for any deterministic strategic order vector U with $\|U\| < \rho$, the exact nonlinear feedback system has a unique induced total-flow path $q(U)$ and hence a unique wedge-price path $\Delta P(U) = \mathcal{I}q(U)$. Moreover, $q(\cdot)$ is differentiable at $U = 0$ and

$$Dq(0) = (I + \mathcal{K}(z; W))^{-1}.$$

Proof. Given a candidate total-flow path q, wedge prices are $\Delta P = \mathcal{I}q$. Feeding the implied transaction-price path through the exact rule (16)–(17) and the exact demand map $X(\cdot)$ delivers the constrained-sector position path, and first differences deliver forced flow. Denote this exact forced-flow mapping by $v = \Phi(q)$. Then

$$q = U + \Phi(q).$$

Under (i) and (ii), Φ is continuously differentiable near the reference point and satisfies $\Phi(0) = 0$. Its derivative at the reference point is $D\Phi(0) = -\mathcal{K}(z; W)$. Define

$$F(q, U) = q - U - \Phi(q).$$

Then $\partial_q F(0, 0) = I + \mathcal{K}(z; W)$, which is invertible by (iii). The implicit function theorem gives local existence and uniqueness of $q(U)$ for $\|U\|$ small. Differentiating $F(q(U), U) = 0$ at $U = 0$ yields $Dq(0) = (I + \mathcal{K}(z; W))^{-1}$. □

Proposition 9 (A violated test implies a genuinely profitable small round trip)**.** *Fix (z, W) and impose the assumptions of Lemma 12. Suppose the linearized test fails: there exists a deterministic round trip direction $\bar{U}$ in the relevant round-trip subspace (that is, $\bar{U} \in \mathcal{R}_T$ in the single-asset case, and $\bar{U} \in \mathbf{R}_T$ in the multi-asset case) such that*

$$\bar{U}^\top \widehat{H}(z; W)\, \bar{U} < 0.$$

Then there exists $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon})$ the exact nonlinear model admits the scaled round trip $U^\varepsilon = \varepsilon \bar{U}$ (and it is admissible for ε small), and this trade has strictly positive expected profit:

$$\mathbb{E}[\Pi(U^\varepsilon)] > 0.$$

Proof. Fix $\bar{U}$ with $\bar{U}^\top \widehat{H}(z;W)\bar{U} < 0$ and consider the scaled round trip $U^\varepsilon = \varepsilon \bar{U}$. On any round trip, the unaffected-price component contributes zero in expectation (Lemma 1 in the single-asset case and Lemma 2 in the multi-asset case). Expected profit is therefore pinned down by the execution wedge:

$$\mathbb{E}[\Pi(U^\varepsilon)] = -\,U^{\varepsilon\top}\Delta P(U^\varepsilon) = -\,U^{\varepsilon\top}\mathcal{I}q(U^\varepsilon).$$

Lemma 12 implies that the induced total-flow map $q(\cdot)$ is differentiable at the reference point, with derivative $Dq(0) = (I + \mathcal{K}(z;W))^{-1}$. Hence, for ε small,

$$q(U^\varepsilon) = (I + \mathcal{K}(z;W))^{-1}U^\varepsilon + o(\|U^\varepsilon\|).$$

Substituting into the profit expression gives a second-order expansion:

$$\mathbb{E}[\Pi(U^\varepsilon)] = -\,U^{\varepsilon\top}\mathcal{I}(I + \mathcal{K}(z;W))^{-1}U^\varepsilon + o(\|U^\varepsilon\|^2) = -\frac{1}{2}\,U^{\varepsilon\top}\widehat{H}(z;W)U^\varepsilon + o(\|U^\varepsilon\|^2),$$

where the last equality uses (30). The leading term equals $-\frac{1}{2}\,\varepsilon^2\,\bar{U}^\top \widehat{H}(z;W)\bar{U} > 0$, so it is strictly positive and of order ε^2. The remainder is negligible relative to ε^2 for sufficiently small ε, hence there exists $\bar{\varepsilon} > 0$ such that $\mathbb{E}[\Pi(U^\varepsilon)] > 0$ for all $\varepsilon \in (0, \bar{\varepsilon})$. Admissibility holds for ε small because U^ε is deterministic (hence predictable) and its per-date magnitudes scale linearly in ε (Definition 3 and its multi-asset analogue). □

D Proofs and derivations for Section 5

D.1 Stochastic dynamic program (state augmentation and conditional expectations)

Sections 5.1–6.1 work conditional on a deterministic unaffected-price path so that the Bellman recursion is deterministic on the state vectors used there. This subsection records the corresponding formulation when the unaffected price is stochastic.

Let $(\mathcal{F}_t)_{t=-1}^{T}$ be the filtration from Section 2 and maintain the discrete-time predictability convention: u_t is $\mathcal{F}_{t-1}$-measurable for $t = 0, \dots, T-1$. Suppose the unaffected price is driven by a Markov state Z_t that is $\mathcal{F}_{t-1}$-measurable (for example, $Z_t = S_{t-1}$ when (S_t) is Markov). Augment the main-text state by this driver:

$$\tilde{s}_t = (s_t, Z_t).$$

Given $\tilde{s}_t$ and an action $u_t \in \mathcal{U}_t(x_t)$, the next state $\tilde{s}_{t+1}$ is determined by the same accounting identities as in (38) (single asset) or the display defining $\mathcal{T}$ in Section 6.1 (multi asset), together with the Markov update $Z_{t+1} = \Phi(Z_t, \varepsilon_{t+1})$ for i.i.d. innovations (ε_{t+1}).

The stochastic value function is

$$\widetilde{V}_t(\tilde{s}) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \mathbb{E}\left[\sum_{\tau=t}^{T-1} \pi_\tau(\tilde{s}_\tau, u_\tau) \,\middle|\, \tilde{s}_t = \tilde{s}\right], \tag{81}$$

and the Bellman recursion is

$$\widetilde{V}_t(\tilde{s}) = \max_{u \in \mathcal{U}_t(x)} \left\{\pi_t(\tilde{s}, u) + \mathbb{E}\left[\widetilde{V}_{t+1}(\tilde{s}_{t+1}) \,\middle|\, \tilde{s}_t = \tilde{s},\ u_t = u\right]\right\}, \tag{82}$$

with the terminal condition $\widetilde{V}_T(\tilde{s}) = 0$ if $x_T = 0$ and $-\infty$ otherwise.

At an interior optimum, and under regularity conditions justifying differentiation under the conditional expectation, the within-period first-order conditions take the same form as in Propositions 3 and 7, except that the shadow values are conditional expectations of next-period marginal values. For the single-asset problem, define

$$\mu_{t+1} = \mathbb{E}\left[\partial_x \widetilde{V}_{t+1}(\tilde{s}_{t+1}) \,\middle|\, \tilde{s}_t,\ u_t^\star\right], \qquad \nu_{t+1} = \mathbb{E}\left[\frac{d}{dP_t}\widetilde{V}_{t+1}(\tilde{s}_{t+1}) \,\middle|\, \tilde{s}_t,\ u_t^\star\right],$$

then the Euler equation is $\eta(2u_t^\star + v_t) = \mu_{t+1} + \eta\nu_{t+1}$. For the multi-asset problem, with

$$\mu_{t+1} = \mathbb{E}\left[\nabla_x \widetilde{V}_{t+1}(\tilde{s}_{t+1}) \,\middle|\, \tilde{s}_t,\ u_t^\star\right], \qquad \nu_{t+1} = \mathbb{E}\left[\nabla_{P_t} \widetilde{V}_{t+1}(\tilde{s}_{t+1}) \,\middle|\, \tilde{s}_t,\ u_t^\star\right],$$

the Euler equation is $A(2u_t^\star + v_t) = \mu_{t+1} + A\nu_{t+1}$.

D.2 Proof of Proposition 3

Proof. Fix s_t and write $s_{t+1}(u) = \mathcal{T}(s_t, u)$. Let $\mu_{t+1}(u) = \partial_x V_{t+1}(s_{t+1}(u))$, $\nu_{t+1}(u) = \frac{d}{dP_t} V_{t+1}(s_{t+1}(u))$, where $\frac{d}{dP_t}$ denotes the total derivative through all components of s_{t+1} that depend on P_t under the transition map. Evaluate these at $u = u_t^\star$ and denote the resulting values by μ_{t+1} and ν_{t+1}. The

Bellman objective at date t is $\pi_t(s_t, u) + V_{t+1}(\mathcal{T}(s_t, u))$. By (37), $\partial_u \pi_t(s_t, u) = -\eta(2u + v_t)$. Under the stated differentiability conditions, u enters $\mathcal{T}(s_t, u)$ only through $x_{t+1} = x_t + u$ and $P_t = S_t + \eta(u + v_t)$; all other next-state components in (38) are functions of these and predetermined objects in s_t. Therefore, by the chain rule,

$$\frac{d}{du} V_{t+1}(\mathcal{T}(s_t, u)) = \partial_x V_{t+1}(s_{t+1}) \cdot \frac{dx_{t+1}}{du} + \frac{d}{dP_t} V_{t+1}(s_{t+1}) \cdot \frac{dP_t}{du} = \mu_{t+1} + \eta \nu_{t+1},$$

evaluated at $u = u_t^\star$. At an interior optimum,

$$0 = \frac{d}{du}\Big[\pi_t(s_t, u) + V_{t+1}(\mathcal{T}(s_t, u))\Big]_{u=u_t^\star} = -\eta(2u_t^\star + v_t) + \mu_{t+1} + \eta \nu_{t+1},$$

which gives (44) and hence (45). Finally, (46) follows from Lemma 3 and the identity

$$u_t^\star + \tfrac{v_t}{2} = \frac{1}{2\eta}\big(\mu_{t+1} + \eta \nu_{t+1}\big),$$

which is obtained by rearranging (44). □

D.3 Proof of Corollary 1

Proof. Since $0 \in \mathcal{U}_t(x_t)$ by assumption, $u = 0$ is an admissible deviation in (41). If $v_t = 0$, then Lemma 3 gives $\pi_t(s_t, u) \le 0 = \pi_t(s_t, 0)$ for all u. Optimality of $u_t^\star$ in (41) implies

$$\pi_t(s_t, u_t^\star) + V_{t+1}(\mathcal{T}(s_t, u_t^\star)) \ge \pi_t(s_t, 0) + V_{t+1}(\mathcal{T}(s_t, 0)) = V_{t+1}(\mathcal{T}(s_t, 0)),$$

so if $u_t^\star \neq 0$ then the continuation term must strictly outweigh the contemporaneous loss $\pi_t(s_t, u_t^\star) \le 0$. Under the conditions of Proposition 3, setting $v_t = 0$ in (45) yields the stated reduction. Finally, if $\mu_{t+1} = \nu_{t+1} = 0$ at an interior optimum, then (45) gives $u_t^\star = -v_t/2$. □

