

# Dynamic Arbitrage from Price-Based Risk Constraints

Francesco Nicolai\*  
BI Norwegian Business School

Simona Risteska<sup>†</sup>  
Warwick Business School

February 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 create dynamic arbitrage. Given only the rulebook and an impact calibration, we deliver a no-dynamic-arbitrage test and a capacity bound for the constrained sector. We solve the optimal manipulator strategy under single-asset and portfolio constraints; in the portfolio case the attack reduces to trading two portfolios. In volatility-managed portfolios, admissible scale is well below one day of ADV, and vulnerability rises sharply once linked notional reaches roughly one to two days of ADV. Manipulability is strongest in low-volatility states, through feedback between measured risk and rule-driven trading, not balance-sheet constraints.

**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, G28, G18, G11.

---

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

<sup>†</sup>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 next-period feasibility via margin requirements, leverage limits, collateral haircuts, or mandated exposure. Prices clear trades today and, through the constraint, mechanically determine tomorrow’s feasible positions. When the constrained sector is large enough to have market impact<sup>1</sup>, a strategic trader can profit by tilting the sampled statistic that sets future constraints and by trading against the induced forced flow.

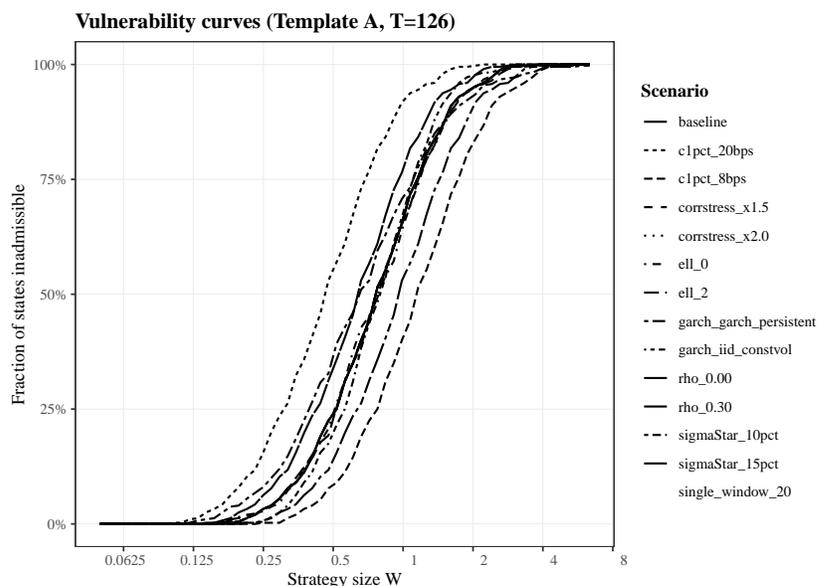
This paper shows that the classical [Huberman and Stanzl \(2004\)](#) restrictions on price impact, which rule out price manipulation by requiring that every admissible round trip has nonnegative expected execution cost, are not sufficient once price-based constraints create deterministic feedback. A market can satisfy these classical impact restrictions and still admit profitable round trips because price-based constraints generate predictable forced trading that can be manipulated. When constraints depend on a sampled risk statistic (e.g. realized variance, VaR, expected shortfall), an attacker can trade within the sampling window to tilt the measured statistic, induce mechanical rebalancing by the constrained sector, then unwind into the resulting predictable flow. The mechanism operates with arbitrarily small trades and arbitrarily small price moves, and it does not rely on making a slack constraint bind. We then derive additional restrictions that restore no-dynamic-arbitrage in the presence of price-based constraints (Theorem 1).

Beyond the refined no-arbitrage restrictions, the paper provides an ex ante screen for any proposed price-based rulebook or index methodology. The inputs are directly observable: a local market-impact estimate from execution-cost benchmarks and the disclosed mapping from sampled transaction prices into the statistic that sets next-period margin requirements, haircuts, leverage limits, or target exposures. We convert the implied rebalancing into an adjusted cost matrix  $\hat{H}$  and apply Theorem 1. The screen is transparent: the market is manipulation-free if and only if  $\hat{H}$  satisfies a restricted-eigenvalue condition. The same construction yields a conservative, liquidity-scaled capacity bound and vulnerability curves mapping sector scale into the share of states with profitable round trips.

Section 7 implements the screen for volatility-managed indices embedded in structured products and indexed annuities to show that actionable limits can be obtained from public information alone. Using only a market-impact estimate and the published volatility-control rule, we compute  $\hat{H}$ , deliver a pass/fail determination, and, when needed, produce a conservative upper bound on admissible scale and the associated vulnerability curve. Interpreted as an ex ante stress test, the diagnostic maps linked notional into the fraction of stress states in which a profitable dynamic manipulation strategy exists over horizon  $T$ . We implement the stress test for two product templates: a single-asset volatility-managed rule (Template A) and a portfolio rule (Template B). We vary liq-

---

<sup>1</sup>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](#)).



**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  $W_{\max}(z) = \sup\{W \geq 0 : \lambda_{\min}^{\mathcal{R}_T}(\hat{H}(z; W)) \geq 0\}$ , 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.

uidity, target volatility, implementation lags, volatility dynamics, cross-impact, and return correlations (Appendix F.4 and Appendix H). Figure 1 previews the diagnostic for a standard Template A volatility-control rule: for several plausible calibration and implementation choices, each line plots the fraction of simulated stress states in which the admissibility screen fails, as a function of sector size  $W$ . At  $T = 126$ , the conservative capacity bound  $W_{\max}(\mathcal{Z})$  ranges from 0.12 to 0.23 days of ADV (Template A: 0.165; Template B: 0.125 for  $N = 2$ , 0.122 for  $N = 4$ , and 0.230 for  $N = 8$ ). At the benchmark scale  $W = 1$  (one day of ADV), the admissibility screen fails in 66.8% of simulated stress states in Template A; in Template B it fails in 88.8% ( $N = 2$ ), 94.0% ( $N = 4$ ), and 78.8% ( $N = 8$ ).<sup>2</sup>

The paper also characterizes optimal attacks in Sections 5 and 6. The attacker trades for three reasons: to earn wedge profits by leaning against predictable rule-induced orders, to manage inventory so the strategy ends flat, and to tilt the sampled statistic by trading inside the sampling

<sup>2</sup>Volatility-managed products typically allocate between a risky ETF and a low-risk asset (cash, T-bills, or a bond index). The ADV and weight here refer to the risky ETF leg. If the risky allocation averages 0.80 of AUM and the risky ETF has \$5 billion ADV, then  $W_{\max}(\mathcal{Z}) = 0.165$  implies a risky notional of  $0.165 \times \$5 \text{ billion} = \$0.825 \text{ billion}$  and total linked AUM of  $\$0.825/0.80 = \$1.03 \text{ billion}$ . At  $W = 1$ , the risky notional is one ADV (\$5 billion) and total linked AUM is  $\$5/0.80 = \$6.25 \text{ billion}$ ; at that scale, the screen fails in 66.8% of stress states. When the screen fails, the linearized closed loop admits a positive-profit trigger-and-reverse direction. The attack need not be implemented at large volume: in our calibration, profitable round trips can be executed with daily turnover on the order of 1% of ADV. The attack is inventory-light, in the sense that the trigger can be implemented via a short-lived spike-and-revert that raises measured volatility while largely undoing the price-level change, so the attacker can unwind quickly and avoid warehousing a large directional position (Section 7.2.5).

window because the resulting statistic mechanically reshapes future requirements and future forced flow. Propositions 3–4 and 7–8 show that the manipulation motive is pinned down by two elasticities, the sensitivity of the sampled statistic to current prints and the slope of the requirement or exposure schedule in the binding region. For volatility-control designs, convexity of the statistic in returns implies that triggering can be inventory-light, and rolling windows mechanically time the onset and decay of induced tightening, creating a predictable entry and exit that can be harvested. This channel is economically distinct from standard models of corners, short squeezes, and standard price-manipulation stories: profits do not come from sustaining a distorted price level or restricting supply, but from shifting a disclosed statistic that deterministically reallocates demand across time.

In multiple assets, when the risk engine marks one portfolio direction while liquidation occurs along another, Theorem 2 shows that optimal exploitation lies in the two-dimensional span of these two portfolios. Cross-impact matters only through three scalars, so the diagnostic and the implied hedge instruments are identified directly from the marked and liquidation directions.

Timing matters. Theorem 3 shows that 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 prints inside the same sampling window (instantaneous updating). With a lag (and standard smoothing), the rule defines a stable closed-loop map from prices into trades. Without a lag, the rule can approach a fixed-point singularity: arbitrarily small print changes can imply arbitrarily large mandated rebalancing. Lags and smoothing therefore act as viability conditions for the rulebook, not only as devices to reduce procyclicality.

A central implication is that fragility need not coincide with high volatility, because the risk input itself is most manipulable in tranquil states. For target-volatility rules in their interior region, both links in the feedback loop are steepest when measured volatility is low: inverse-vol scaling makes exposure most sensitive to the volatility statistic, and any incremental contribution to realized variance produces the largest percentage change in the volatility estimate when the baseline is low. Consequently, calm markets are precisely when a small spike-and-revert can shift measured volatility enough to trigger material, 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, with requirements tightening because collateral values fall or funding conditions worsen. Here amplification operates through measurement: requirements are a deterministic function of a sampled risk statistic computed from recent prints, and that statistic can be shifted by trading even absent changes in fundamentals, collateral values, or intermediary capital.

## 1.1 Contributions to the Literature

Our first contribution is to extend the classical notion of market viability (no dynamic arbitrage) to settings in which transaction prices enter binding feasibility maps. In the benchmark of Huberman and Stanzl (2004) and the subsequent no-dynamic-arbitrage literature, a linear price-impact rule is

manipulation-free when every admissible round trip has nonnegative expected execution cost (Huberman and Stanzl, 2004; Gatheral, 2010; Schneider and Lillo, 2019). This restriction is sufficient because demand is exogenous from the trader’s perspective: trades move prices, but they do not change tomorrow’s constraints or tomorrow’s order flow. Price-based constraints add exactly the missing feedback. Sampled prints enter next-period requirements; when those requirements bind, they generate predictable rebalancing. Classical impact restrictions can therefore hold while dynamic arbitrage appears: a trader can profit from a trigger-and-reverse round trip that tilts the sampled statistic and then unwinds into the induced forced flow. Theorem 1 derives the sharp viability condition in this feedback economy. In the linearized closed loop, profitable round trips exist if and only if the adjusted cost matrix  $\hat{H}$  has a negative direction on the round-trip subspace. Equivalently, admissibility reduces to a single statistic, the restricted eigenvalue  $\lambda_{\min}^{\mathcal{R}_T}(\hat{H})$  (Theorem 4), computed mechanically from local rulebook derivatives and a local impact estimate. Appendix C shows that failure of the screen is not a linearization artifact: away from kinks, if  $\lambda_{\min}^{\mathcal{R}_T}(\hat{H}) < 0$  at a binding configuration, then the exact nonlinear system admits a genuinely profitable trigger-and-reverse round trip for sufficiently small admissible perturbations (Proposition 9). Together, these results identify an arbitrage channel that is invisible to classical no-manipulation screens and provide the additional restriction needed to restore market viability when feasibility is a deterministic function of transaction prices.

Our second contribution is to provide an implementable supervisory test. In the spirit of systemic-risk diagnostics that translate observable inputs into actionable oversight objects such as rankings, capital shortfalls, and vulnerability measures (Adrian and Brunnermeier, 2016; Acharya et al., 2012, 2017; Brownlees and Engle, 2017; Greenwood et al., 2015), we deliver a pass/fail test for price-based rulebooks. Given a disclosed rulebook and a local impact calibration, we construct the adjusted cost matrix  $\hat{H}$  that internalizes rule-induced flow and apply Theorem 1. The output is transparent: the rule is manipulation-free if and only if a restricted-eigenvalue condition holds, and the same construction yields a conservative capacity bound and a vulnerability curve mapping sector scale into the probability of exploitability.

Our third contribution is to characterize optimal strategic trading when forced flow is both predictable and manipulable, and to clarify how this channel differs from the existing predatory-trading canon. In predatory-trading and order-anticipation models, the imbalance is typically exogenous from the strategist’s perspective (Brunnermeier and Pedersen, 2005; Carlin et al., 2007; Attari et al., 2005; Rostek and Weretka, 2015; Sannikov and Skrzypacz, 2016; Fardeau, 2021; Lou et al., 2013): conditional on state variables, the trader optimally intermediates a given liquidation need or anticipated demand shock, and any feedback operates primarily through contemporaneous price effects on the victim’s balance-sheet constraint. Here the target is a disclosed, price-sampled risk engine. Transaction prices feed into a publicly specified statistic  $\Gamma_t$ , the rule maps  $\Gamma_t$  into next-period requirements  $M_{t+1} = g(\Gamma_t)$ , and when the requirement binds this mapping pins down induced rebalancing. Sections 5 and 6 solve the finite-horizon optimal round trip. Propositions 3, 4, 7, and 8 decompose optimal trading into harvesting induced flow, inventory management imposed by the round-trip constraint, and a distinct rule-manipulation motive pinned down by two observable local rule derivatives: the sensitivity of the sampled statistic to current prints and the local slope of the schedule  $g$ .

Mechanism-wise, the closest literature is contract-based manipulation and benchmark manipulation, because it shares the same causal chain: transaction prices determine a reference statistic; the statistic enters a publicly specified rule; and the rule mechanically induces payoffs or forced behavior. In [Kumar and Seppi \(1992\)](#), a trader takes a derivatives position and then trades the underlying to move the spot price used for cash settlement; in [Hillion and Suominen \(2004\)](#), the target is explicitly a reference statistic (the close) used for marking and evaluation; and [Dutt and Harris \(2005\)](#) and [Onur and Reiffen \(2018\)](#) study how settlement-rule design and position limits reshape the incentive to concentrate trading in the pricing window. [Zhang \(2022\)](#) provides an equilibrium synthesis in which contract structure disciplines manipulation incentives. Our contribution is different in object and in implication: we replace settlement or evaluation rules with a risk-engine constraint that deterministically generates future forced flow, and we solve the optimal trigger-and-reverse round trip in the resulting closed loop, alongside the supervision-ready admissibility diagnostic. Related feedback mechanisms arise when prices trigger leverage constraints or real actions ([Brunnermeier and Oehmke, 2014](#); [Goldstein and Guembel, 2008](#); [Gao et al., 2025](#)), but those channels do not feature a disclosed, statistic-driven schedule that mechanically maps sampled prints into required positions. Finally, when  $\Gamma$  is convex, as for realized variance and for standard tail-risk measures such as expected shortfall, triggering can be fundamentally inventory-light: the attacker can raise measured risk via short-lived spike-and-revert price oscillations that largely undo the price-level change, so the binding discipline on exploitation is execution costs and rule design, not the ability to warehouse a large directional position.

Our fourth contribution is methodological: we provide a transparent way to convert any disclosed, deterministic (or state-contingent) rulebook into the market’s effective execution-cost technology once the rulebook induces predictable flow. A large microstructure literature treats price impact as an equilibrium object pinned down by information frictions, intermediaries’ inventory and balance-sheet constraints, or limit-order-book dynamics and resilience.<sup>3</sup> Our approach is complementary and deliberately reduced form. We take as input any baseline local impact calibration, summarized by an operator  $\mathcal{I}$  that maps order flow into execution wedges, and we then internalize the rule-induced response using only public rule derivatives. Concretely, composing  $\mathcal{I}$  with the local loop operator  $\mathcal{K}$  implied by the rulebook yields an effective wedge operator  $\tilde{\mathcal{I}} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}$  and the corresponding adjusted symmetric cost matrix  $\hat{H} = \tilde{\mathcal{I}} + \tilde{\mathcal{I}}^\top$ .

Our fifth contribution is to make the question “how large can the sector be?” literal and supervision-ready. Given a disclosed rulebook and a baseline local impact calibration, we compute an admissible capacity bound and vulnerability curves: for each state  $z$ ,

$$W_{\max}(z) = \sup \left\{ W \geq 0 : \lambda_{\min}^{\mathcal{R}_T}(\hat{H}(z; W)) \geq 0 \right\},$$

and summarize by  $W_{\max}(\mathcal{Z}) = \inf_{z \in \mathcal{Z}} W_{\max}(z)$  and  $W \mapsto \Pr_z[W > W_{\max}(z)]$ . The object being bounded is not generic price pressure, but manipulability of the induced closed loop: beyond  $W_{\max}$  the admissibility screen fails and profitable round trips exist.

---

<sup>3</sup>Representative traditions include information-based equilibrium impact ([Kyle, 1985](#); [Glosten and Milgrom, 1985](#)), inventory-based liquidity provision ([Ho and Stoll, 1981](#); [Grossman and Miller, 1988](#)), and limit-order-book models emphasizing resilience and recovery ([Obizhaeva and Wang, 2013](#)).

Related “how large can X be?” literatures typically bound different objects. Inelastic-markets and demand-system work disciplines price effects of flows (Koijen and Yogo, 2019; Gabaix and Koijen, 2021; Haddad et al., 2025; Vayanos and Woolley, 2013; Ben-David et al., 2018; Da and Shive, 2018; Israeli et al., 2017); systemic-risk and fire-sale diagnostics map size and overlap into vulnerability or shortfalls (Greenwood et al., 2015; Acharya et al., 2012, 2017; Adrian and Brunnermeier, 2016; Brownlees and Engle, 2017); risk-constraint work studies amplification and procyclicality (Basak and Shapiro, 2001; Danielsson et al., 2004; Glasserman and Wu, 2018; BIS, 2010; ESRB, 2020); predictable-trading papers measure event-time price pressure (Lou et al., 2013; Bessembinder et al., 2016; Ni et al., 2005). We instead take the rulebook as primitive, internalize its induced flow into  $\widehat{H}$ , and deliver a computable no-manipulation capacity bound  $W_{\max}$ .

Finally, this paper is closely related to Nicolai (2026), which provides an impossibility result for price-based risk constraints and uses it to study mechanism design of update schedules. In binding high-amplification states, he shows that no price-based, risk-sensitive rule can simultaneously deliver liquidity continuity and round-trip manipulation-proofness, and he characterizes the implied design responses (pooling regions and bounded short-horizon pass-through). Our contribution is complementary: we take the institutional rulebook as given and, for any baseline local impact calibration, internalize rule-induced flow into an adjusted cost matrix  $\widehat{H}$  that delivers a computable pass/fail admissibility screen, together with conservative capacity bounds  $W_{\max}$  and vulnerability curves.

**Roadmap.** Section 2 sets up the baseline impact environment and the classical no-manipulation condition; Section 2.2 extends to multiple assets and cross-impact. Section 3 introduces deterministic price-based constraints and closes the feedback loop, and Section 4 derives the closed-loop screen (Theorem 1) and the timing result (Theorem 3). Sections 5 and 6 solve the optimal attack and establish the two-portfolio reduction (Theorem 2). Section 7 implements the rulebook-first 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*

$$x_{t+1} = x_t + u_t, \quad x_0 = 0. \quad (1)$$

*A strategy is admissible if it is predictable and uniformly bounded: there exists  $\bar{u} > 0$  such that  $|u_t| \leq \bar{u}$  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 almost surely,*

$$x_T = 0, \quad a.s. \quad (2)$$

*Equivalently,  $\sum_{t=0}^{T-1} u_t = 0$  a.s.*

### 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  $(\mathcal{F}_t)_{t \geq -1}$ .

**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 (Kyle, 1985; Almgren and Chriss, 2001; Gatheral, 2010).

**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*

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

The term  $\eta G_{t-s} u_s$  is the contribution of the time- $s$  trade to the time- $t$  execution price. The kernel  $(G_\ell)$  controls persistence: small  $G_\ell$  at large lags means impact decays quickly. A useful benchmark is purely temporary impact,  $G_0 = 1$  and  $G_\ell = 0$  for all  $\ell \geq 1$ , 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. The following lemma records the standard zero-expectation identity.

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

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

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

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

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

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

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.<sup>4</sup>

**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:

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

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.

### 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  $K \in \mathbb{R}^{T \times T}$  by

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

---

<sup>4</sup>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.

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. Define the symmetric matrix

$$H = \eta(K + K^\top). \quad (8)$$

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

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

For a stochastic round trip  $u$ , the realized trade vector satisfies  $u(\omega) \in \mathcal{R}_T$  a.s.