D.4 Proof of Corollary 2

Proof. From (45), $u_t^\star = -v_t/2 + \chi_{t+1}/(2\eta)$. If $v_t \neq 0$, then $u_t^\star v_t > 0$ is equivalent to

$$\left(-\frac{v_t}{2} + \frac{\chi_{t+1}}{2\eta}\right) v_t > 0 \quad \Longleftrightarrow \quad \chi_{t+1} v_t > \eta v_t^2,$$

which holds if and only if $\operatorname{sign}(\chi_{t+1}) = \operatorname{sign}(v_t)$ and $|\chi_{t+1}| > \eta |v_t|$. For (ii), the payoff identity (46) implies

$$\pi_t(s_t, u_t^\star) = \frac{\eta}{4} v_t^2 - \frac{1}{4\eta} \chi_{t+1}^2,$$

which is negative if and only if $|\chi_{t+1}| > \eta |v_t|$. □

D.5 Proof of Proposition 4

Proof. View V_{t+1} as a differentiable function of the next state

$$s_{t+1} = (x_{t+1}, y_{t+1}, P_t, M_{t+1}, X_{t+1}, X_t).$$

When differentiating with respect to P_t, the component $x_{t+1} = x_t + u_t$ is fixed. The components that depend on P_t are $y_{t+1} = F(y_t, P_t)$ and (M_{t+1}, X_{t+1}), where

$$M_{t+1} = g(\Gamma_t), \qquad \Gamma_t = \Gamma(y_t, P_{t-1}, P_t), \qquad X_{t+1} = \frac{W}{M_{t+1}}.$$

By the chain rule,

$$\begin{aligned}\nu_{t+1} &= \frac{d}{dP_t} V_{t+1}(s_{t+1}) \\ &= \partial_P V_{t+1}(s_{t+1}) + \big(\nabla_y V_{t+1}(s_{t+1})\big) \cdot \frac{\partial y_{t+1}}{\partial P_t} + \partial_M V_{t+1}(s_{t+1}) \cdot \frac{dM_{t+1}}{dP_t} + \partial_X V_{t+1}(s_{t+1}) \cdot \frac{dX_{t+1}}{dP_t}.\end{aligned}$$

Since $y_{t+1} = F(y_t, P_t)$, we have $\frac{\partial y_{t+1}}{\partial P_t} = \partial_P F(y_t, P_t)$. Moreover,

$$\frac{dM_{t+1}}{dP_t} = g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t).$$

Finally, $X_{t+1} = W/M_{t+1}$ implies

$$\frac{dX_{t+1}}{dP_t} = -\frac{W}{M_{t+1}^2} \frac{dM_{t+1}}{dP_t} = -B_{t+1} \frac{dM_{t+1}}{dP_t}.$$

Substituting the last two displays into the expansion for ν_{t+1} and grouping the (∂_M, ∂_X) terms yields

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + \big(\nabla_y V_{t+1}(s_{t+1})\big) \cdot \partial_P F(y_t, P_t) + \big(\partial_M V_{t+1}(s_{t+1}) - B_{t+1} \partial_X V_{t+1}(s_{t+1})\big) \frac{dM_{t+1}}{dP_t},$$

which is (48) after inserting the definition (47). □

D.6 Proof of Corollary 3

Proof. The decomposition (48) is a sum of three terms. The only term that depends on (g, Γ) is the last product $\Lambda_{t+1} g'(\Gamma_t) \partial_P \Gamma(y_t, P_{t-1}, P_t)$, so it vanishes when either factor $g'(\Gamma_t)$ or $\partial_P \Gamma(y_t, P_{t-1}, P_t)$ is zero. The sign claim follows directly from the same product representation. □

D.7 Proof of Lemma 4

Proof. Fix $t \geq T - 1 - \ell$. For any $\tau \in \{t+1, \ldots, T-1\}$ we have $M_\tau = g(\Gamma_{\tau-1-\ell})$ with $\tau - 1 - \ell \leq T - 2 - \ell \leq t - 1$, so M_τ (and $X_\tau = W/M_\tau$) depends only on prices up to date $t-1$ and is fixed at time t. Hence the within-horizon forced flows $(v_{t+1}, \ldots, v_{T-1})$ are predetermined at time t and cannot be influenced by actions $(u_t, \ldots, u_{T-1})$. Once this sequence is fixed, no constraint or payoff from date $t+1$ onward depends on P_t, so $V_{t+1}(s_{t+1})$ is locally constant in the P_t direction and $\nu_{t+1} = 0$

whenever the derivative exists. For the last-trigger-date claim, if $t \le T-2-\ell$ then $t+\ell+1 \le T-1$, so Γ_t enters $M_{t+\ell+1} = g(\Gamma_t)$ within the horizon and can affect forced flow. If $t \ge T-1-\ell$, the first part shows that no within-horizon requirement $(M_{t+1}, \dots, M_{T-1})$ depends on Γ_t. □

D.8 Proof of Proposition 5

Proof. When future forced flows are exogenous, the continuation objective from date t is

$$\max_{(u_\tau)_{\tau=t}^{T-1}} \; -\eta \sum_{\tau=t}^{T-1} \left(u_\tau^2 + u_\tau v_\tau\right) \qquad \text{s.t.} \qquad \sum_{\tau=t}^{T-1} u_\tau = -x_t,$$

where the constraint is equivalent to $x_T = 0$ given $x_{t+1} = x_t + u_t$. The objective is a strictly concave quadratic and the constraint set is affine, so the maximizer is unique. Introduce a Lagrange multiplier λ on the inventory constraint. The Lagrangian is

$$\mathcal{L} = -\eta \sum_{\tau=t}^{T-1} \left(u_\tau^2 + u_\tau v_\tau\right) + \lambda \left(\sum_{\tau=t}^{T-1} u_\tau + x_t \right).$$

The first-order condition for each τ is

$$-2\eta u_\tau - \eta v_\tau + \lambda = 0 \qquad \Longrightarrow \qquad u_\tau = -\frac{v_\tau}{2} + \frac{\lambda}{2\eta}.$$

Summing over $\tau = t, \dots, T-1$ and imposing $\sum_{\tau=t}^{T-1} u_\tau = -x_t$ yields

$$-\frac{1}{2} \sum_{\tau=t}^{T-1} v_\tau + \frac{n\lambda}{2\eta} = -x_t \qquad \Longrightarrow \qquad \frac{\lambda}{2\eta} = \frac{\bar{v}_t}{2} - \frac{x_t}{n},$$

which gives (49).

To obtain the value, use Lemma 3 date by date:

$$-\eta u_\tau (u_\tau + v_\tau) = -\eta \Big(u_\tau + \tfrac{v_\tau}{2}\Big)^2 + \frac{\eta}{4} v_\tau^2.$$

Under (49),

$$u_\tau^\star + \tfrac{v_\tau}{2} = \frac{\bar{v}_t}{2} - \frac{x_t}{n},$$

which is constant across τ. Therefore

$$V_t^{\text{harv}} = \sum_{\tau=t}^{T-1} \left[\frac{\eta}{4} v_\tau^2 - \eta \left(\frac{\bar{v}_t}{2} - \frac{x_t}{n} \right)^2 \right] = \frac{\eta}{4} \sum_{\tau=t}^{T-1} v_\tau^2 - \eta n \left(\frac{\bar{v}_t}{2} - \frac{x_t}{n} \right)^2.$$

Expanding the last square and using $\sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2 = \sum_{\tau=t}^{T-1} v_\tau^2 - n\bar{v}_t^2$ yields (50). □

D.9 Proof of Corollary 4

Proof. The first display is exactly (50). If $v_\tau = c$, then $\bar{v}_t = c$ and $\sum_{\tau=t}^{T-1}(v_\tau - \bar{v}_t)^2 = 0$, so (50) gives $V_t^{\mathrm{harv}} = \eta x_t c - \frac{\eta}{n} x_t^2$. Equation (49) yields $u_\tau^\star = -\frac{1}{2}(c - c) - \frac{x_t}{n} = -\frac{x_t}{n}$ for all τ. Finally, if $x_t = 0$ then $u_\tau^\star = 0$ and $V_t^{\mathrm{harv}} = 0$. □

D.10 Proof of Corollary 5

Proof. Under the stated forced-flow path, $\bar{v}_t = \frac{1}{n}\sum_{\tau=t}^{T-1} v_\tau = v/n$. Applying (49) with $x_t = 0$ gives, for $\tau = t, \dots, T-2$,

$$u_\tau^\star = -\frac{1}{2}\left(0 - \frac{v}{n}\right) = \frac{v}{2n},$$

and for $\tau = T-1$,

$$u_{T-1}^\star = -\frac{1}{2}\left(v - \frac{v}{n}\right) = -\frac{n-1}{2n} v.$$

Since $v < 0$, the terminal trade is positive. Finally, with $x_t = 0$, Corollary 4 gives $V_t^{\mathrm{harv}} = \frac{\eta}{4}\sum_{\tau=t}^{T-1}(v_\tau - \bar{v}_t)^2$. Here there are $n-1$ terms equal to $(0 - v/n)^2 = v^2/n^2$ and one term equal to $(v - v/n)^2 = v^2(n-1)^2/n^2$, so

$$\sum_{\tau=t}^{T-1}(v_\tau - \bar{v}_t)^2 = \frac{v^2}{n^2}\big((n-1) + (n-1)^2\big) = \frac{n-1}{n} v^2,$$

which yields the stated value. □

D.11 Proof of Proposition 6

Proof. The objective is a strictly concave quadratic in $(u_\tau)_{\tau=t}^{T-1}$ and the constraints define a nonempty compact convex set under feasibility, so a unique maximizer exists. Introduce a multiplier λ on the inventory constraint and multipliers $\alpha_\tau, \beta_\tau \geq 0$ on the lower and upper bounds. The Lagrangian is

$$\mathcal{L} = -\eta \sum_{\tau=t}^{T-1}(u_\tau^2 + u_\tau v_\tau) + \lambda\left(\sum_{\tau=t}^{T-1} u_\tau + x_t\right) + \sum_{\tau=t}^{T-1}\alpha_\tau(u_\tau + \bar{u}) + \sum_{\tau=t}^{T-1}\beta_\tau(\bar{u} - u_\tau).$$

The KKT conditions imply, for each τ,

$$-2\eta u_\tau - \eta v_\tau + \lambda + \alpha_\tau - \beta_\tau = 0, \qquad \alpha_\tau(u_\tau + \bar{u}) = 0, \qquad \beta_\tau(\bar{u} - u_\tau) = 0,$$

together with $\alpha_\tau, \beta_\tau \geq 0$, $-\bar{u} \leq u_\tau \leq \bar{u}$, and the equality constraint $\sum_{\tau=t}^{T-1} u_\tau = -x_t$. These conditions are equivalent to the pointwise form

$$u_\tau = \mathrm{clip}_{[-\bar{u},\bar{u}]}\left(-\frac{v_\tau}{2} + \frac{\lambda}{2\eta}\right).$$

Setting $c_t = \lambda/(2\eta)$ yields (51).

It remains to justify that c_t can be chosen to satisfy the inventory identity. Define the continuous,

nondecreasing function

$$\Phi(c) = \sum_{\tau=t}^{T-1} \mathrm{clip}_{[-\bar{u},\bar{u}]}\left(-\frac{v_\tau}{2} + c\right).$$

Feasibility implies $-x_t \in [-(T-t)\bar{u}, (T-t)\bar{u}]$, while $\Phi(c) \to -(T-t)\bar{u}$ as $c \to -\infty$ and $\Phi(c) \to (T-t)\bar{u}$ as $c \to +\infty$. Hence there exists c_t with $\Phi(c_t) = -x_t$, and the corresponding $(u_\tau^\star)$ satisfies the equality constraint. If $|u_{\bar{\tau}}^\star| < \bar{u}$ for some $\bar{\tau}$, then in a neighborhood of c_t the term $\mathrm{clip}_{[-\bar{u},\bar{u}]}(-v_{\bar{\tau}}/2 + c)$ is locally equal to $-v_{\bar{\tau}}/2 + c$ and thus varies strictly with c. Therefore Φ is strictly increasing at c_t, which implies that the equation $\Phi(c) = -x_t$ has a unique solution; hence c_t is unique. If all bounds bind, then Φ is locally flat at c_t and the multiplier (hence c_t) need not be unique, even though $(u_\tau^\star)$ is unique by strict concavity.