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

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

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,

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

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

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

which proves (10).

(ii) $\Rightarrow$ (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) $\Rightarrow$ (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).  $\square$

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  $t = 0, 1, \dots, T-1$ , 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

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

A vector round trip starts and ends flat,

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

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  $\|u_t\| \leq \bar{u}$  for all  $t$  and all  $\omega$ . If  $\sum_{t=0}^{T-1} u_t = 0$  a.s., then*

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

*Proof.* Let  $x_t = \sum_{s=0}^{t-1} u_s$  be inventory, so  $u_t = x_{t+1} - x_t$ . Uniform boundedness of  $u$  and  $\mathbb{E}[\|S_t\|] < \infty$  imply integrability of the terms below. Moreover,

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

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  $\mathbb{E}[S_t | \mathcal{F}_{t-1}] = S_{t-1}$ , hence for each  $t \geq 1$ ,

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

Summing over  $t$  gives the claim. □

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

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

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  $\ell \geq 1$ . 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  $U \in \mathbb{R}^{NT}$ ,

$$U = (u_0^\top, \dots, u_{T-1}^\top)^\top.$$

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

$$\mathbf{K}_{t,s} = G_{t-s} \mathbf{1}_{\{t \geq s\}}, \quad t, s \in \{0, 1, \dots, T-1\}. \quad (13)$$

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  $P_t(u) = S_t + (\mathbf{K}U)_t$ . Define the symmetric matrix

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

Collect deterministic vector round trips in

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

where the constraint is a vector equality in  $\mathbb{R}^N$ . For a stochastic round trip  $u$ , the realized stacked vector satisfies  $U(\omega) \in \mathbf{R}_T$  a.s.

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

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

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

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

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

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

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  $\mathbf{R}_T$ .  $\square$

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: temporary symmetric cross-impact

Consider purely temporary cross-impact: set  $G_0 = A$  and  $G_\ell = 0$  for all  $\ell \geq 1$ . Then  $P_t(u) = S_t + Au_t$ , so only the contemporaneous trade affects the contemporaneous execution price. For a round trip, Lemma 2 implies  $\mathbb{E}[\sum_t u_t^\top S_t] = 0$ , hence

$$\mathbb{E} \left[ \sum_{t=0}^{T-1} u_t^\top P_t(u) \right] = \mathbb{E} \left[ \sum_{t=0}^{T-1} u_t^\top Au_t \right]. \quad (16)$$

Since  $q^\top Aq = q^\top \frac{A+A^\top}{2} q$  for all  $q \in \mathbb{R}^N$ , only the symmetric part of  $A$  matters. Let  $A_s = (A + A^\top)/2$ . If  $A_s$  is symmetric and p.s.d., then  $u_t^\top Au_t = u_t^\top A_s u_t \geq 0$  for each  $t$ , so every admissible round trip has nonnegative expected execution cost. If instead there exists  $q \in \mathbb{R}^N$  with  $q^\top A_s q < 0$ , then the two-date round trip  $u_0 = q, u_1 = -q$  yields negative expected execution cost.

### 2.2.3 Example: two-asset arbitrage

A concrete setting 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

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

The diagonal terms  $\lambda_F, \lambda_C$  are own-impact. The off-diagonal terms are cross-impact: trading asset 1 shifts the execution price of asset 2 with slope  $\Lambda$ , while trading asset 2 shifts the execution price of asset 1 with slope  $\epsilon$ . Expected execution cost depends only on the symmetric part

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

A profitable round trip exists if and only if  $A_s$  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

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

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  $v$ ,

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

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

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 risk regulation that rescales exposure when measured risk rises (Basak and Shapiro, 2001; Danielsson et al., 2004). 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; 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. The common structure is: 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,

$$\Gamma_t = \Gamma(P_{t-m}, \dots, P_t), \quad (17)$$

and sets the next-date requirement

$$M_{t+1} = g(\Gamma_t), \quad (18)$$

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

$$|X_t(i)| M_t \leq E(i) \quad \text{for all constrained agents } i, \quad (19)$$

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  $X_t = X(M_t)$ . 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

$$q_t = u_t + v_t.$$

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*

$$P_t = S_t + \mathcal{A}q_t, \tag{20}$$

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  $z^\top A z \geq 0$  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 sequel is not impact-only manipulation. It is created by a distinct channel: transaction prices enter the rule (17)–(18), 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  $\Gamma_t$ , 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

$$\Pi(u) = - \sum_{t=0}^{T-1} u_t P_t. \tag{21}$$

For a round trip, Lemma 1 implies

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

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

$$\mathbb{E}[\Pi(u)] = -\eta \mathbb{E} \left[ \sum_{t=0}^{T-1} u_t^2 \right] - \eta \mathbb{E} \left[ \sum_{t=0}^{T-1} u_t v_t \right]. \quad (22)$$

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,

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

where the requirement  $M_{t+1} = g(\Gamma_t)$  is computed from past transaction prices  $\Gamma_t = \Gamma(P_{t-m}, \dots, P_t)$ . 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  $v_{t+1}, v_{t+2}, \dots$ .

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

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 trigger-and-reverse round trips. The constraint maps sampled transaction prices into next-period requirements; when requirements bind, marginal changes in requirements force predictable rebalancing by the constrained sector. A strategic trader can therefore trade within the sampling window to tilt the measured statistic, tighten tomorrow’s requirement, and then unwind against the resulting forced flow.

The 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 strategist (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 sharp screen for the linearized closed loop: margin-feedback manipulation is absent if and only if the symmetric closed-loop cost matrix is positive semidefinite on the round-trip subspace. The test depends on the impact map  $\mathcal{I}$ , the statistic Jacobian  $J$ , the local schedule slope  $s = g'(\Gamma_0)$ , the requirement-to-position sensitivity  $B$ , and the timing operators  $L$  and  $D$ . Although the condition is derived from a linearization, Appendix C shows that failure of the screen implies a truly profitable sufficiently small round trip in the exact nonlinear model (Proposition 9).

## 4.1 A linearized model

Fix a horizon  $T$  and a reference configuration at which the requirement binds. This subsection studies small trade-induced deviations around that reference, holding the unaffected price path  $S$  fixed. For any object  $Y$ , let  $\delta Y$  denote its deviation from the reference value induced by trading; in particular,  $\delta S = 0$ . Let  $u \in \mathbb{R}^T$  be the strategic trader's deviation in trades and let  $v \in \mathbb{R}^T$  be the resulting deviation in forced flow from the constrained sector. Total flow is

$$q = u + v.$$

Let  $S$  denote the unaffected price path and  $P$  the transaction price path. The impact block maps total flow into the wedge between transaction prices and unaffected prices. In deviations,

$$\delta P = \delta S + \mathcal{I}q, \tag{23}$$

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

$$\delta \Delta P = \delta P - \delta S.$$

Because deviations are taken at fixed  $S$ , we have  $\delta S = 0$  and hence  $\delta \Delta P = \delta P$ . Substituting into (23) yields the within-horizon wedge-price relation

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

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

$$\delta \Gamma = J \delta \Delta P, \tag{24}$$

where  $J$  is the Jacobian of  $\Gamma(P)$  evaluated at the reference configuration, holding  $S$  fixed. To keep requirements aligned with the trading dates  $t = 0, \dots, T-1$ , 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  $t = 1, \dots, T - 1$ ,

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

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 on paths, defined by  $(Lx)_0 = 0$  and  $(Lx)_t = x_{t-1}$  for  $t = 1, \dots, T - 1$ .

#### 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

$$v = -B D \delta M, \tag{25}$$

where  $D \in \mathbb{R}^{T \times T}$  is the first-difference operator on paths, defined by  $(Dx)_0 = x_0$  and  $(Dx)_t = x_t - x_{t-1}$  for  $t = 1, \dots, T - 1$ . Since  $(\delta M)_0 = 0$  under the within-horizon rule update, rule-induced forced flow can begin no earlier than  $t = 1$ .

#### 4.1.2 Closing the loop

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

$$v = -B s D L J \delta \Delta P. \tag{26}$$

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

$$q = u - B s D L J \mathcal{I}q. \tag{27}$$

Define the loop operator

$$\mathcal{K} = B s D L J \mathcal{I}. \tag{28}$$

Then  $(\text{Id} + \mathcal{K})q = u$ . We stack time-indexed paths in chronological order as column vectors  $x = (x_0, \dots, x_{T-1})^\top$ . We impose the restriction that current wedge prices depend only on current and past flows, and that  $\Gamma_t$  depends only on prices observed up to date  $t$ . In matrix form, this means

$\mathcal{I}_{t,s} = 0$  and  $J_{t,s} = 0$  for all  $s > t$ , so  $J\mathcal{I}$  is lower triangular. Under lagged updating  $M_{t+1} = g(\Gamma_t)$ , the operator  $DL$  is strictly lower triangular, hence  $\mathcal{K} = BsDLJ\mathcal{I}$  is strictly lower triangular on a finite horizon. Therefore  $\text{Id} + \mathcal{K}$  is automatically invertible and the fixed point selects a unique total-flow path for each strategic deviation  $u$ :

$$q = (\text{Id} + \mathcal{K})^{-1}u, \quad \delta\Delta P = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}u. \quad (29)$$

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.

## 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 closed loop, (29) implies

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

Define the effective impact operator

$$\tilde{\mathcal{I}} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}.$$

This operator 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

$$\mathbb{E} \left[ \sum_{t=0}^{T-1} u_t P_t \right] = \mathbb{E} \left[ u^\top (P - S) \right] = \mathbb{E} \left[ u^\top \tilde{\mathcal{I}} u \right]. \quad (30)$$

Only the symmetric part matters on a quadratic form, so define the augmented symmetric cost matrix

$$\hat{H} = \tilde{\mathcal{I}} + \tilde{\mathcal{I}}^\top = \mathcal{I}(\text{Id} + \mathcal{K})^{-1} + (\text{Id} + \mathcal{K})^{-\top} \mathcal{I}^\top. \quad (31)$$

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

$$\mathbb{E}[\Pi(u)] = -\mathbb{E} \left[ u^\top \tilde{\mathcal{I}} u \right] = -\frac{1}{2} \mathbb{E} \left[ u^\top \hat{H} u \right]. \quad (32)$$

A profitable round trip exists exactly when  $\hat{H}$  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 (23)–(29) around a binding state, and assume  $\text{Id} + \mathcal{K}$  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

$$v = -Bs (D \otimes \text{Id}_N)(\text{Id}_T \otimes b) L J \delta\Delta P.$$

Define the loop operator

$$\mathcal{K} = Bs (D \otimes \text{Id}_N)(\text{Id}_T \otimes b) L J \mathcal{I} \in \mathbb{R}^{NT \times NT},$$

and the symmetric closed-loop cost matrix

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

Then the combined system is margin-feedback manipulation free in the linearized class if and only if  $\widehat{\mathbf{H}}$  is positive semidefinite on the multi-asset round-trip subspace  $\mathbf{R}_T$ .