Finally, if no bound binds at the optimizer, then $\mathrm{clip}_{[-\bar{u},\bar{u}]}$ is the identity and stationarity gives $u_\tau^\star = -v_\tau/2 + c_t$ with a common offset c_t. Imposing $\sum_{\tau=t}^{T-1} u_\tau^\star = -x_t$ yields $c_t = \bar{v}_t/2 - x_t/(T-t)$, which recovers (49). □

E Proofs for Section 6

This appendix collects proofs and technical details for Section 6.

E.1 Pseudoinverse and projections when A is p.s.d.

Let $A \in \mathbb{R}^{N\times N}$ be symmetric and positive semidefinite. Denote by $A^\dagger$ the Moore-Penrose pseudoinverse and define

$$\Pi_A = A^\dagger A = AA^\dagger,$$

the orthogonal projector onto $\mathrm{Range}(A)$; see Boyd and Vandenberghe (2004, Appendix A.5.4).

Lemma 13. *For symmetric and p.s.d. A and Π_A as above:*

1. *For all $z \in \mathbb{R}^N$, $z^\top Az = (\Pi_A z)^\top A(\Pi_A z)$.*
2. *For any $w \in \mathrm{Range}(A)$, the solution set to $Az = w$ is $\{A^\dagger w + n : \ n \in \ker(A)\}$ and every solution satisfies $z^\top Az = w^\top A^\dagger w$.*

Proof. These are standard properties of the Moore-Penrose pseudoinverse and the associated orthogonal projectors for symmetric p.s.d. matrices; see Boyd and Vandenberghe (2004, Appendix A.5.4 and Example 4.5). □

E.2 Proof of Proposition 7

Fix t and s_t, and let $u_t^\star$ be an interior maximizer in (64). Write the date-t Bellman objective as

$$\Psi(u) = \pi_t(s_t, u) + V_{t+1}(\mathcal{T}(s_t, u)).$$

Step 1: first-order condition. Under (60) and symmetry of A,

$$\nabla_u \pi_t(s_t, u) = -A(2u + v_t).$$

The control enters the continuation state through $x_{t+1} = x_t + u$ and through the current print $P_t = S_t + A(u + v_t)$; all other next-state components depend on u only through P_t. By the chain rule,

$$\nabla_u V_{t+1}(\mathcal{T}(s_t, u)) = \nabla_x V_{t+1}(s_{t+1}) \cdot \frac{\partial(x_t + u)}{\partial u} + \frac{d}{dP_t} V_{t+1}(s_{t+1}) \cdot \frac{\partial P_t}{\partial u} = \mu_{t+1} + A\nu_{t+1},$$

where $\mu_{t+1} = \nabla_x V_{t+1}(s_{t+1})$ and $\nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1})$. At an interior optimum, $\nabla_u \Psi(u_t^\star) = 0$, hence

$$-A(2u_t^\star + v_t) + \mu_{t+1} + A\nu_{t+1} = 0,$$

which is (67).

Step 2: decomposition. Since A may be singular, (67) identifies $u_t^\star$ only up to additions in $\ker(A)$. Choose the canonical representative with $u_t^\star \in \mathrm{Range}(A)$, so $\Pi_A u_t^\star = u_t^\star$. Multiplying (67) by $A^\dagger$ gives

$$\Pi_A(2u_t^\star + v_t) = A^\dagger \mu_{t+1} + \Pi_A \nu_{t+1}.$$

Using linearity and $\Pi_A u_t^\star = u_t^\star$, we have

$$\Pi_A(2u_t^\star + v_t) = 2u_t^\star + \Pi_A v_t,$$

hence

$$2u_t^\star = -\Pi_A v_t + A^\dagger \mu_{t+1} + \Pi_A \nu_{t+1},$$

which is (68).

Step 3: payoff identity. By Lemma 5,

$$\pi_t(s_t, u_t^\star) = \frac{1}{4} v_t^\top A v_t - \Big(u_t^\star + \tfrac{v_t}{2}\Big)^\top A \Big(u_t^\star + \tfrac{v_t}{2}\Big) = \frac{1}{4} v_t^\top A v_t - \frac{1}{4}(2u_t^\star + v_t)^\top A (2u_t^\star + v_t).$$

Using (67), $A(2u_t^\star + v_t) = \mu_{t+1} + A\nu_{t+1} \in \mathrm{Range}(A)$. By Lemma 13(2) (see Boyd and Vandenberghe (2004, Example 4.5)),

$$(2u_t^\star + v_t)^\top A (2u_t^\star + v_t) = (\mu_{t+1} + A\nu_{t+1})^\top A^\dagger (\mu_{t+1} + A\nu_{t+1}).$$

Substituting yields (69).

E.3 Proof of Proposition 8

Fix (s_t, u_t) and let $s_{t+1} = \mathcal{T}(s_t, u_t)$. In the transition map, P_t enters the next state directly as the lagged price component, and it also affects $(y_{t+1}, M_{t+1}, X_{t+1})$ via $y_{t+1} = F(y_t, P_t)$, $M_{t+1} = g(\Gamma_t)$, and $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$. Let $\nu_{t+1} = \frac{d}{dP_t} V_{t+1}(s_{t+1})$. By the multivariate chain rule,

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + \big(\partial_P F(y_t, P_t)\big)^\top \nabla_y V_{t+1}(s_{t+1}) + \partial_M V_{t+1}(s_{t+1}) \frac{\partial M_{t+1}}{\partial P_t} + \nabla_X V_{t+1}(s_{t+1})^\top \frac{\partial X_{t+1}}{\partial P_t}.$$

Since $M_{t+1} = g(\Gamma_t)$ with $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$,

$$\frac{\partial M_{t+1}}{\partial P_t} = g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t).$$

Also, $X_{t+1} = (W/M_{t+1})b$, so $\frac{\partial X_{t+1}}{\partial M_{t+1}} = -B_{t+1}$ and therefore

$$\frac{\partial X_{t+1}}{\partial P_t} = -B_{t+1}\, g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t)^\top.$$

Substituting implies

$$\nabla_X V_{t+1}(s_{t+1})^\top \frac{\partial X_{t+1}}{\partial P_t} = -\big(B_{t+1}^\top \nabla_X V_{t+1}(s_{t+1})\big)\, g'(\Gamma_t)\, \partial_P \Gamma(y_t, P_{t-1}, P_t),$$

so grouping the $g'(\Gamma_t)\partial_P \Gamma$ terms yields (71) with Λ_{t+1} defined in (70).

E.4 Proof of Theorem 2

Assume the rule is portfolio marked in direction d as in the theorem statement. Let $\mathcal{W} = \mathrm{span}\{b, d\}$ and write $B = [b\, d] \in \mathbb{R}^{N\times 2}$. Define the 2×2 Gram matrix $G = B^\top AB$ and let $G^\dagger$ be its Moore-Penrose pseudoinverse.

For any $u \in \mathbb{R}^N$ define $\Pi_{\mathcal{W}} u$ as the minimizer of

$$\min_{w\in\mathcal{W}} (u-w)^\top A(u-w).$$

Writing $w = Bc$, this is the convex quadratic problem

$$\min_{c\in\mathbb{R}^2} (u-Bc)^\top A(u-Bc).$$

A standard pseudoinverse solution (Boyd and Vandenberghe, 2004, Appendix A.5.4 and Example 4.5) yields the canonical minimizer

$$c^\star = G^\dagger B^\top Au, \qquad \Pi_{\mathcal{W}} u = Bc^\star = B\, G^\dagger B^\top Au,$$

so $\Pi_{\mathcal{W}} u \in \mathcal{W}$. By optimality, the residual $u - \Pi_{\mathcal{W}} u$ is A-orthogonal to $\mathcal{W}$, i.e.

$$w^\top A(u - \Pi_{\mathcal{W}} u) = 0 \qquad \text{for all } w \in \mathcal{W}. \tag{83}$$

In particular, taking $w = b$ and $w = d$ gives

$$b^\top A\, \Pi_{\mathcal{W}} u = b^\top Au, \qquad d^\top A\, \Pi_{\mathcal{W}} u = d^\top Au. \tag{84}$$

Moreover, (83) implies the Pythagorean identity

$$u^\top Au = (\Pi_{\mathcal{W}} u)^\top A(\Pi_{\mathcal{W}} u) + (u - \Pi_{\mathcal{W}} u)^\top A(u - \Pi_{\mathcal{W}} u), \tag{85}$$

so the projection weakly decreases the A-norm:

$$(\Pi_{\mathcal{W}} u)^\top A(\Pi_{\mathcal{W}} u) \le u^\top Au. \tag{86}$$

Fix any admissible round trip $(u_t)_{t=0}^{T-1}$ for the relaxed problem (terminal constraint $x_T = 0$ and no per-date bounds). Define the projected strategy $\tilde{u}_t = \Pi_{\mathcal{W}} u_t$ for each t. Linearity of $\Pi_{\mathcal{W}}$ implies $\sum_{t=0}^{T-1} \tilde{u}_t = \Pi_{\mathcal{W}} \sum_{t=0}^{T-1} u_t = 0$, so $(\tilde{u}_t)$ is also a round trip. Let $z_t = d^\top P_t$ denote the marked print. We show by induction that z_t is identical under (u_t) and $(\tilde{u}_t)$. At $t = 0$, $v_0 = X_0 - X_{-1}$ is pinned down by the initial conditions and is the same under both strategies. Assume the histories $(z_0, \dots, z_{t-1})$ coincide. Because the rule is portfolio marked, $(z_0, \dots, z_{t-1})$ implies the same (y_t, M_t, X_t) and hence the same predetermined forced flow $v_t = X_t - X_{t-1}$ under both strategies. Under (59),

$$z_t = d^\top S_t + d^\top A(u_t + v_t).$$

Using (84) with $u = u_t$ gives $d^\top A\tilde{u}_t = d^\top Au_t$, hence z_t is the same under both strategies. This closes

the induction. Consequently the requirement and forced-flow sequences coincide under (u_t) and $(\tilde{u}_t)$.

Because v_t is unchanged, compare profits period by period:

$$\pi_t(u) = -u^\top A(u + v_t) = -u^\top Au - u^\top Av_t.$$

Since $v_t \in \mathrm{span}\{b\}$, the interaction term depends on u only through $b^\top Au$. By (84) with $w = b$, this interaction term is preserved when replacing u_t by $\tilde{u}_t$. By (86), the direct impact loss weakly falls: $\tilde{u}_t^\top A\tilde{u}_t \le u_t^\top Au_t$. Therefore $\pi_t(\tilde{u}_t) \ge \pi_t(u_t)$ for each t, and summing yields weakly higher total wedge profits. Thus, from any admissible strategy for the relaxed problem we can construct another admissible strategy trading only in $\mathcal{W} = \mathrm{span}\{b, d\}$ that attains weakly higher total expected wedge profits. Hence the supremum value is unchanged under the span restriction, and if attained an optimizer can be chosen in $\mathcal{W}$. Finally, once $u_t \in \mathcal{W}$ and $v_t \in \mathrm{span}\{b\}$, all quadratic wedge terms depend on A only through its restriction to $\mathcal{W}$, equivalently through the Gram matrix $G = B^\top AB$ with entries $b^\top Ab$, $b^\top Ad$, and $d^\top Ad$.