*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 (29), we have  $P - S = \widetilde{\mathcal{I}}u$ , hence (32) gives

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

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.  $\square$

Appendix B, Theorem 4 rewrites the condition as a restricted eigenvalue test on the round-trip subspace, which is convenient for computation. The economic content is that the relevant object is the effective wedge map  $\widetilde{\mathcal{I}} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}$ , rather than  $\mathcal{I}$  alone: profitability requires the rule-induced response to generate sufficient order-flow predictability, with the right sign and timing, to overcome direct execution costs.

This mechanism 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 that is generated by the counterparty's balance-sheet stress, funding constraints, or liquidation needs; conditional on the relevant state variables, the liquidation

pressure is taken as given from the predator's perspective, and any feedback from the predator's own trades works primarily through moving the price level and thereby worsening the counterparty's financing terms. Here, the imbalance is produced by a disclosed and deterministic mapping from sampled transaction prices into next-period requirements. When requirements bind, the induced rebalancing is mechanical and locally pinned down by observable derivatives of the rule. The strategic trader can therefore shift the magnitude and timing of future forced flow by trading inside the sampling window to move the statistic the rule measures, which makes the profit opportunity a 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 illustrate that an impact block 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,

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

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  $v_1 = k u_0$  with  $k > 0$ . Then  $q_0 = u_0$  and  $q_1 = u_1 + k u_0$ . For the round trip  $u = (-q, q)$ ,

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

so the round trip is profitable whenever  $k > 2$ . This reduced-form rule is consistent with the linearized primitives in Section 4. In deviations,

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

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 the induced loop operator  $\mathcal{K} = BsDLJI$  is

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

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

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

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

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

which matches the positive-profit region above. The example shows why classical no-manipulation, which constrains  $\mathcal{I}$  alone, is not sufficient once deterministic feedback is present. The relevant object is the effective map  $\tilde{\mathcal{I}} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}$ , and the augmented condition tests the symmetric part  $\hat{H}$  on the round-trip subspace.

## 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 for temporary-impact models (Huberman and Stanzl, 2004; Gatheral, 2010; Schneider and Lillo, 2019). This section takes the rule and impact primitives as given and solves the strategic trader's finite-horizon optimal round trip in the single-asset environment. The solution separates two motives. The attacker earns wedge profits by supplying liquidity against predictable constraint-driven orders, as in models of order anticipation and predatory trading (Brunermeier and Pedersen, 2005). At the same time, the attacker can trade inside the sampling window to move the rule's input and thereby reshape future forced flow. The dynamic program makes these incentives explicit through shadow values for inventory and for shifting the current sampled print.

### 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  $t = 0, 1, \dots, T - 1$ . 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. Appendix D.1 states the corresponding stochastic Bellman recursion, which augments the state by  $S_t$  and reinstates conditional expectations. At the start of date  $t$ , before choosing  $u_t$ , the state is

$$s_t = (x_t, y_t, P_{t-1}, M_t, X_t, X_{t-1}). \quad (33)$$

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  $(y_t, P_t)$ :

$$y_{t+1} = F(y_t, P_t), \quad \Gamma_t = \Gamma(y_t, P_{t-1}, P_t), \quad M_{t+1} = g(\Gamma_t), \quad (34)$$

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  $t$ :

$$v_t = X_t - X_{t-1}. \quad (35)$$

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

$$\mathcal{U}_t(x_t) = \{u \in \mathbb{R} : |u| \leq \bar{u}, |x_t + u| \leq (T - 1 - t)\bar{u}\}. \quad (36)$$

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

$$P_t = S_t + \eta q_t = S_t + \eta(u_t + v_t). \quad (37)$$

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 (Appendix D.1 makes this explicit in the stochastic formulation), 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 in wedge units as

$$\pi_t(s_t, u_t) = -u_t(P_t - S_t) = -\eta u_t(u_t + v_t), \quad (38)$$

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 (34) and can therefore move future requirements and future forced flow. Given  $s_t$  and  $u_t$ , the next state is

$$s_{t+1} = \mathcal{T}(s_t, u_t) = (x_t + u_t, F(y_t, P_t), P_t, g(\Gamma_t), W/g(\Gamma_t), X_t), \quad (39)$$

where  $P_t$  is given by (37) and  $\Gamma_t$  by (34). Define the value function

$$V_t(s) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \sum_{\tau=t}^{T-1} \pi_\tau(s_\tau, u_\tau), \quad (40)$$

subject to  $u_\tau \in \mathcal{U}_\tau(x_\tau)$ , the transition (39), and the terminal constraint  $x_T = 0$ . The associated Bellman

recursion is

$$V_T(s) = 0 \text{ if } x_T = 0, \quad V_T(s) = -\infty \text{ otherwise,} \quad (41)$$

$$V_t(s) = \max_{u \in \mathcal{U}_t(x)} \{ \pi_t(s, u) + V_{t+1}(\mathcal{T}(s, u)) \}, \quad t = T - 1, \dots, 0. \quad (42)$$

## 5.2 Trade decomposition: harvesting, inventory management, and triggering

The dynamic program highlights two distinct roles of a trade. First, it determines the current wedge payoff through (38). Second, by moving  $P_t$ , it affects the rule state 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 (38), the one-period wedge payoff satisfies*

$$\pi_t(s_t, u) = -\eta \left( u + \frac{v_t}{2} \right)^2 + \frac{\eta}{4} v_t^2. \quad (43)$$

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

$$u_t^{\text{myopic}} = -\frac{v_t}{2}, \quad \max_{u \in \mathbb{R}} \pi_t(s_t, u) = \frac{\eta}{4} v_t^2. \quad (44)$$

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

*Proof.* Completing the square in (38) gives (43). The square term is minimized at  $u = -v_t/2$ , yielding the unique maximizer and the maximal value in (44). If  $v_t = 0$ , then  $\pi_t(s_t, u) = -\eta u^2 \leq 0$ , with equality only at  $u = 0$ .  $\square$

Lemma 3 provides the harvesting benchmark. When  $v_t$  is predetermined, wedge gains arise only from trading against the forced order, and the myopic optimum absorbs exactly half of it. The quadratic form (43) makes the interpretation explicit: the attacker is compensated for providing intermediacy against a predictable imbalance, while paying the usual temporary-impact cost on own orders. Optimal attacks depart from this 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  $u_t^*$  in (42) is interior. Assume that  $V_{t+1}$  is differentiable at  $s_{t+1}^* = \mathcal{T}(s_t, u_t^*)$  and that  $\mathcal{T}(s_t, u)$  is differentiable at  $u_t^*$ . Define*

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

Then  $u_t^*$  satisfies the Euler equation

$$\eta(2u_t^* + v_t) = \mu_{t+1} + \eta\nu_{t+1}, \quad (45)$$

or equivalently

$$u_t^* = -\frac{v_t}{2} + \frac{1}{2\eta}\mu_{t+1} + \frac{1}{2}\nu_{t+1}. \quad (46)$$

Moreover, the resulting current-period payoff satisfies

$$\pi_t(s_t, u_t^*) = \frac{\eta}{4}v_t^2 - \eta\left(u_t^* + \frac{v_t}{2}\right)^2 = \frac{\eta}{4}v_t^2 - \frac{1}{4\eta}(\mu_{t+1} + \eta\nu_{t+1})^2. \quad (47)$$

*Proof.* See Appendix D.2. □

The decomposition (46) isolates three forces. The term  $-v_t/2$  is the myopic harvesting benchmark from Lemma 3. The inventory shadow value  $\mu_{t+1}$  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  $\nu_{t+1}$  is specific to the feedback environment because  $P_t$  is an input to the requirement update; when a marginal change in the current print shifts future requirements and thus future forced flow,  $\nu_{t+1}$  is the continuation-value benefit of moving  $P_t$ .

Equation (47) expresses the implied tradeoff in wedge units. The term  $\frac{\eta}{4}v_t^2$  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  $u_t^*$  that is deliberately off the harvesting benchmark, often making  $\pi_t(s_t, u_t^*)$  negative on impact, because the continuation-value gains embodied in  $\mu_{t+1}$  and/or  $\nu_{t+1}$  dominate.

**Corollary 1** (Trigger dates versus harvest dates). *Fix a date  $t$  with  $v_t = 0$ . Lemma 3 implies  $\pi_t(s_t, u) \leq 0$  for all  $u$ . Assume  $0 \in \mathcal{U}_t(x_t)$  (equivalently,  $|x_t| \leq (T - 1 - t)\bar{u}$ ), so that  $u = 0$  is feasible. Then any optimal action with  $u_t^* \neq 0$  must be chosen for its effect on continuation value in (42), since it cannot raise the current-period wedge payoff.*

*If, in addition, the conditions of Proposition 3 hold at  $(s_t, u_t^*)$ , then the decomposition (46) reduces to*

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

*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,  $u_t^* = -v_t/2$ .*

*Proof.* See Appendix D.3. □

Dates with nonzero  $v_t$  are natural harvest dates because wedge 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 wedge 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 print 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  $(s_t, u_t^*)$ . Define the combined continuation wedge

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

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

*Proof.* See Appendix D.4. □

Corollary 2 isolates the knife-edge for procyclicity. Trading with the forced order requires the continuation wedge  $\chi_{t+1}$  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  $\text{sign}(\chi_{t+1}) = \text{sign}(v_t)$ . In that region the attacker necessarily accepts a negative current wedge 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  $\nu_{t+1}$  as the continuation-value gain from moving the current transaction price through the rule update. The next result expresses  $\nu_{t+1}$  directly in terms of the rule primitives  $(F, \Gamma, g)$  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  $\nu_{t+1}$ ). Fix  $t \in \{0, \dots, T - 1\}$  and consider any state-action pair  $(s_t, u_t)$ . Let  $P_t$  be given by (37) 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  $\Gamma$  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

$$\Lambda_{t+1} = \partial_M V_{t+1}(s_{t+1}) - B_{t+1} \partial_X V_{t+1}(s_{t+1}), \quad (48)$$

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

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + (\nabla_y V_{t+1}(s_{t+1})) \cdot \partial_P F(y_t, P_t) + \Lambda_{t+1} g'(\Gamma_t) \partial_P \Gamma(y_t, P_{t-1}, P_t). \quad (49)$$

*Proof.* See Appendix D.5. □

The first two terms in (49) 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  $\Lambda_{t+1}$  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  $\text{sign}(\Lambda_{t+1}g'(\Gamma_t)\partial_P \Gamma)$ .