E.5 KKT conditions under the ℓ_1 per-date trade bound

This subsection records the first-order conditions for the constrained one-period maximization in (64) when the ℓ_1 constraints bind. Fix t and s_t and assume V_{t+1} is differentiable at the relevant continuation state. Let

$$G_t(u) = \pi_t(s_t, u) + V_{t+1}(\mathcal{T}(s_t, u)).$$

Its gradient is

$$\nabla_u G_t(u) = -A(2u + v_t) + \mu_{t+1}(u) + A\nu_{t+1}(u),$$

where $\mu_{t+1}(u) = \nabla_x V_{t+1}(\mathcal{T}(s_t, u))$ and $\nu_{t+1}(u) = \nabla_{P_t} V_{t+1}(\mathcal{T}(s_t, u))$. The feasible set is the intersection of two ℓ_1 balls,

$$\|u\|_1 \le \bar{u}, \qquad \|x_t + u\|_1 \le (T - 1 - t)\bar{u}.$$

A (possibly) binding optimum $u_t^\star$ admits multipliers $\lambda_t, \kappa_t \ge 0$ and subgradients $\zeta_t \in \partial \|u_t^\star\|_1$, $\xi_T \in \partial \|x_t + u_t^\star\|_1$ such that

$$0 \in \nabla_u G_t(u_t^\star) + \lambda_t \zeta_t + \kappa_t \xi_T, \tag{87}$$

$$\lambda_t \left(\|u_t^\star\|_1 - \bar{u}\right) = 0, \qquad \kappa_t \left(\|x_t + u_t^\star\|_1 - (T - 1 - t)\bar{u}\right) = 0. \tag{88}$$

Here $\partial \|z\|_1$ is the standard ℓ_1 subdifferential: componentwise, $\zeta_i = \mathrm{sign}(z_i)$ if $z_i \neq 0$ and $\zeta_i \in [-1, 1]$ if $z_i = 0$. Equation (87) reduces to the interior Euler equation (67) when both constraints are slack.

F Volatility-managed funds

Volatility-controlled (risk-control, target-volatility) indices publish a deterministic exposure rule: given recent realized returns, they set the next-day risky exposure, typically subject to caps, floors, and implementation lags. When these indices underlie structured products and index-linked annuities, dealer replication turns the published rule into mechanical hedging flow in the underlying. This is exactly the environment the linearized admissibility condition in Theorem 1 is meant to evaluate.

F.1 Institutional setting

Volatility-controlled indices are widely used as underlyings for retail structured products and insurance wrappers. In these contracts the payoff is indexed to the published index level, so the issuer or insurer must replicate the index exposure to hedge its liabilities. The market is relatively large: U.S. retail annuity sales were on the order of hundreds of billions of dollars in 2024, and fixed indexed annuities (FIAs) are a major and growing segment (LIMRA, 2025; American Academy of Actuaries, 2026b). Index providers explicitly design and license risk-control indices for product embedding, including structured products and indexed investment products (MSCI, 2020, 2021; S&P Dow Jones Indices, 2025). MSCI also documents a broad set of U.S. indexed annuities (FIA/RILA/IUL) that reference MSCI indexes, illustrating how index methodologies translate into balance-sheet exposures (MSCI, 2025). Regulators treat the index rule and crediting methodology as core disclosures for index-linked annuities (Securities and Exchange Commission, 2024).

F.1.1 From rulebooks to trading flow

Three parties map the disclosed rulebook into trading flow relevant for our admissibility test.

1. *Index sponsor and calculation agent.* The sponsor specifies the methodology (sampling convention, volatility estimator, exposure rule, caps/floors, fees, and any lags) and publishes the official index level. These documents pin down the feedback map from prices to next-day exposure (MSCI, 2021; S&P Dow Jones Indices, 2025; BlackRock, 2023).

2. *Issuer/insurer.* The issuer sells a contract whose payoff references the index level (structured note) or credits interest using an index-linked option (FIA/RILA). The liability is therefore indexed to the rule (American Academy of Actuaries, 2026b; Securities and Exchange Commission, 2024).

3. *Dealer replication and hedge instruments.* Replication is implemented in futures, ETFs, swaps, and cash. When the index rebalances, replication requires trading these hedge instruments. Conditional on the published rule, the direction and timing of the flow are mechanical.

F.1.2 Single-asset versus multi-asset indices

Most marketed designs share a common architecture: a reference portfolio and a cash account, with a volatility-control constraint that scales exposure to the reference portfolio to target a stated risk level

(MSCI, 2021; S&P Dow Jones Indices, 2025, 2021a; Bloomberg Index Services Limited, 2025). They differ mainly in the construction of the reference portfolio.

- *Single-underlying (Template A).* The reference portfolio is a single parent index. The risk-control rule varies the weight on the parent index versus cash as an inverse-volatility function of a rolling volatility estimator (MSCI, 2021; S&P Dow Jones Indices, 2025).
- *Multi-underlying, fixed basket (Template B).* The reference portfolio is a fixed multi-asset basket (often equity and rates, implemented with ETFs or futures), and the rule scales the basket versus cash to hit a volatility target (S&P Dow Jones Indices, 2025; BlackRock, 2023).
- *Multi-underlying, signal-driven basket (Template C).* The reference portfolio is itself rule-based and time-varying (e.g., rotation, momentum/trends), and a volatility-control rule then scales the resulting portfolio versus cash (J.P. Morgan, 2015; BNP Paribas Indices, 2025; UBS Investment Bank, 2025).

We focus on Templates A and B.

F.2 Formulae and timing conventions

Our baseline formulas are taken from public index methodology documents and representative offering materials for structured products and annuity crediting options. Methodologies specify the estimator, caps/floors, turnover buffers, fees, the cash component, and the calculation day. Disclosure documents typically restate the same mechanics and, importantly for our timing, make explicit the lag between volatility measurement and exposure implementation (MSCI, 2024; S&P Dow Jones Indices, 2021b; BlackRock, 2023; Barclays Bank, 2021; Morgan Stanley, 2025).

F.2.1 Template A: single-underlying daily clipped target-vol

The single-underlying template allocates between a parent index (e.g., MSCI World Index) and cash. Let B_t be the parent index level used for the rule, sampled once per business day, and let $r_t = \log(B_t/B_{t-1})$ denote the daily log return used for volatility estimation. An annualized realized volatility estimate is constructed as a function of recent returns. For example, MSCI computes a short-horizon realized volatility and a long-horizon realized volatility using equally weighted daily gross total returns and sets Parent Index volatility as the maximum of the two (with $N = 20$ and $N = 60$ trading days) (MSCI, 2024). S&P documents an analogous structure, allowing simple-weighted or exponentially weighted volatility and using the maximum of short- and long-term estimates (S&P Dow Jones Indices, 2021b). Given a target volatility σ^*, the exposure to the parent index is set as a ratio of target to estimated volatility, subject to a maximum leverage cap $L_{\max}$,

$$L_t = \min\left\{L_{\max},\ \frac{\sigma^*}{\widehat{\sigma}_{t-\ell}}\right\},$$

with the residual weight $1 - L_t$ allocated to a cash component (and, when $L_t > 1$, borrowing costs applied to the leveraged portion). The lag ℓ is explicit in some rulebooks and disclosures. For instance,

MSCI determines index leverage for an effective date using volatility estimated two trading days before the effective date and applies a turnover buffer so leverage is updated only when changes exceed a threshold (MSCI, 2024). Representative offering documents for risk-control underliers also describe a two-day lag between leverage-factor calculation and implementation and disclose the leverage cap (e.g., 150%) (Morgan Stanley, 2025; Barclays Bank, 2021).

F.2.2 Template B: multi-underlying risk-control

In the multi-underlying template, the reference portfolio is a basket of tradable constituents, and the rule scales the basket versus cash to track a volatility target. The basket composition is fixed and only the overall risk budget varies. Formally, let $P_t \in \mathbb{R}^N$ denote the vector of constituent price levels used for index calculation (sampled once per business day), and let $d \in \mathbb{R}^N$ be fixed risky-basket weights. The reference-basket return is $r_t^{\text{ref}} = d^\top r_t$ where r_t are constituent returns. The rule computes a realized volatility estimate $\widehat{\sigma}_t$ from the history of r^{ref} (or an equivalent basket-level statistic) and sets a scalar leverage factor L_t as above, subject to caps/floors and a lag. The resulting risky weights are $w_t = L_t d$, and the cash weight is the residual $1 - \mathbf{1}^\top w_t$. A concrete industry example is the BlackRock Adaptive U.S. Equity Index Series, which specifies a universe of ETF constituents (including an equity ETF and Treasury ETFs) plus a cash constituent, calculates index levels from constituent closing prices, and applies constituent weights with an explicit one-business-day lag (BlackRock, 2023). The methodology also documents rounding conventions for weights and an excess-return index return calculation relative to an interest rate and an index fee (BlackRock, 2023).

F.2.3 Timing conventions

Three timing details matter for the no-arbitrage and stress-test computations.

1. Methodologies commonly use end-of-day closing prices to compute index levels and the statistics entering the rebalancing rule (BlackRock, 2023).

2. Many designs implement leverage and weights with a one- or two-business-day lag. This lag is explicit in methodology documents (MSCI, 2024; Morgan Stanley, 2025).

3. Caps, floors, turnover buffers, rounding, and exchange-holiday conventions create kinks and discrete updates (MSCI, 2024; BlackRock, 2023).

F.3 Derivatives and construction of the admissibility test

An index methodology gives a deterministic rule from recent returns to next-day risky weight, including caps, floors, and any implementation lag. We provide the map with the terms used in the paper.

F.3.1 (Γ, g) and $s = g'(\Gamma)$

For a standard target-volatility rule,

$$w_{t+1} = \mathrm{clip}\left(w_{\min},\, w_{\max},\, \frac{\sigma^*}{\widehat{\sigma}_{t-\ell}}\right), \qquad M_{t+1} = \frac{1}{w_{t+1}}, \tag{89}$$

with $\widehat{\sigma}_t = \sqrt{\Gamma_t}$. Many methodologies add a turnover buffer: the weight is held fixed unless the raw update exceeds a threshold τ,

$$w^{\mathrm{raw}}_{t+1} = \mathrm{clip}\left(w_{\min},\, w_{\max},\, \frac{\sigma^*}{\widehat{\sigma}_{t-\ell}}\right), \qquad w_{t+1} = \begin{cases} w_t, & \text{if } |w^{\mathrm{raw}}_{t+1} - w_t| \le \tau, \\ w^{\mathrm{raw}}_{t+1}, & \text{if } |w^{\mathrm{raw}}_{t+1} - w_t| > \tau, \end{cases} \qquad M_{t+1} = \frac{1}{w_{t+1}}.$$

When the buffer binds (first case), exposure is locally flat in $\widehat{\sigma}_{t-\ell}$, so $s_t = 0$. On the interior region (cap/floor inactive and buffer inactive),

$$M_{t+1} = \frac{\widehat{\sigma}_{t-\ell}}{\sigma^*} = \frac{\sqrt{\Gamma_{t-\ell}}}{\sigma^*}, \qquad s_t = \frac{\partial M_{t+1}}{\partial \Gamma_{t-\ell}} = \frac{1}{2\sigma^* \widehat{\sigma}_{t-\ell}}. \tag{90}$$

F.3.2 J for Template A

Template A uses a single price series P_t and a rolling realized-variance statistic. Define log returns and the statistic by

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right), \qquad \Gamma_t = \frac{a_{\mathrm{ann}}}{m} \sum_{i=0}^{m-1} r^2_{t-i}, \tag{91}$$

with window length m and annualization factor $a_{\mathrm{ann}} > 0$. Some rulebooks compute realized variance at multiple horizons (e.g., 20 and 60 days) and set Γ_t equal to the maximum across horizons. For a finite set $\mathcal{M}$, define

$$\Gamma^{(m)}_t = \frac{a_{\mathrm{ann}}}{m} \sum_{i=0}^{m-1} r^2_{t-i}, \qquad m \in \mathcal{M}, \qquad \Gamma_t = \max_{m \in \mathcal{M}} \Gamma^{(m)}_t.$$

Away from ties, let m^*_t be the unique maximizing window and take the Jacobian row at date t from that active window. At ties we choose the smallest window. Differentiating (91) gives, for any date j,

$$\frac{\partial \Gamma_t}{\partial P_j} = \frac{2a_{\mathrm{ann}}}{m\, P_j} \Big(\mathbf{1}\{t - m + 1 \le j \le t\}\, r_j - \mathbf{1}\{t - m \le j \le t - 1\}\, r_{j+1} \Big). \tag{92}$$

On a T-date horizon, writing $\delta\Gamma = J\, \delta\Delta P$, the matrix J is the $T \times T$ matrix with (t, j) entry given by (92) evaluated at the reference state. Only prices inside the rolling window affect Γ_t.

F.3.3 J for Template B

Template B scales exposure to a fixed risky basket against cash. Let $P_t \in \mathbb{R}^N$ collect the constituent prices used by the methodology, and define constituent log returns

$$r_t^k = \log\left(\frac{P_t^k}{P_{t-1}^k}\right), \qquad k = 1, \dots, N. \tag{93}$$

Fix basket weights $d \in \mathbb{R}^N$ and write the basket return as $\bar{r}_t = d^\top r_t$. Realized-variance is

$$\Gamma_t = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} \bar{r}_{t-i}^2. \tag{94}$$

As in Template A, some rulebooks take the maximum across horizons. For implementation we break ties in favor of the smallest window. Differentiating (94) with respect to a constituent price P_j^k yields

$$\frac{\partial \Gamma_t}{\partial P_j^k} = \frac{2a_{\text{ann}}}{m\, P_j^k}\, d_k \Big(\mathbf{1}\{t-m+1 \le j \le t\}\, \bar{r}_j - \mathbf{1}\{t-m \le j \le t-1\}\, \bar{r}_{j+1} \Big). \tag{95}$$

Stack deviations over dates and assets as

$$\delta\Delta P = \big((\delta\Delta P_0)^\top, \dots, (\delta\Delta P_{T-1})^\top \big)^\top \in \mathbb{R}^{NT}.$$

Then (95) defines$J$ such that $\delta\Gamma = J\, \delta\Delta P$, as in (23)..

F.3.4 $\mathcal{K}$ and $\widehat{H}$

Given J and s, requirement deviations satisfy

$$\delta M = s\, L^{\ell+1}\, \delta\Gamma = s\, L^{\ell+1}\, J\, \delta\Delta P, \tag{96}$$

where L is the one-step lag operator and ℓ is the disclosed implementation lag, so the total shift is $\ell + 1$. Forced flow comes from the change in constrained demand. Assume the constrained sector trades along a fixed liquidation direction $b \in \mathbb{R}^N$ (for the single-underlying template, $N = 1$ and $b = 1$), so $X_t = X(M_t)\, b$ with local sensitivity $B = -X'(M_0) > 0$ as in (24). Stacking over dates,

$$\delta X = -B\, (I_T \otimes b)\, \delta M, \qquad v = (D \otimes I_N)\, \delta X = -Bs\, (D \otimes I_N)(I_T \otimes b)\, L^{\ell+1}\, J\, \delta\Delta P, \tag{97}$$

where D is the first-difference operator on the T-date horizon (applied blockwise via $D \otimes I_N$). Wedge prices are generated by total flow: $\delta\Delta P = \mathcal{I} q$, with $q = u + v$. Substituting (97) yields

$$q = u - \mathcal{K} q, \qquad \mathcal{K} = Bs\, (D \otimes I_N)(I_T \otimes b)\, L^{\ell+1} J \mathcal{I}. \tag{98}$$

If $I + \mathcal{K}$ is invertible, then

$$q = (I + \mathcal{K})^{-1} u, \qquad \delta\Delta P = \mathcal{I}(I + \mathcal{K})^{-1} u. \tag{99}$$

Define $\widetilde{\mathcal{I}} = \mathcal{I}(I+\mathcal{K})^{-1}$. Theorem 1 applies to

$$\widehat{H} = \widetilde{\mathcal{I}} + \widetilde{\mathcal{I}}^{\top} = \mathcal{I}(I+\mathcal{K})^{-1} + \Big((I+\mathcal{K})^{-1}\Big)^{\top}\mathcal{I}^{\top}, \tag{100}$$

and the manipulation-free condition is that $\widehat{H}$ is positive semidefinite.

F.3.5 A one-step closed-form diagnostic

To gain some intuition, this subsection gives a one-step version of the admissibility condition for the daily target-volatility rule. Assume a single asset and $\delta\Delta P_t = \eta q_t$. Let

$$\Gamma_t = \frac{a_{\text{ann}}}{m}\sum_{j=0}^{m-1} r_{t-j}^2, \qquad r_t = \log(P_t/P_{t-1}),$$

and caps, floors etc do not bind so $M_{t+1} = \widehat{\sigma}_t/\sigma^*$ with $\widehat{\sigma}_t = \sqrt{\Gamma_t}$. Then

$$s_t = \frac{\partial M_{t+1}}{\partial \Gamma_t} = \frac{1}{2\sigma^*\widehat{\sigma}_t}, \qquad \frac{\partial \Gamma_t}{\partial P_t} = \frac{2a_{\text{ann}}}{m}\frac{r_t}{P_t}.$$

We have $B = W/M_0^2$. Evaluating at $M_0 = M_{t+1} = \widehat{\sigma}_t/\sigma^*$ gives

$$k_t = B\, s_t\, \eta\, \frac{\partial \Gamma_t}{\partial P_t} = \eta\, \frac{W}{P_t}\, \frac{\sigma^* a_{\text{ann}}}{m\,\widehat{\sigma}_t^3}\, r_t, \tag{101}$$

Corollary 6 (One-step inequality for the daily target-vol rule)**.** *Consider $T = 2$ trading dates and the immediate $t \to t+1$ channel under temporary impact $\mathcal{I} = \eta I_2$. In this one-step specialization, admissibility on round trips is equivalent to*

$$2 + k_t \geq 0, \tag{102}$$

where k_t is given by (101). Equivalently, writing $r_t^- = \max\{-r_t, 0\}$,

$$\eta\, \frac{W}{P_t}\, \frac{\sigma^* a_{\text{ann}}}{m\,\widehat{\sigma}_t^3}\, r_t^- \leq 2. \tag{103}$$

Proof. In the one-step specialization, $\mathcal{K}$ is 2×2 strictly lower triangular with single entry k_t in the $(2,1)$ position. With $\mathcal{I} = \eta I_2$,

$$\widehat{H} = \mathcal{I}(I+\mathcal{K})^{-1} + \big((I+\mathcal{K})^{-1}\big)^{\top}\mathcal{I}^{\top}.$$

A direct calculation gives $\widehat{H} = \eta \begin{pmatrix} 2 & -k_t \\ -k_t & 2 \end{pmatrix}$. Restricting to the two-date round-trip subspace $\{(u,-u) : u \in \mathbb{R}\}$ yields a quadratic form proportional to $2 + k_t$, which gives (102). Rearranging yields (103). □

Rearranging (103) gives the one-step capacity implied by the immediate channel:

$$W_{\max}^{(1\text{-}step)}(t) = \frac{2m\, P_t\, \widehat{\sigma}_t^3}{\eta\, \sigma^* a_{\text{ann}}\, r_t^-}, \qquad W_{\max}^{(1\text{-}step)}(t) = +\infty \text{ if } r_t^- = 0. \tag{104}$$

The one-step bound loosens with a longer window m, and it is irrelevant whenever the exposure rule is locally flat (caps, floors, or buffers bind) or when an additional lag removes dependence on r_t.

F.4 Stress-test design and empirical implementation details

The stress test needs two inputs: the sensitivity of trades with respect to price changes (which is obtained form the volatility index rules), and an estimate of market impact. Fix a state z. At that point we can compute both, We assume price impact per fraction of volume traded is constant across states, therefore at that point we have a varying sensitivity of trades with respect to price changes.

We construct $\mathcal{Z}$ from simulated return paths. For Template A we simulate a single daily log-return process. For Template B we simulate an N-vector of daily log returns with a realistic correlation structure. The baseline generator is multivariate GARCH(1,1):

$$r_t^k = \sqrt{h_t^k}\,\varepsilon_t^k, \qquad h_{t+1}^k = \omega_k + \alpha\,(r_t^k)^2 + \beta\,h_t^k, \qquad k = 1, \dots, N,$$

with $\alpha \geq 0$, $\beta \geq 0$, and $\alpha+\beta < 1$. Innovations $\varepsilon_t \in \mathbb{R}^N$ are i.i.d. over t with mean zero, unit variances, and a specified correlation matrix; the baseline uses a multivariate Student-t distribution with ν degrees of freedom, scaled to unit variance componentwise. For correlation stress in Template B, we multiply off-diagonal correlations by $\kappa \geq 1$, cap absolute correlations at 0.99, and project to the nearest valid correlation matrix before simulation.