**Corollary 3.** *Under the conditions of Proposition 4, the margin-update component of  $\nu_{t+1}$  in (49) is*

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

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

$$\text{sign}(\nu_{t+1}^M) = \text{sign}(\Lambda_{t+1}g'(\Gamma_t)\partial_P \Gamma(y_t, P_{t-1}, P_t)).$$

*Proof.* See Appendix D.6. □

Manipulation through the margin update requires both links in the feedback chain to be locally operative: the current print 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 locally 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  $(v_\tau)$ . When both links are active, the induced sensitivity of next-period constrained demand to the current print is

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

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

## 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.

**Lemma 4** (Terminal unwind). *At the terminal trade date  $t = T - 1$ , the admissible action set (36) satisfies*

$$\mathcal{U}_{T-1}(x_{T-1}) = \begin{cases} \{-x_{T-1}\}, & \text{if } |x_{T-1}| \leq \bar{u}, \\ \emptyset, & \text{otherwise.} \end{cases}$$

Consequently, on any feasible terminal state (so that  $\mathcal{U}_{T-1}(x_{T-1}) \neq \emptyset$ ), the unique admissible trade is  $u_{T-1}^* = -x_{T-1}$  and

$$V_{T-1}(s_{T-1}) = \pi_{T-1}(s_{T-1}, -x_{T-1}).$$

*Proof.* At  $t = T - 1$ , the feasibility condition in (36) becomes  $|x_{T-1} + u| \leq 0$ , hence  $u = -x_{T-1}$ . This trade is admissible if and only if  $|u| \leq \bar{u}$ , equivalently  $|x_{T-1}| \leq \bar{u}$ . When the action set is nonempty, the Bellman recursion (42) together with the terminal condition (41) implies  $V_{T-1}(s_{T-1}) = \pi_{T-1}(s_{T-1}, -x_{T-1})$ .  $\square$

**Corollary 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  $v_{t+1} = 0$  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.  $\square$

Corollary 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  $\tau = t, \dots, T - 1$ ,

$$u_\tau^* = -\frac{1}{2}(v_\tau - \bar{v}_t) - \frac{x_t}{n}. \quad (50)$$

Moreover, the associated maximal continuation profit is

$$V_t^{\text{harv}} = \frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2 + \eta x_t \bar{v}_t - \frac{\eta}{n} x_t^2. \quad (51)$$

*Proof.* See Appendix D.8. □

Proposition 5 shows that with a pinned-down forced-flow path, the attacker 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  $v_\tau$  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 5** (Only dispersion generates timing rents). *In the setting of Proposition 5,*

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

The only component of  $V_t^{\text{harv}}$  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

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

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

*Proof.* See Appendix D.9. □

**Corollary 6** (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*

$$u_\tau^* = \frac{v}{2n} < 0 \quad \text{for } \tau = t, \dots, T - 2, \quad u_{T-1}^* = -\frac{n-1}{2n} v > 0.$$

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. □

Corollaries 5–6 isolate a point that matters for feedback. Holding the forced-flow path fixed, continuation value increases with its temporal dispersion rather than with its average level. Design features that bunch rebalancing into a narrow window therefore raise dispersion and increase the rents available from subsequent intermediation. Corollary 6 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  $(v_\tau)$  is induced, the remaining problem is exactly the harvesting-and-unwind continuation characterized above.

Finally, per-date speed limits matter for implementation. When  $|u_\tau| \leq \bar{u}$  binds, the unconstrained intermediation rule (50) cannot be implemented date by date. The optimum then clips the unconstrained target, saturating at  $\pm\bar{u}$  on the most extreme dates and adjusting 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  $|u_\tau| \leq \bar{u}$ :*

$$\max_{(u_\tau)_{\tau=t}^{T-1}} -\eta \sum_{\tau=t}^{T-1} (u_\tau^2 + u_\tau v_\tau) \quad \text{s.t.} \quad \sum_{\tau=t}^{T-1} u_\tau = -x_t, \quad |u_\tau| \leq \bar{u} \quad \forall \tau.$$

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

$$u_\tau^* = \text{clip}_{[-\bar{u}, \bar{u}]} \left( -\frac{v_\tau}{2} + c_t \right), \quad (52)$$

where  $\text{clip}_{[-\bar{u}, \bar{u}]}(z) = \min\{\bar{u}, \max\{-\bar{u}, z\}\}$  and  $c_t$  can be chosen so that  $\sum_{\tau=t}^{T-1} u_\tau^* = -x_t$ . If at least one bound is slack at the optimizer (so  $|u_\tau^*| < \bar{u}$  for some  $\tau$ ), then  $c_t$  is unique. If the bounds are slack for all  $\tau$ , then (52) reduces to (50).

*Proof.* See Appendix D.11. □

The single-asset dynamic program yields a clean mapping from the optimal round trip to economically interpretable motives and corresponding design levers. Fix a state  $s_t$ . If forced flow  $v_t$  were pinned down, Lemma 3 implies that wedge profits arise only from supplying liquidity against it; the myopic benchmark is  $u_t = -v_t/2$  and the one-period harvesting rent is  $\eta v_t^2/4$ . Proposition 3 shows how an optimal policy departs from this benchmark once the terminal round-trip constraint and rule feedback are internalized. At an interior optimum,

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

so the trade splits into harvesting, an inventory term  $\mu_{t+1} = \partial_x V_{t+1}$  that allocates the remaining unwind across dates under quadratic impact, and a rule-manipulation term  $\nu_{t+1} = \frac{d}{dP_t} V_{t+1}$  that values moving the current print because it enters future requirements and thus future forced trading. Corollary 2 sharpens the implication: trading with forced flow, and accepting a negative current

wedge payoff, occurs only when the combined continuation wedge  $\chi_{t+1} = \mu_{t+1} + \eta\nu_{t+1}$  is large enough to overturn harvesting, so the date- $t$  trade is chosen primarily to improve continuation value.

Proposition 4 and Corollary 3 identify which features of the rule generate the manipulation motive. The margin-update component of  $\nu_{t+1}$  equals  $\Lambda_{t+1}g'(\Gamma_t)\partial_P\Gamma$ , where  $\partial_P\Gamma$  is the local sensitivity of the statistic to the current print,  $g'(\Gamma_t)$  is the local slope of the posted schedule, and  $\Lambda_{t+1}$  is the shadow value of tightening next-period feasibility along  $X_{t+1} = W/M_{t+1}$ . This channel is operative only when  $g'(\Gamma_t)\partial_P\Gamma \neq 0$ ; if either link is locally flat, the rule cannot be moved at the margin and the attacker faces a fixed within-horizon forced-flow path, reverting to harvesting and inventory management.

Endgame structure further disciplines behavior. Lemma 4 pins down the final trade, and Corollary 4 shows that with an implementation lag  $\ell$  there is a last trigger date  $t_{\text{last}} = T - 2 - \ell$ , after which  $\nu_{t+1} = 0$  and the continuation problem becomes pure intermediation and unwind against predetermined forced flow. Propositions 5 and 6 solve that continuation problem. With slack bounds, the optimal policy is linear and the value splits into a dispersion term driven by temporal variation in ( $v_\tau$ ) and an inventory term driven by ( $x_t, \bar{v}_t$ ). With binding speed limits, the unconstrained target is clipped pointwise and the remaining slack is used to satisfy the terminal inventory constraint. Taken together, the section delivers a tractable description of the attacker's toolkit and the rule features that govern it, namely the local derivatives  $g'(\Gamma_t)$  and  $\partial_P\Gamma$  and the implementation primitives (lags and speed limits) that determine when, and how strongly, a current print can reshape within-horizon forced trading. Section 6 extends the same logic to multiple assets with cross-impact and a common statistic.

## 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 wedge 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 engine computes a scalar statistic from the joint price path, 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

$$s_t = (x_t, y_t, P_{t-1}, M_t, X_t, X_{t-1}), \quad (53)$$

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

$$y_{t+1} = F(y_t, P_t), \quad \Gamma_t = \Gamma(y_t, P_{t-1}, P_t), \quad M_{t+1} = g(\Gamma_t). \quad (54)$$

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

$$X_t = \frac{W}{M_t} b. \quad (55)$$

The associated local sensitivity of holdings to the requirement is

$$-\frac{\partial X_t}{\partial M_t} = \frac{W}{M_t^2} b. \quad (56)$$

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

$$v_t = X_t - X_{t-1}. \quad (57)$$

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

$$\|u_t\|_1 \leq \bar{u}. \quad (58)$$

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

$$\mathcal{U}_t(x_t) = \{u \in \mathbb{R}^N : \|u\|_1 \leq \bar{u}, \|x_t + u\|_1 \leq (T - 1 - t)\bar{u}\}. \quad (59)$$