Numerical calibration is as follows. We simulate $n_{\text{paths}} = 10$ independent paths of length 2500 business days and discard the first 500 days as burn-in. From each path we sample 40 evaluation dates uniformly at random (without replacement), so $|\mathcal{Z}| = 400$ states per template and scenario. Innovations are Student-t with $\nu = 8$. We consider $(\alpha, \beta) \in \{(0,0), (0.05, 0.90), (0.05, 0.94)\}$. With $a_{\text{ann}} = 252$, we set $\omega_k = (1-\alpha-\beta)\,\sigma_{\text{ann},k}^2/a_{\text{ann}}$ so that unconditional annualized volatility equals $\sigma_{\text{ann},k}$. For Template A we set $\sigma_{\text{ann}} = 0.20$. For Template B we consider $N \in \{2, 4, 8\}$ with baseline off-diagonal correlation 0.2 (equicorrelation for $N \in \{4, 8\}$) and stress factors $\kappa \in \{1.0, 1.5, 2.0\}$. Annualized volatilities and fixed weights (d, b) are:

$$\begin{aligned}
N = 2: &\quad \sigma_{\text{ann}} = (0.20, 0.10), \quad d = b = (0.6, 0.4), \\
N = 4: &\quad \sigma_{\text{ann}} = (0.20, 0.15, 0.12, 0.10), \quad d = b = \tfrac{1}{4}\,\mathbf{1}, \\
N = 8: &\quad \sigma_{\text{ann}} = (0.22, 0.20, 0.18, 0.16, 0.14, 0.12, 0.10, 0.08), \quad d = b = \tfrac{1}{8}\,\mathbf{1}.
\end{aligned}$$

We calibrate $\mathcal{I}$ from one-day execution-cost benchmarks expressed in basis points per 1% of average daily volume (ADV). Two complementary references guide the range. Kyle and Obizhaeva (2016) report total costs of 10.71 bps (linear) and 14.16 bps (square-root) for a 1% ADV order (Table V), motivating 12 bps per 1% ADV as a baseline and $[8, 20]$ bps as a conservative range. Frazzini et al. (2018) report mean market impact of 8.90 bps for large-cap trades and 18.95 bps for small-cap trades (Table II, Panel A), supporting the upper end of the range in stressed scenarios. We implement these magnitudes as a local linear impact law, which is the object required by the theory: in the single-asset case, $\mathcal{I}$ is set so that trading 1% of ADV generates a one-day cost of $c_{1\%} \in \{8, 12, 20\}$ bps. In the multi-asset case we use a common $c_{1\%}$ across constituents. Cross-impact is added as a separate sensitivity. Our baseline sets $\mathcal{I}$ diagonal (most conservative and most transparent) and then

considers a symmetric cross-impact parameterization

$$\mathcal{I}_{ij} = \rho \sqrt{\mathcal{I}_{ii}\mathcal{I}_{jj}}\, \mathrm{corr}_{ij} \qquad (i \neq j),$$

with $\rho \in [0, 0.3]$ and baseline $\rho = 0.2$. We report results under $\rho = 0$ and over the cross-impact range. In computation we enforce that stressed correlation matrices are valid by projection to the nearest positive semidefinite correlation matrix. When adding cross-impact terms, we symmetrize the resulting impact matrix and apply a diagonal adjustment if needed to ensure it is positive semidefinite. These numerical safeguards address knife-edge cases and do not change the analytical objects used by the theorem.

G Additional properties of optimal attacks against target-volatility rules

This appendix collects auxiliary derivative and timing facts for the interior target-volatility rule in (72) and the rolling realized-volatility statistic (74). Section 7.2 emphasizes the main mechanisms; here we provide additional structure that is useful for interpreting timing and state dependence.

G.1 Local derivative structure: sparsity, finite memory, and the two-adjacent-returns channel

- From (76), $J_{t,j} \neq 0$ only if $j \in \{t-m, t-m+1, \dots, t\}$. Each row therefore depends on at most the last $m+1$ prints. Equivalently, a print P_j can affect Γ_t only for $t \in \{j, j+1, \dots, j+m\}$ (through r_j and r_{j+1}). With lag ℓ, that print can affect posted requirements only over dates $t+1 \in \{j+\ell+1, \dots, j+\ell+m+1\}$.

- A marginal increase in P_j raises $r_j = \log(P_j/P_{j-1})$ and lowers $r_{j+1} = \log(P_{j+1}/P_j)$. Because Γ is built from squared returns, a spike-and-revert pattern can raise Γ even if the price level ends near its starting point.

- Setting $j = t$ in (76) gives
$$\frac{\partial \Gamma_t}{\partial P_t} = \frac{2a_{\text{ann}}}{m\,P_t}\, r_t.$$
To raise Γ_t locally, P_t must move in the direction of the current return r_t. The $1/P_t$ factor implies that what matters is the relative move $\delta P_t/P_t$.

- If $\Gamma_t = \max_{m \in \mathcal{M}} \Gamma_t^{(m)}$, then away from ties the Jacobian row equals that of the active window, so the same sparsity and finite-memory properties hold with $m = m_t^*$. At ties the map is kinked and J is not unique; the main text focuses on interior regions away from such kinks.

G.2 Timing structure: entry and exit, roll-off reversal, and horizon restrictions

- Fix a date j and consider the single-return perturbation in (75). Since requirements depend on $\Gamma_{t-\ell}$, the induced requirement change is confined to the shifted block $t+1 \in \{j+\ell+1, \dots, j+\ell+m\}$. To see the implied flow, linearize the binding position map $X_{t+1} = W/M_{t+1}$:
$$\delta X_{t+1} \approx -B_{t+1}\,\delta M_{t+1}, \qquad B_{t+1} = \frac{W}{M_{t+1}^2}.$$
If δM_{t+1} is block-constant over the affected dates, then δX_{t+1} is also block-constant, and forced flow is its first difference, $\delta v_{t+1} = \delta X_{t+1} - \delta X_t$. Hence δv_{t+1} is concentrated at the block boundaries: deleveraging when the perturbed return enters the window, and releveraging when it rolls off (both shifted by $\ell+1$ relative to date j).

- A spike at date j typically moves both r_j and r_{j+1} (the same print enters two adjacent returns). In the interior approximation of Section 7.2, $\delta\Gamma_t$ is therefore the sum of two overlapping blocks, supported on $\{j, \dots, j+m-1\}$ and on $\{j+1, \dots, j+m\}$. Passing this through (72) and $X = W/M$ generates a two-step tightening of requirements (and deleveraging flow), followed

by a two-step easing at roll-off. If the horizon includes both exit dates, there are two predictable reversal windows rather than one.

- A return at date j leaves the m-day window at date $j+m$. With lag ℓ, the associated easing in requirements occurs at date $j+\ell+m+1$. To capture both the entry-driven deleveraging and the exit-driven releveraging within a T-date horizon, the trigger must satisfy

$$j+\ell+m+1 \le T-1, \qquad \text{equivalently} \qquad j \le T-2-\ell-m.$$

If $T < \ell+m+1$, the roll-off event lies outside the horizon, so the optimal deviation can only harvest the entry-driven tightening. This sharp threshold is specific to rolling windows; with infinite-memory statistics there is no discrete roll-off date.

- With lag ℓ, the earliest forced flow induced by a date-t manipulation arrives at $t+\ell+1$. A larger ℓ both moves the last feasible trigger date earlier and delays any harvesting. In the terminal region $t \ge T-1-\ell$, the margin-feedback motive disappears, and the continuation reduces to the closed-form intermediation benchmark in Proposition 5.

G.3 State dependence and sign patterns in the optimal trade

- When forced flow is predetermined ($v_t \neq 0$), the myopic harvesting component in (73) is $-v_t/2$, so the attacker intermediates against the mechanical trade. By contrast, the manipulation motive is driven by $\Lambda_{t+1} s_t \, \partial_P \Gamma$ (Proposition 4). Since

$$\partial_{P_t} \Gamma_t = \frac{2 a_{\mathrm{ann}}}{m P_t} \, r_t,$$

raising Γ_t locally requires moving P_t in the direction of the current return: buy after an up move and sell after a down move. Hence, within one optimal policy, trigger trades are locally trend-amplifying (because the statistic is convex in returns), while harvest trades are contrarian (because profits come from providing liquidity against v_t). This trigger-versus-harvest sign flip is a distinctive feature of feedback through a risk statistic, rather than standard price-level manipulation: it is sharp here because Γ is built from squared returns, and it is much weaker for constraints written on price levels or signed returns.

- On the interior of the target-volatility rule,

$$B_{t+1} s_t = \frac{W \sigma^*}{2 \, \Gamma_{t-\ell}^{3/2}},$$

so the local gain rises with W and σ^* and falls with $\Gamma_{t-\ell}$. Holding the impact environment fixed, the loop operator and the admissibility test in Section 7.3 are therefore most stringent in low-volatility states: when $\widehat{\sigma}_{t-\ell}$ is small, the same linked notional implies larger risky-unit positions and a larger marginal pass-through from measured risk to forced flow. This is why conservative stress-test capacity is often pinned down by "quiet" states.

- The target σ^* scales both the position level and the gain through $B_{t+1}s_t$, so higher σ^* increases vulnerability holding W fixed. Window length m enters J as a $1/m$ attenuation, but it also delays roll-off of a manipulated return. Thus larger m makes the statistic harder to move on impact, while increasing the value of a long horizon that can harvest the roll-off reversal. Both effects come directly from the rulebook structure $M \propto \widehat{\sigma}$ and the hard rolling-window design.

H Additional results for the volatility-control stress tests

This appendix complements Section 7.2 by reporting (i) $W_{\max}(\mathcal{Z})$ for every scenario and horizon and (ii) additional diagnostic figures for the cross-sectional distribution of $W_{\max}(z)$ and related objects.

Table 4 Template A: conservative admissible scale $W_{\max}(\mathcal{Z})$ by scenario and horizon. Entries are $W_{\max}(\mathcal{Z}) = \min_{z\in\mathcal{Z}} W_{\max}(z)$ in ADV-normalized units.

Scenario	$T = 63$	$T = 126$	$T = 252$
baseline	0.191	0.165	0.081
$\sigma^* = 10\%$	0.239	0.206	0.101
$\sigma^* = 15\%$	0.159	0.137	0.068
$c_{1\%} = 8$ bps	0.287	0.247	0.122
$c_{1\%} = 20$ bps	0.115	0.099	0.049
iid const vol	0.254	0.181	0.132
GARCH persistent	0.110	0.120	0.035
lag $\ell = 0$	0.171	0.150	0.087
lag $\ell = 2$	0.178	0.153	0.086
window 20 only	0.153	0.141	0.086
cross-impact $\rho = 0$	0.191	0.165	0.081
cross-impact $\rho = 0.3$	0.191	0.165	0.081
corr stress $\kappa = 1.5$	0.191	0.165	0.081
corr stress $\kappa = 2.0$	0.191	0.165	0.081

Table 5 Template B, $N = 2$: conservative admissible scale $W_{\max}(\mathcal{Z})$ by scenario and horizon. Entries are $W_{\max}(\mathcal{Z}) = \min_{z\in\mathcal{Z}} W_{\max}(z)$ in ADV-normalized units.

Scenario	$T = 63$	$T = 126$	$T = 252$
baseline	0.158	0.125	0.115
$\sigma^* = 10\%$	0.197	0.156	0.144
$\sigma^* = 15\%$	0.127	0.128	0.113
$c_{1\%} = 8$ bps	0.237	0.188	0.173
$c_{1\%} = 20$ bps	0.095	0.075	0.069
iid const vol	0.162	0.186	0.146
GARCH persistent	0.120	0.094	0.099
lag $\ell = 0$	0.135	0.148	0.110
lag $\ell = 2$	0.125	0.134	0.107
window 20 only	0.103	0.095	0.082
cross-impact $\rho = 0$	0.163	0.129	0.119
cross-impact $\rho = 0.3$	0.155	0.123	0.113
corr stress $\kappa = 1.5$	0.168	0.128	0.123
corr stress $\kappa = 2.0$	0.173	0.132	0.099

Table 6 Template B, $N = 4$: conservative admissible scale $W_{\max}(\mathcal{Z})$ by scenario and horizon. Entries are $W_{\max}(\mathcal{Z}) = \min_{z \in \mathcal{Z}} W_{\max}(z)$ in ADV-normalized units.

Scenario	$T = 63$	$T = 126$	$T = 252$
baseline	0.129	0.122	0.209
$\sigma^* = 10\%$	0.229	0.181	0.218
$\sigma^* = 15\%$	0.149	0.115	0.234
cross-impact $\rho = 0$	0.144	0.135	0.229
cross-impact $\rho = 0.3$	0.123	0.116	0.200
corr stress $\kappa = 1.5$	0.219	0.143	0.169
corr stress $\kappa = 2.0$	0.161	0.143	0.211

Table 7 Template B, $N = 8$: conservative admissible scale $W_{\max}(\mathcal{Z})$ by scenario and horizon. Entries are $W_{\max}(\mathcal{Z}) = \min_{z \in \mathcal{Z}} W_{\max}(z)$ in ADV-normalized units.