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

$$P_t = S_t + Aq_t = S_t + A(u_t + v_t). \quad (60)$$

On a round trip, the unaffected-price component has zero expected contribution under Assumption 3 (Lemma 2), so expected profits are pinned down by the execution wedge. Using (60), the date- $t$  wedge cashflow is

$$\pi_t(s_t, u_t) = -u_t^\top (P_t - S_t) = -u_t^\top A(u_t + v_t). \quad (61)$$

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 geometry encoded in  $A$ . Define the transition map  $s_{t+1} = \mathcal{T}(s_t, u_t)$

by

$$\mathcal{T}(s_t, u_t) = \left( x_t + u_t, F(y_t, P_t), P_t, g(\Gamma_t), X_{t+1}, X_t \right), \quad (62)$$

where  $P_t$  is given by (60),  $\Gamma_t = \Gamma(y_t, P_{t-1}, P_t)$ , and

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

The value function is

$$V_t(s) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \sum_{\tau=t}^{T-1} \pi_\tau(s_\tau, u_\tau), \quad (63)$$

subject to  $u_\tau \in \mathcal{U}_\tau(x_\tau)$ , the transition map (62), and  $x_T = 0$ . The associated Bellman recursion is

$$V_T(s) = 0 \text{ if } x_T = 0, \quad V_T(s) = -\infty \text{ otherwise}, \quad (64)$$

$$V_t(s) = \max_{u \in \mathcal{U}_t(x)} \{ \pi_t(s, u) + V_{t+1}(\mathcal{T}(s, u)) \}, \quad t = T-1, \dots, 0. \quad (65)$$

Relative to (42), 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: wedge 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 (61), the one-period payoff satisfies*

$$\pi_t(s_t, u) = -\left(u + \frac{v_t}{2}\right)^\top A \left(u + \frac{v_t}{2}\right) + \frac{1}{4} v_t^\top A v_t. \quad (66)$$

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

$$u_t^{\text{myopic}} = -\frac{v_t}{2} + n, \quad A n = 0, \quad (67)$$

and the maximal wedge 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 (61) 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

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

which rearranges to (66). Since  $A$  is positive semidefinite,  $z^\top Az$  is convex in  $z$ , so  $-(u + \frac{v_t}{2})^\top A(u + \frac{v_t}{2})$  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 (66) yields the maximal value. If  $A$  is positive definite, then  $An = 0$  implies  $n = 0$ , so the maximizer is unique.  $\square$

Lemma 5 is the vector analogue of Lemma 3. When  $v_t$  is predetermined, wedge gains arise only from intermediating that flow, and the instantaneous harvesting rent is  $\frac{1}{4}v_t^\top Av_t$ . 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 in the cross-impact geometry.

**Proposition 7.** Fix  $t \in \{0, \dots, T-1\}$  and a state  $s_t$ . Suppose the optimizer  $u_t^*$  in (65) is interior,  $u_t^* \in \text{int } \mathcal{U}_t(x_t)$ , and assume that  $V_{t+1}$  is differentiable at the continuation state  $s_{t+1}^* = \mathcal{T}(s_t, u_t^*)$ . Define

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

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  $u_t^*$  satisfies the Euler equation

$$A(2u_t^* + v_t) = \mu_{t+1} + A\nu_{t+1}. \quad (68)$$

If  $A$  is positive definite, (68) uniquely determines the interior optimizer:

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

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 print under (60) and does not generate wedge 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 (68) pins down  $\Pi_A u_t^*$ , and the canonical choice that sets the impact-neutral component to zero satisfies

$$\Pi_A u_t^* = -\frac{1}{2}\Pi_A v_t + \frac{1}{2}A^\dagger \mu_{t+1} + \frac{1}{2}\Pi_A \nu_{t+1}. \quad (69)$$

If  $A$  is positive definite, then  $A^\dagger = A^{-1}$  and  $\Pi_A = I$ , so (69) reduces to the previous display. Moreover, the resulting current-period wedge payoff can be written as

$$\pi_t(s_t, u_t^*) = \frac{1}{4}v_t^\top Av_t - \left(u_t^* + \frac{v_t}{2}\right)^\top A\left(u_t^* + \frac{v_t}{2}\right) = \frac{1}{4}v_t^\top Av_t - \frac{1}{4}(\mu_{t+1} + A\nu_{t+1})^\top A^\dagger(\mu_{t+1} + A\nu_{t+1}). \quad (70)$$

*Proof.* See Appendix E.2.  $\square$

The decomposition (69) has the same logic as (46), but cross-impact changes what it means to lean against forced flow and what it means to move the print. 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 wedge 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,  $\mu_{t+1}$ , 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  $\nu_{t+1}$  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. When the per-date  $\ell_1$  constraints in (59) bind, the interior Euler equation is replaced by KKT conditions with  $\ell_1$  subgradients; Appendix E.5 records these conditions.

### 6.3 Chain-rule decomposition of the rule-manipulation motive

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

**Proposition 8** (Chain-rule decomposition of  $\nu_{t+1}$ ). *Fix  $t \in \{0, \dots, T-1\}$  and consider any state-action pair  $(s_t, u_t)$ . Let  $P_t$  be given by (60) 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  $\Gamma$  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*

$$\Lambda_{t+1} = \partial_M V_{t+1}(s_{t+1}) - B_{t+1}^\top \nabla_X V_{t+1}(s_{t+1}), \quad (71)$$

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

$$\nu_{t+1} = \partial_P V_{t+1}(s_{t+1}) + (\partial_P F(y_t, P_t))^\top \nabla_y V_{t+1}(s_{t+1}) + \Lambda_{t+1} g'(\Gamma_t) \partial_P \Gamma(y_t, P_{t-1}, P_t). \quad (72)$$

*Proof.* See Appendix E.3. □

The margin-update term in (72) isolates the economics of manipulation in three primitives. The vector  $\partial_P \Gamma(y_t, P_{t-1}, P_t)$  identifies which directions in the price vector move the statistic the rule measures. The scalar  $g'(\Gamma_t)$  captures how strongly the posted requirement responds locally to that statistic. The scalar  $\Lambda_{t+1}$  converts a marginal tightening of next-period requirements into continuation value along the constrained-demand benchmark. Relative to the single-asset case, the key new feature is directional: 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 local 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 price prints 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 (statistic targeting)). *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 (60), 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*

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

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

*Proof.* A marginal trade perturbation  $\delta u$  moves the print 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 in wedge units. Under temporary impact, the instantaneous wedge cost of  $\delta u$  is the quadratic form  $\delta u^\top A \delta u$ , so we minimize this cost subject to delivering the required exposure:

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

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  $\delta u$  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

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

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

Lemma 6 isolates the cheapest way to move the rule's input. Locally, only the statistic gradient  $j_t$

matters: the rule reacts to the current print through the single exposure  $j_t^\top P_t$ , so the lowest-cost way to change  $\Gamma_t$  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. In the common marked-portfolio case where the rule depends on  $d^\top P_t$ , 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*

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

for all relevant  $t$ . Consider the multi-asset dynamic program in Section 6.1 under temporary total-flow impact (60) and constrained demand (55). Consider the relaxed problem in which the per-date trade bound (58) (and hence the admissible set (59)) 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,

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

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

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

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 wedge 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 single marked print  $d^\top P_t$ , 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 neither: it does not help absorb the forced  $b$ -imbalance, and it does not move the marked print the rule uses. Such trades can therefore be deleted without changing the best achievable objective. The optimal attack can be implemented using only  $b$  (the portfolio where liquidation arrives, hence where harvesting rents live) and  $d$  (the portfolio the risk engine effectively watches, hence where triggering power lives). Once attention is restricted to combinations of  $b$  and  $d$ , cross-impact affects outcomes

only through three transparent scalars:

$$b^\top Ab, \quad d^\top Ad, \quad b^\top Ad.$$

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

**Corollary 7.** *Let  $A = U \text{diag}(\lambda_1, \dots, \lambda_N) U^\top$  be an eigen-decomposition with orthonormal  $U$  and  $\lambda_i \geq 0$ . Write the liquidation and marked portfolios in the impact-factor basis as  $\tilde{b} = U^\top b$  and  $\tilde{d} = U^\top d$ . Then*

$$b^\top Ab = \sum_{i=1}^N \lambda_i \tilde{b}_i^2, \quad d^\top Ad = \sum_{i=1}^N \lambda_i \tilde{d}_i^2.$$

*In particular, each impact factor  $i$  contributes to harvesting rents and to triggering power through the same impact intensity  $\lambda_i$ .*

*Proof.* Substitute the eigen-decomposition into the quadratic forms. □

The corollary is easiest to read as a factor decomposition of execution costs. The matrix  $A$  defines impact factors (the columns of  $U$ ) and their intensities ( $\lambda_i$ ). Loading on factor  $i$  moves prices more per unit of trade when  $\lambda_i$  is large. That same intensity scales both sides of the attack problem. It scales harvesting rents because forced liquidation that loads on a high-impact factor generates a larger wedge rent when intermediated. It scales triggering power because moving the marked price requires moving prices, and prices move most per unit of trade in high-impact factors. Optimal attacks therefore concentrate on impact factors where the liquidation portfolio  $b$  and the marked portfolio  $d$  both have sizable loadings, with  $\lambda_i$  governing how valuable it is to operate through that factor.

## 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 exact design problem at the center of this paper. The underlying methodology is publicly disclosed and mechanically implemented: a short return history is mapped into a risk statistic and then into next-day exposure, often with caps, floors, smoothing, and implementation lags (Krein and Fernandez, 2012; 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 not a theoretical abstraction; it is a contractual trading rule that scales with product issuance. As a result, when realized volatility rises and exposure is cut, replication requires predictable rebalancing in the hedge instruments, and when volatility falls the rule predictably re-levers. This is the disclosed feedback loop studied in the previous sections, now operating in an economically meaningful environment in which the rule is written down ex ante and cannot be assumed away.

This setting is valuable because volatility-control indices are both large in practice and unusually transparent in design, so one can evaluate the stability of the rule before scale builds. The methodology is disclosed and fixes the mapping from sampled transaction prices into the risk statistic and then into next-day exposure, which pins down  $(\Gamma, g)$  and the local derivatives that enter the test. The remaining input is the execution-cost technology in the hedge instruments used for replication, summarized by an impact operator  $\mathcal{I}$ ; no holdings, dealer inventories, or ownership data are needed. Given a candidate methodology and an impact calibration (own- and cross-impact), Theorem 1 delivers an explicit capacity bound, namely the largest linked notional for which the induced closed-loop pricing map remains manipulation free in the linearized system.

The application also connects to evidence that demand is not perfectly elastic and that predictable, mechanical trading pressure can move prices even without information (Gabaix and Koijen, 2021; Haddad et al., 2025; Ben-David et al., 2018; Bretscher et al., 2025). Here the source of mechanical demand is the disclosed methodology itself, which fixes both the timing of rebalancing and the local sensitivity of next-day exposure to sampled prints. This transparency makes the diagnostic operational: for a given rule and an impact calibration, one can test whether the induced closed loop admits profitable local round trips and compute how large the rule-following sector can become before the admissibility screen is violated.