Scenario	$T = 63$	$T = 126$	$T = 252$
baseline	0.311	0.230	0.286
$\sigma^* = 10\%$	0.360	0.265	0.340
$\sigma^* = 15\%$	0.373	0.339	0.407
cross-impact $\rho = 0$	0.395	0.287	0.360
cross-impact $\rho = 0.3$	0.281	0.209	0.259
corr stress $\kappa = 1.5$	0.340	0.242	0.341
corr stress $\kappa = 2.0$	0.376	0.256	0.358

H.1 Vulnerability curves

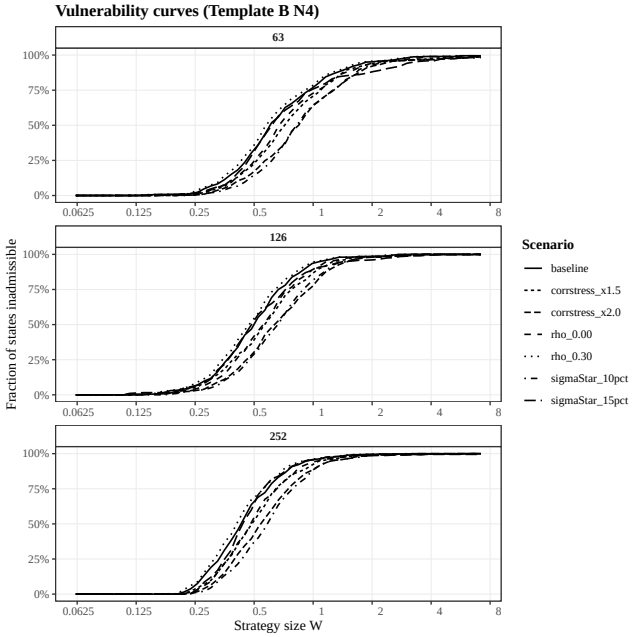

Figure 4 Vulnerability curves for Template B with $N = 4$.
Panels correspond to horizons $T \in \{63, 126, 252\}$; lines correspond to the scenarios implemented for $N = 4$.

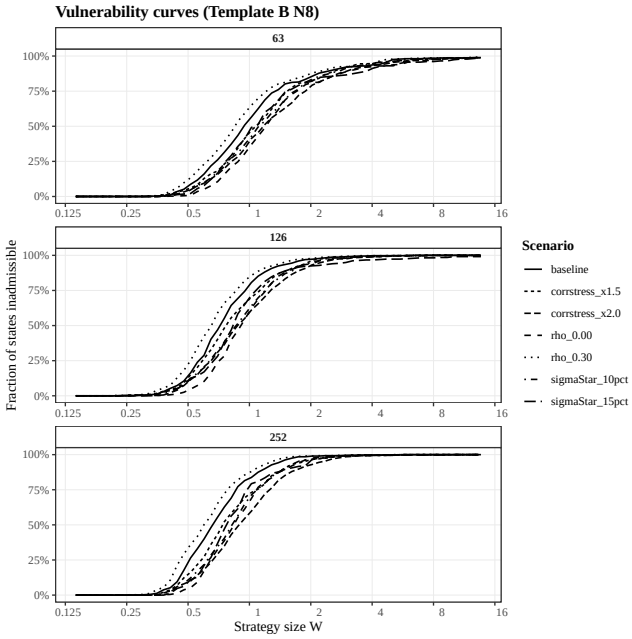

Figure 5 Vulnerability curves for Template B with $N = 8$.
Panels correspond to horizons $T \in \{63, 126, 252\}$; lines correspond to the scenarios implemented for $N = 8$.

H.2 Distribution of statewise capacities $W_{\max}(z)$

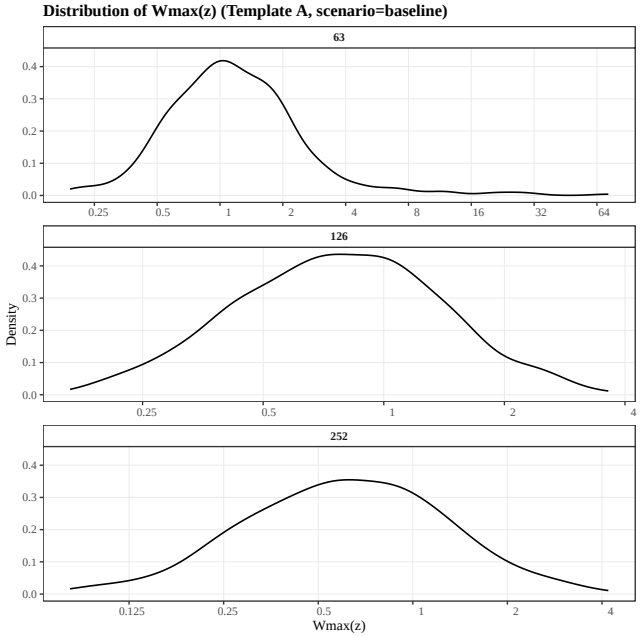

Figure 6 Kernel density estimate of $W_{\max}(z)$ across stress states for Template A (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

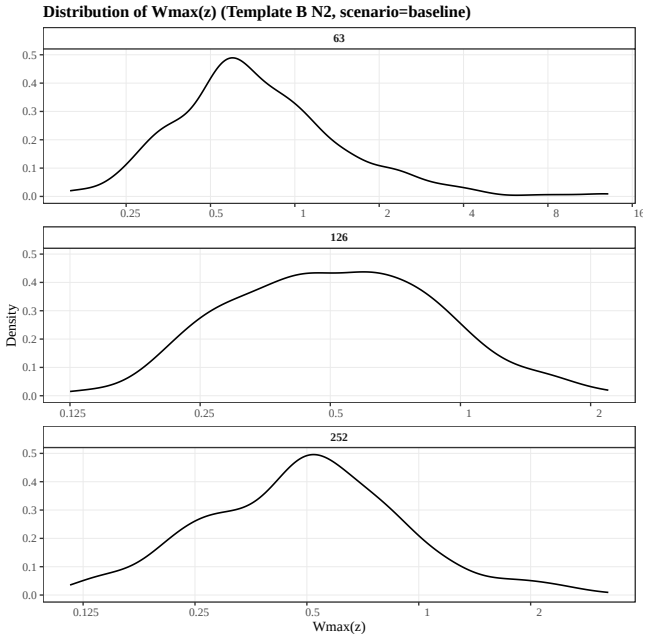

Figure 7 Kernel density estimate of $W_{\max}(z)$ across stress states for Template B with $N = 2$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

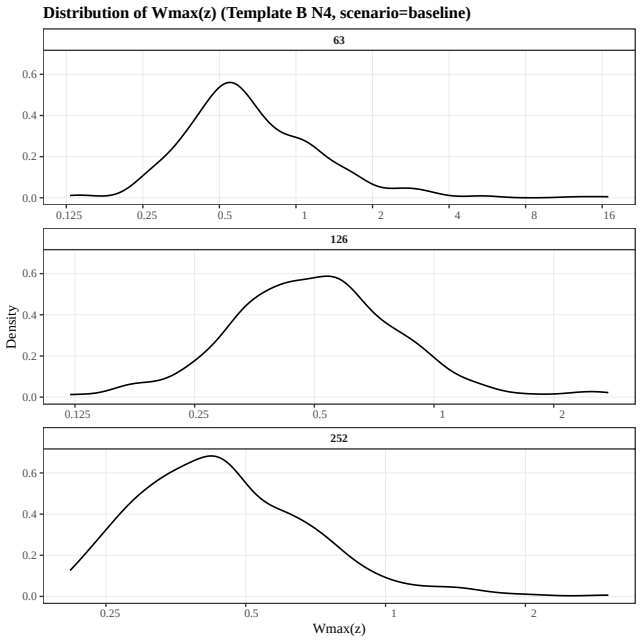

Figure 8 Kernel density estimate of $W_{\max}(z)$ across stress states for Template B with $N = 4$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

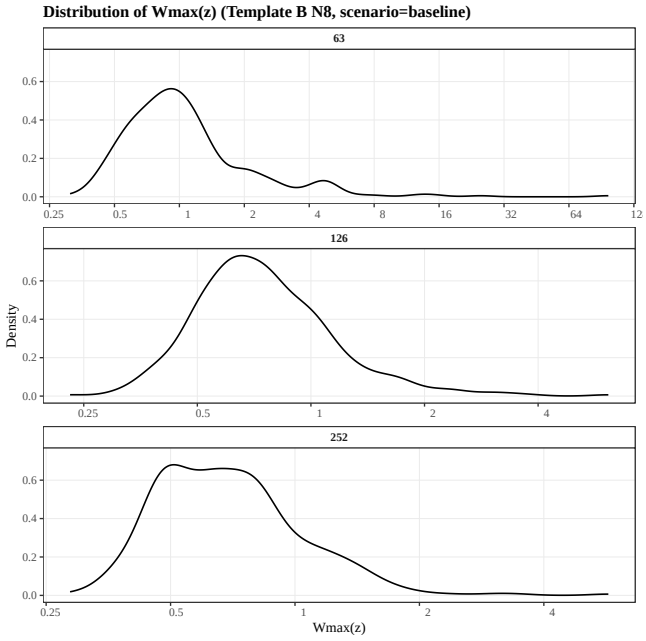

Figure 9 Kernel density estimate of $W_{\max}(z)$ across stress states for Template B with $N = 8$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

H.3 Local slope diagnostics

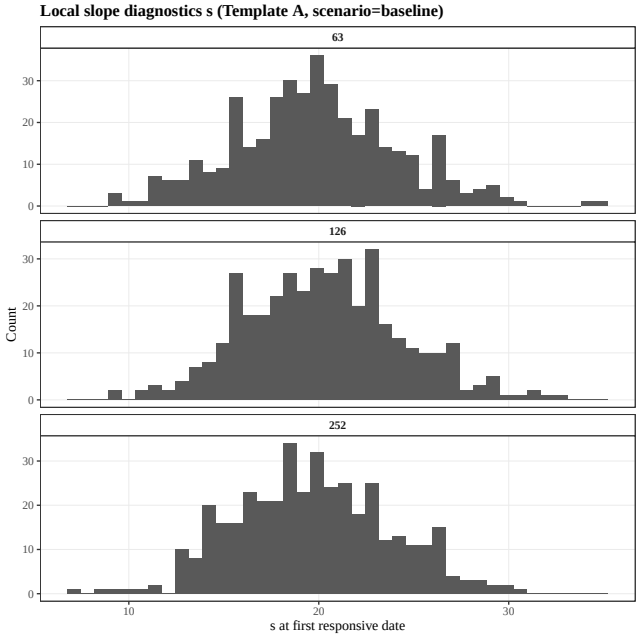

Figure 10 Histogram of the local slope diagnostic s at the first responsive date for Template A (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