## 7.1 Product templates

The methodology is stated at the index level, but the feedback runs through the hedge instruments used to deliver index-linked payoffs. Three parties matter. 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. A hedge desk then replicates the liability using liquid instruments (ETFs, futures, swaps, and cash). When the methodology calls for a rebalance, replication requires a corresponding rebalance in the hedge instruments. A rule written on an index therefore becomes a trading rule in the underlying markets. Appendix F details this institutional chain 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, and a volatility-control rule adjusts the risky weight to keep risk near a stated target. We focus on two templates that capture the feedback channel in our model:

1. *Template A (single underlying)*. 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. *Template B (multi-underlying, fixed basket)*. 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 overlay is then applied to

the selected portfolio. This adds state variables, but the price-to-demand feedback studied here still operates through the volatility-control mapping. For clarity, the main text studies only the volatility-control overlay, treating the reference portfolio as given; Appendix F then shows, step by step, how the published methodology maps into the model inputs (the statistic  $\Gamma$ , the schedule  $g$ , and the implied update timing and exposures).

## 7.2 Mapping to the model and stress-test outputs

Fix a template and a daily sampling convention. The published methodology defines (i) a statistic computed from recent transaction prices,  $\Gamma_t = \Gamma(P_{t-m}, \dots, P_t)$ , and (ii) a schedule that maps this statistic into the next applied requirement or exposure,  $M_{t+1} = g(\Gamma_t)$ , possibly with an explicit implementation lag  $\ell \geq 0$  so that  $M_{t+1} = g(\Gamma_{t-\ell})$ . The methodology therefore pins down the two rule primitives in our model,  $(\Gamma, g)$ , and their local linearization at the binding point: the pass-through slope  $s = g'(\Gamma_0)$  and the statistic Jacobian  $J$  defined by

$$\delta\Gamma = J \delta\Delta P.$$

Market structure supplies the remaining primitives: the constrained-sector sensitivity  $B$  that maps requirement changes into demand changes (cf. (25)) and an execution-cost model in the hedge instruments summarized by an impact operator  $\mathcal{I}$  via the wedge-price relation  $\delta\Delta P = \mathcal{I}q$ . Timing is captured by the lag operator  $L$  and the first-difference operator  $D$ . With the baseline one-day timing from  $\Gamma_t$  to  $M_{t+1}$ , within-horizon requirement deviations satisfy  $\delta M = s L \delta\Gamma$ . If the methodology adds an extra  $\ell$ -day lag, then  $\delta M = s L^{\ell+1} \delta\Gamma$ , where  $\ell+1$  includes the baseline one-day delay. Substituting these mappings into the forced-flow relation and the impact relation yields the loop operator

$$\mathcal{K} = B s D L^{\ell+1} J \mathcal{I},$$

which is (28) with  $L$  replaced by  $L^{\ell+1}$ . Given  $\mathcal{K}$ , Theorem 1 evaluates admissibility by testing whether the augmented symmetric cost matrix  $\hat{H}$  is positive semidefinite on the round-trip subspace  $\mathcal{R}_T$ . Appendix F works out  $(s, J)$  for Templates A and B and shows how we translate this screen into the capacity objects  $W_{\max}(z)$  and their aggregates.

### 7.2.1 Outputs

Fix a state  $z$  and a linked notional  $W$ . The mapping above delivers the augmented symmetric cost matrix  $\hat{H}(z; W)$ . By Theorem 1, the linearized closed loop is admissible at  $(z; W)$  if and only if  $\hat{H}(z; W)$  is positive semidefinite on the round-trip subspace. We implement the test numerically by restricting  $\hat{H}(z; W)$  to the round-trip subspace and checking its smallest eigenvalue. This yields a statewise capacity bound,

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

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

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

and a vulnerability curve that reports, for each scaled size  $W$ , the fraction of states  $z \in \mathcal{Z}$  at which admissibility fails. We construct  $\mathcal{Z}$  from simulated price paths rather than estimating it from holdings or order flow. This choice matches the objective of the exercise: 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, and execution costs enter only through an externally calibrated impact environment (own-impact and cross-impact). Given these inputs, we simulate realistic return paths, compute  $\hat{H}(z; W)$  state by state, and obtain capacity bounds without using holdings, dealer inventories, order-level data, or proprietary position information. Appendix F details the simulation design and the impact calibration used for the reported results.

### 7.2.2 Optimal attacks against pure target-volatility rules

This subsection 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 and the feedback channel is purely mechanical. We therefore abstract from caps, floors, and turnover buffers in the main text and treat those design features separately. In the interior region, the risky weight is the target volatility divided by the rolling realized volatility estimate, evaluated with an implementation lag  $\ell$ :

$$w_{t+1} = \frac{\sigma^*}{\hat{\sigma}_{t-\ell}}, \quad M_{t+1} = \frac{1}{w_{t+1}} = \frac{\hat{\sigma}_{t-\ell}}{\sigma^*} = \frac{\sqrt{\Gamma_{t-\ell}}}{\sigma^*}, \quad s_t = \frac{\partial M_{t+1}}{\partial \Gamma_{t-\ell}} = \frac{1}{2\sigma^* \hat{\sigma}_{t-\ell}}. \quad (73)$$

Replication in the hedge instrument implies that the rule-following 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 wedge payoff (38). The economic tradeoff is the same as in the general model: the attacker harvests predictable rebalancing by trading against  $v_t$ , while also valuing trades that shift the current print because the print enters future requirements and thereby reshapes future forced flow. At an interior optimum, Proposition 3 implies the decomposition

$$u_t^* = -\frac{v_t}{2} + \frac{1}{2\eta} \mu_{t+1} + \frac{1}{2} \nu_{t+1}, \quad (74)$$

with  $(\mu_{t+1}, \nu_{t+1})$  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 print because the print shifts the rule input and therefore changes future  $M$  and  $v$ . Proposition 4 expresses  $\nu_{t+1}$  through the rule primitives. In a target-volatility methodology, the objects entering that chain rule are explicit: the schedule slope is given by (73), the state update  $F$  is the deterministic shift associated with a rolling window, and the statistic derivative with respect to the current print is available in closed form. As a

result, the manipulation motive, and hence the optimal policy, can be read directly from the disclosed local pair  $(s_t, J)$  together with the timing parameters  $(m, \ell)$ .

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

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right), \quad \Gamma_t = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} r_{t-i}^2. \quad (75)$$

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

### 7.2.3 A predictable flow reversal

Fix a calendar date  $j$  and consider a perturbation that changes  $r_j$  while leaving all other returns unchanged. Because (75) is a rolling average of squared returns, the term  $r_j^2$  enters the statistic at every date  $t$  whose  $m$ -day window still contains  $j$ . Holding other returns fixed, the induced perturbation to the statistic is therefore constant across that block:

$$\delta\Gamma_t = \frac{a_{\text{ann}}}{m} \delta(r_j^2) \cdot \mathbf{1}\{t \in \{j, j+1, \dots, j+m-1\}\}. \quad (76)$$

It then drops to zero at  $t = j+m$  when the observation  $r_j$  rolls out of the window. Under the interior target-volatility mapping (73) with an implementation lag  $\ell$ , this block shift in  $\Gamma$  generates a block of tighter requirements  $M$  (and therefore smaller positions  $X = W/M$ ) starting when the affected statistic first feeds into the applied requirement. Concretely, the tightening shows up for the applied requirements  $(M_{t+1})$  with  $t - \ell \in \{j, \dots, j+m-1\}$ , that is for dates

$$t+1 \in \{j+\ell+1, \dots, j+\ell+m\},$$

and it disappears one day later when  $r_j$  rolls out of the window and no longer enters  $\Gamma_{t-\ell}$ . Since forced flow is the first difference of positions,  $v_{t+1} = X_{t+1} - X_t$ , this produces a mechanically predictable reversal: deleveraging when the tightened requirements begin to apply, followed by releveraging when the tightened block ends (shifted by the baseline one-day delay and the disclosed lag  $\ell$ ). This entry-exit pattern is a distinctive feature of hard rolling-window estimators. It is also a transparent instance of the broader procyclicality logic emphasized in the margining and risk-management literature, where backward-looking risk measures mechanically tighten constraints after volatility spikes and can induce forced selling (Glasserman and Wu, 2018; BIS, 2010; ESRB, 2020). Related evidence that predictable hedging demand can affect spot prices is in Ni et al. (2005). The additional ingredient in our setting is that the statistic is computed from transaction prices and has a known Jacobian, so a strategic trader can move the measured volatility input and then intermediate the induced flow. With exponentially weighted or infinite-memory estimators, the same logic applies but

the discrete roll-off is replaced by gradual decay, which blurs the timing of the reversal.

### 7.2.4 The fragility of low-volatility states

The central object for manipulability is not volatility per se; it is the marginal pass-through from a sampled print into next-day forced flow. In the interior region of a target-volatility rule,  $g(\Gamma) = \sqrt{\Gamma}/\sigma^*$ , so the local pass-through slope  $s_t = g'(\Gamma_{t-\ell})$  in (73) is steepest when the measured variance  $\Gamma_{t-\ell}$  is small. Under the fully constrained benchmark, the linked sector holds

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

which implies

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

The product  $B_{t+1}s_t$  is the mechanical gain that converts a marginal statistic change 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 when measured risk is lowest.

The attacker does not choose  $\Gamma$  directly; she chooses trades that move sampled transaction prices, which move returns, which move  $\Gamma$ . For rolling-window variance, the Jacobian  $\partial_P \Gamma$  is largest after quiet spells: a small perturbation to the sampled print produces a larger marginal change in the measured variance when the recent window has low dispersion. Putting the two links together, the marginal effect of a sampled print on next-day forced flow is proportional to the product

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

This is the object that governs incentives. Proposition 4 shows that the margin-update component of the price shadow value is  $\nu_{t+1}^M = \Lambda_{t+1}s_t\partial_P \Gamma$ , so the incentive to move the current print is steepest precisely in states where that product is largest.