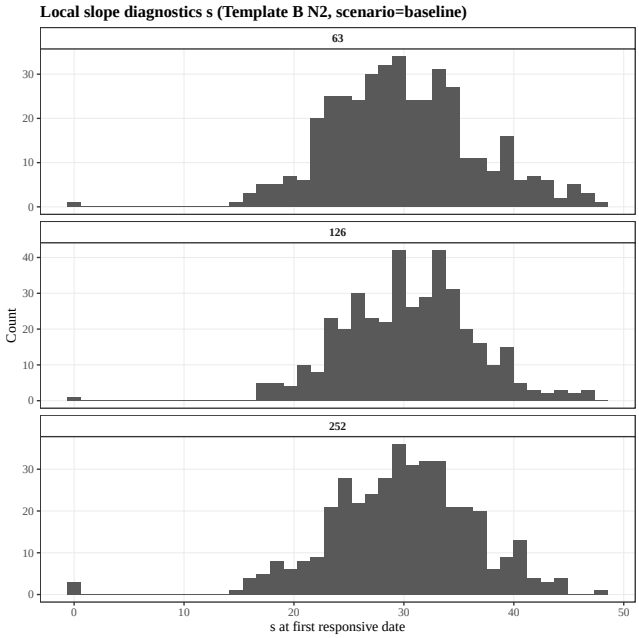

Figure 11 Histogram of the local slope diagnostic s at the first responsive date for Template B with $N = 2$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

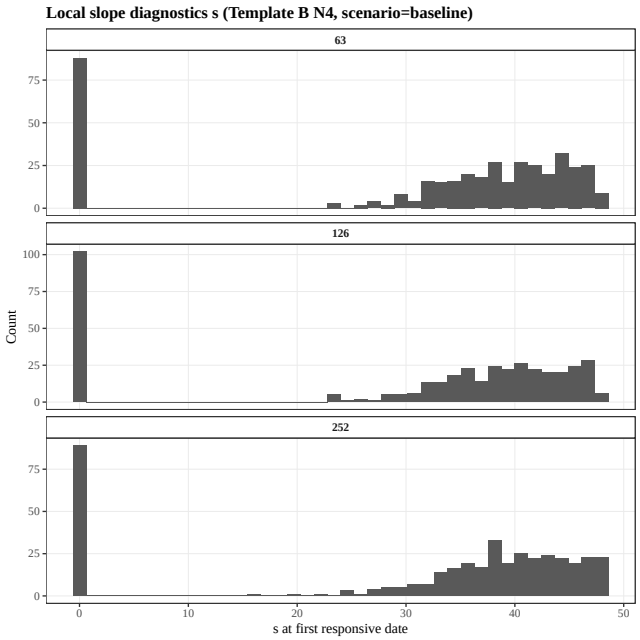

Figure 12 Histogram of the local slope diagnostic s at the first responsive date for Template B with $N = 4$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

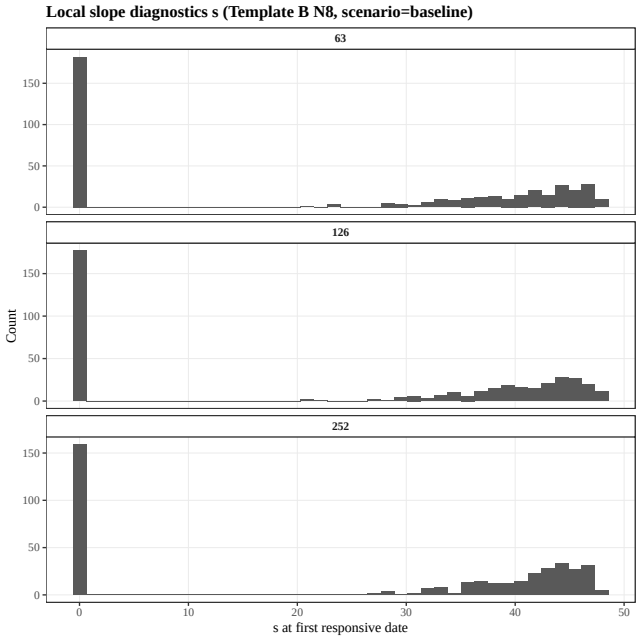

Figure 13 Histogram of the local slope diagnostic s at the first responsive date for Template B with $N = 8$ (baseline scenario), shown separately for $T \in \{63, 126, 252\}$.

H.4 Sensitivity summaries

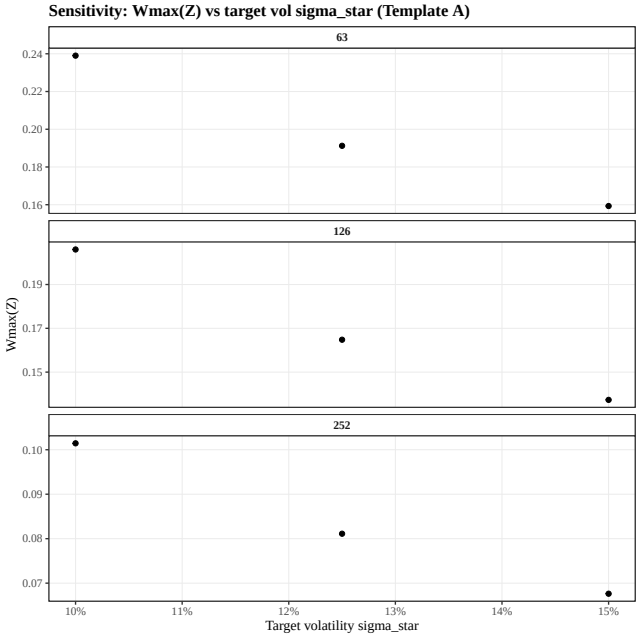

Figure 14 Sensitivity of $W_{\max}(\mathcal{Z})$ to the target volatility σ^* for Template A, shown across horizons $T \in \{63, 126, 252\}$.

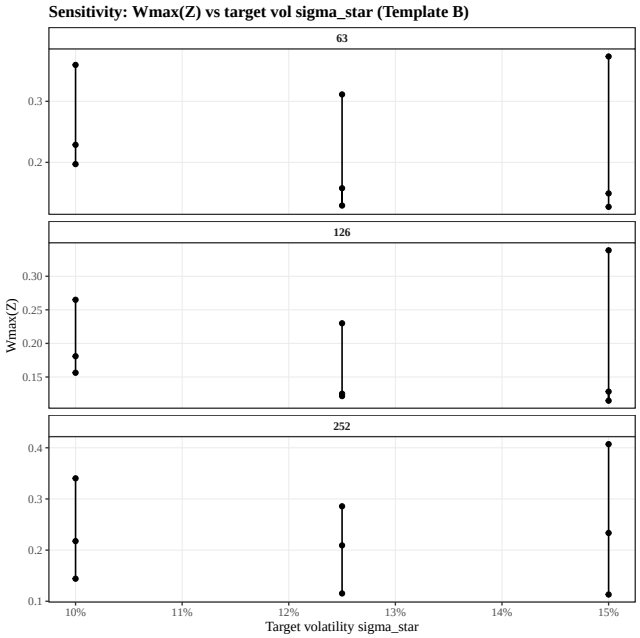

Figure 15 Sensitivity of $W_{\max}(\mathcal{Z})$ to the target volatility σ^* for Template B. The figure reports $W_{\max}(\mathcal{Z})$ for $N \in \{2, 4, 8\}$ and horizons $T \in \{63, 126, 252\}$.

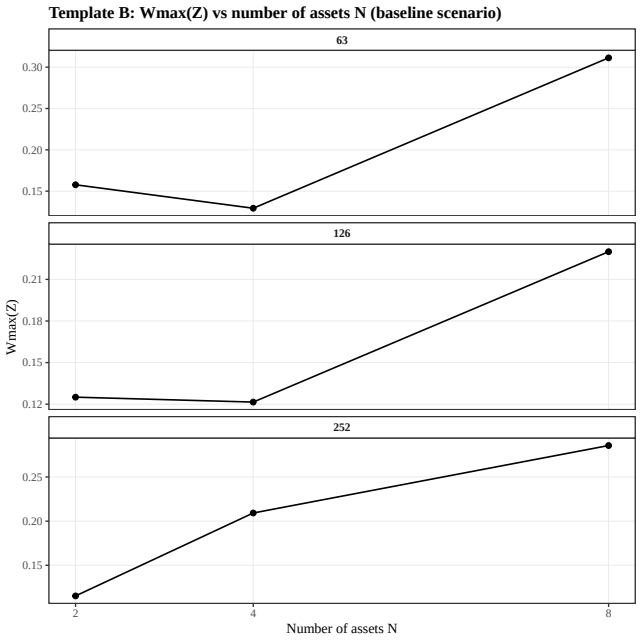

Figure 16 N-sensitivity for Template B (baseline scenario): $W_{\max}(\mathcal{Z})$ as a function of the number of hedge assets N, shown across horizons $T \in \{63, 126, 252\}$.

H.5 Cap-regime incidence at the first responsive date

Table 8 summarizes cap-regime incidence at the first responsive date for Template B in the baseline scenario (which has $\sigma^* = 12.5\%$) and in the two target-volatility scenarios $\sigma^* \in \{10\%, 15\%\}$. For each (N, σ^*, T), the table reports the number of stress states with regime `cap` at the first responsive date, the share of the $|\mathcal{Z}| = 400$ stress set, and the conditional failure probability at $W_{\text{med}} = \text{median}_{z \in \mathcal{Z}} W_{\max}(z)$.

Table 8 Incidence of cap-regime at the first responsive date for Template B (from the failure decomposition output). The "share of $\mathcal{Z}$" column reports $\#\{z : \text{cap}\}/|\mathcal{Z}|$ in percent.

N	Scenario	T	$\#\{z : \text{cap}\}$	share of $\mathcal{Z}$	$\Pr[\text{fail} \mid \text{cap}]$ at W_{med}
2	baseline	63	1	0.25%	100.0%
2	baseline	126	1	0.25%	100.0%
2	baseline	252	3	0.75%	100.0%
2	$\sigma^* = 10\%$	63	0	0.00%	–
2	$\sigma^* = 10\%$	126	0	0.00%	–
2	$\sigma^* = 10\%$	252	0	0.00%	–
2	$\sigma^* = 15\%$	63	31	7.75%	100.0%
2	$\sigma^* = 15\%$	126	19	4.75%	94.7%
2	$\sigma^* = 15\%$	252	28	7.00%	100.0%
4	baseline	63	88	22.00%	83.0%
4	baseline	126	102	25.50%	75.5%
4	baseline	252	89	22.25%	91.0%
4	$\sigma^* = 10\%$	63	2	0.50%	100.0%
4	$\sigma^* = 10\%$	126	6	1.50%	100.0%
4	$\sigma^* = 10\%$	252	2	0.50%	100.0%
4	$\sigma^* = 15\%$	63	249	62.25%	51.4%
4	$\sigma^* = 15\%$	126	248	62.00%	56.9%
4	$\sigma^* = 15\%$	252	246	61.50%	62.6%
8	baseline	63	182	45.50%	62.6%
8	baseline	126	177	44.25%	68.4%
8	baseline	252	159	39.75%	67.3%
8	$\sigma^* = 10\%$	63	38	9.50%	100.0%
8	$\sigma^* = 10\%$	126	28	7.00%	100.0%
8	$\sigma^* = 10\%$	252	21	5.25%	95.2%
8	$\sigma^* = 15\%$	63	314	78.50%	31.5%
8	$\sigma^* = 15\%$	126	313	78.25%	46.6%
8	$\sigma^* = 15\%$	252	319	79.75%	47.6%