This leads to the following idea of fragility: quiet states are maximally manipulable. They are the states in which (i) the rule-following sector is largest, so a given marginal tightening produces a larger dollar flow, (ii) the schedule is steepest, so the same marginal change in measured risk generates the largest change in next-day requirements, and (iii) the volatility estimator is most sensitive at the margin to the sampled print because recent realized variance is low. These are not crisis states. They are the states in which the rule is designed to take the most risk and to be most responsive. The implication is therefore about incentives: in tranquil states the attacker gets the highest return, per unit of impact cost, from shifting the statistic and then harvesting the induced reversal.

This perspective is distinct from the standard 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 fragility is a property of the disclosed mapping from prices into requirements: even holding the rule fixed, the local derivatives that govern incentives are largest in low-

volatility states. This is the rule-based “paradox of financial instability” applied to a concrete, priced-to-trade mechanism (Borio and Drehmann, 2009; Nicolai, 2026).

For volatility-managed strategies (Moreira and Muir, 2017; Barroso and Santa-Clara, 2015; Kirby and Ostdiek, 2012), the point is sharper than the idea that they simply de-risk after spikes. The strategy intentionally loads up after quiet spells, and those are the states in which the mapping from sampled prints to future forced flow is most sensitive at first order. That is why, in our admissibility and capacity calculations, the binding constraints typically come from low-volatility configurations rather than only from stress episodes. A practical implication is that the test can be used as a forward-looking stress test for benign states: given a disclosed methodology and standard impact inputs, one can quantify how close low-volatility, high-exposure configurations are to violating the admissibility condition, even though those configurations may look safest under backward-looking risk measurement.

### 7.2.5 Inventory-light volatility triggering

Appendix F derives the statistic Jacobian for (75). For any calendar date  $j$ ,

$$\frac{\partial \Gamma_t}{\partial P_j} = \frac{2a_{\text{ann}}}{m P_j} \left( \mathbf{1}\{t - m + 1 \leq j \leq t\} r_j - \mathbf{1}\{t - m \leq j \leq t - 1\} r_{j+1} \right), \quad (77)$$

where  $r_j = \log(P_j/P_{j-1})$ . The structural implication is that a single print  $P_j$  loads into two adjacent returns with opposite signs: it raises  $r_j$  and lowers  $r_{j+1}$ . Because  $\Gamma_t$  aggregates squared returns, this two-adjacent-returns channel makes it possible to move the volatility input materially using short-lived variation in the print, even if the price level is quickly restored.

To see the mechanism transparently, consider a two-day “up-then-down” disturbance that changes returns by  $+a$  on day  $j$  and by  $-a$  on day  $j + 1$ , so the level is largely reversed over two days. The impact on the realized-variance statistic is second order because it operates through squared returns. Holding the rest of the window fixed, the incremental contribution of these two returns to  $\Gamma_t$  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 raise the statistic because the transformation is convex. The attacker therefore spends execution cost to manufacture a volatility signal, but can unwind the level quickly, keeping net inventory small relative to the gross trading used to generate the two returns. In the dynamic program this matters because  $J$  enters the margin-update component of the price shadow value through  $\nu_{t+1}$  (Proposition 4): the attacker can increase continuation value by moving  $\Gamma$  without having to carry a large directional position.

This channel is distinct from the 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 optimally 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 the liquidation schedule is pinned down by a disclosed mapping from sampled transaction prices into next-day exposure, so a trader can move the schedule at the margin by trading inside the sampling window. The key feature of target-volatility indices is that the input is a convex function of returns rather than the price level. A short-lived reversal can therefore create two large-magnitude returns that both raise measured volatility while largely undoing the level change, allowing the trigger to be implemented with little net inventory. The economic contribution is that the relevant discipline on manipulation shifts from balance-sheet and inventory risk to execution costs and the local derivatives of the disclosed methodology: the ability to manufacture a volatility spike without warehousing a directional position is exactly what makes quiet states attractive for trigger-and-harvest attacks.

### 7.2.6 Lag, horizon, and rulebook comparative statics.

Disclosed target-volatility methodologies often impose an implementation lag  $\ell$  (Appendix F). In the dynamic program, Corollary 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 margin-feedback motive disappears (formally,  $\nu_{t+1} = 0$  whenever the derivative in Proposition 3 exists), and the continuation problem reduces to intermediation and 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 mechanically determined by harvesting and inventory unwind, with no incentive to move the statistic.

The main methodology parameters  $(\sigma^*, m, \ell)$  map directly into the strength and timing of feedback. The target volatility  $\sigma^*$  scales requirements through  $M_{t+1} = \hat{\sigma}_{t-\ell}/\sigma^*$  and steepens the local pass-through through  $s_t = 1/(2\sigma^*\hat{\sigma}_{t-\ell})$ . Holding linked notional  $W$  fixed, the resulting local 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  $\sigma^*$ : higher targets mechanically raise exposure and amplify the marginal flow response to a given perturbation of the statistic. Window length  $m$  enters through the statistic Jacobian (77). A longer window attenuates the marginal effect of any single return or print on  $\Gamma$  through the explicit  $1/m$  scaling, but it also extends the period over which a given return remains in the window, delaying the roll-off and stretching the induced flow reversal. These effects shift optimal triggering earlier in the horizon and increase the value of having sufficient remaining time to harvest the eventual exit event. Appendix G develops these timing refinements and the associated trade-sign implications.

**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%

### 7.3 Results

We implement the admissibility screen for the two volatility-control templates. 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 linearized closed-loop map implied by the rule and the corresponding state-wise capacity bound  $W_{\max}(z)$ , expressed in days of ADV of the risky leg. For a candidate sector scale  $W$ , the screen fails in state  $z$  whenever  $W > W_{\max}(z)$ .

Two summary objects organize the results. First, the conservative (worst-state) admissible scale is

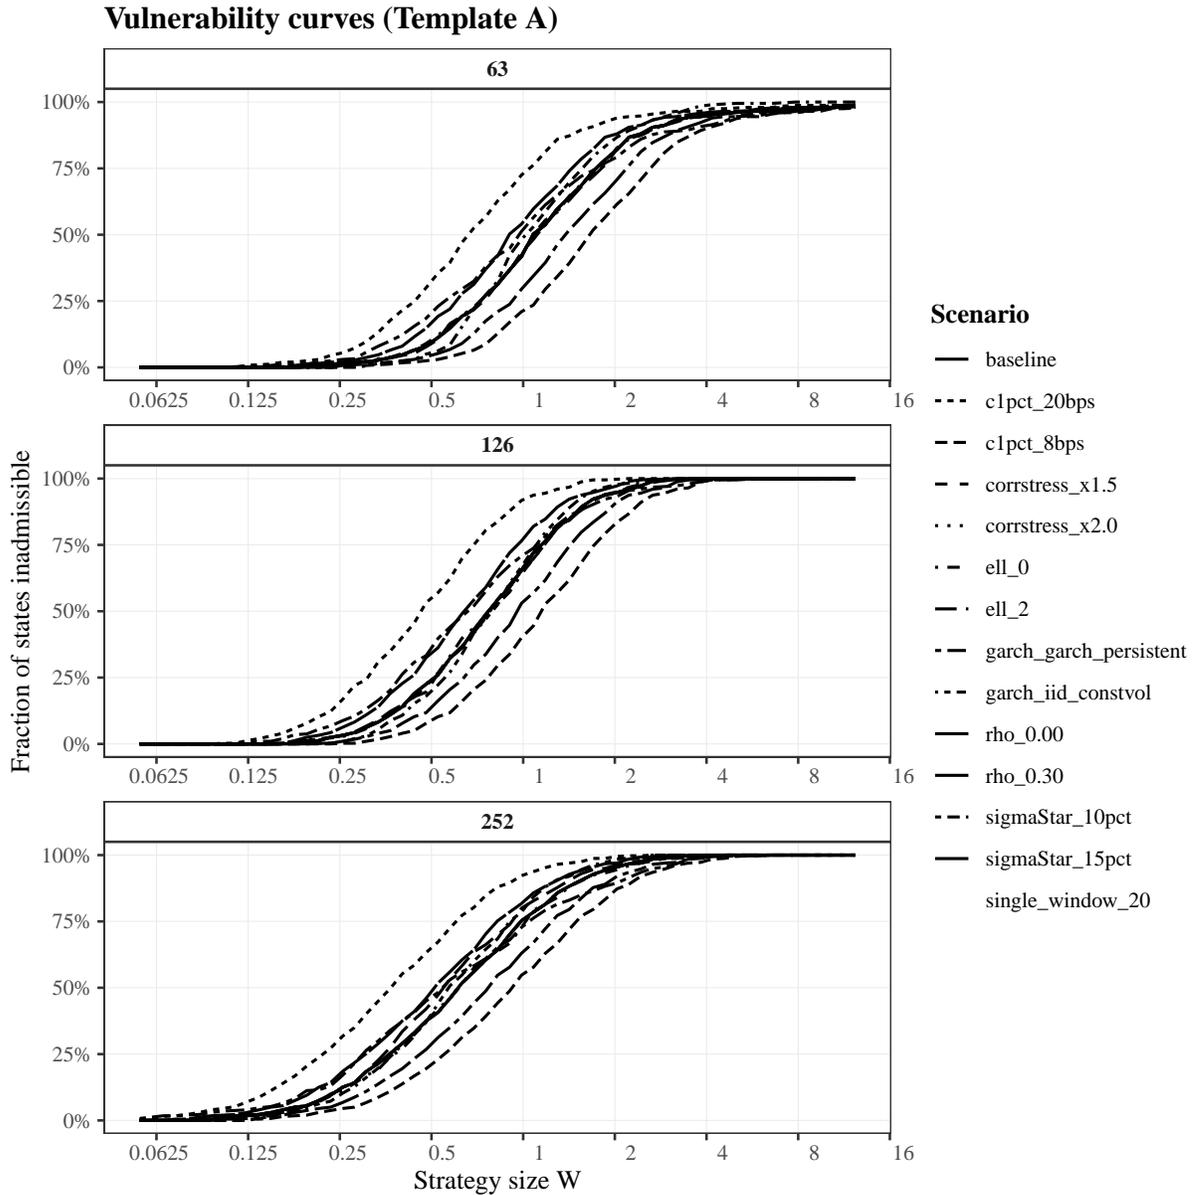
$$W_{\max}(\mathcal{Z}) = \min_{z \in \mathcal{Z}} W_{\max}(z).$$

Second, the vulnerability curve is

$$\Pr_{z \in \mathcal{Z}}[W > W_{\max}(z)],$$

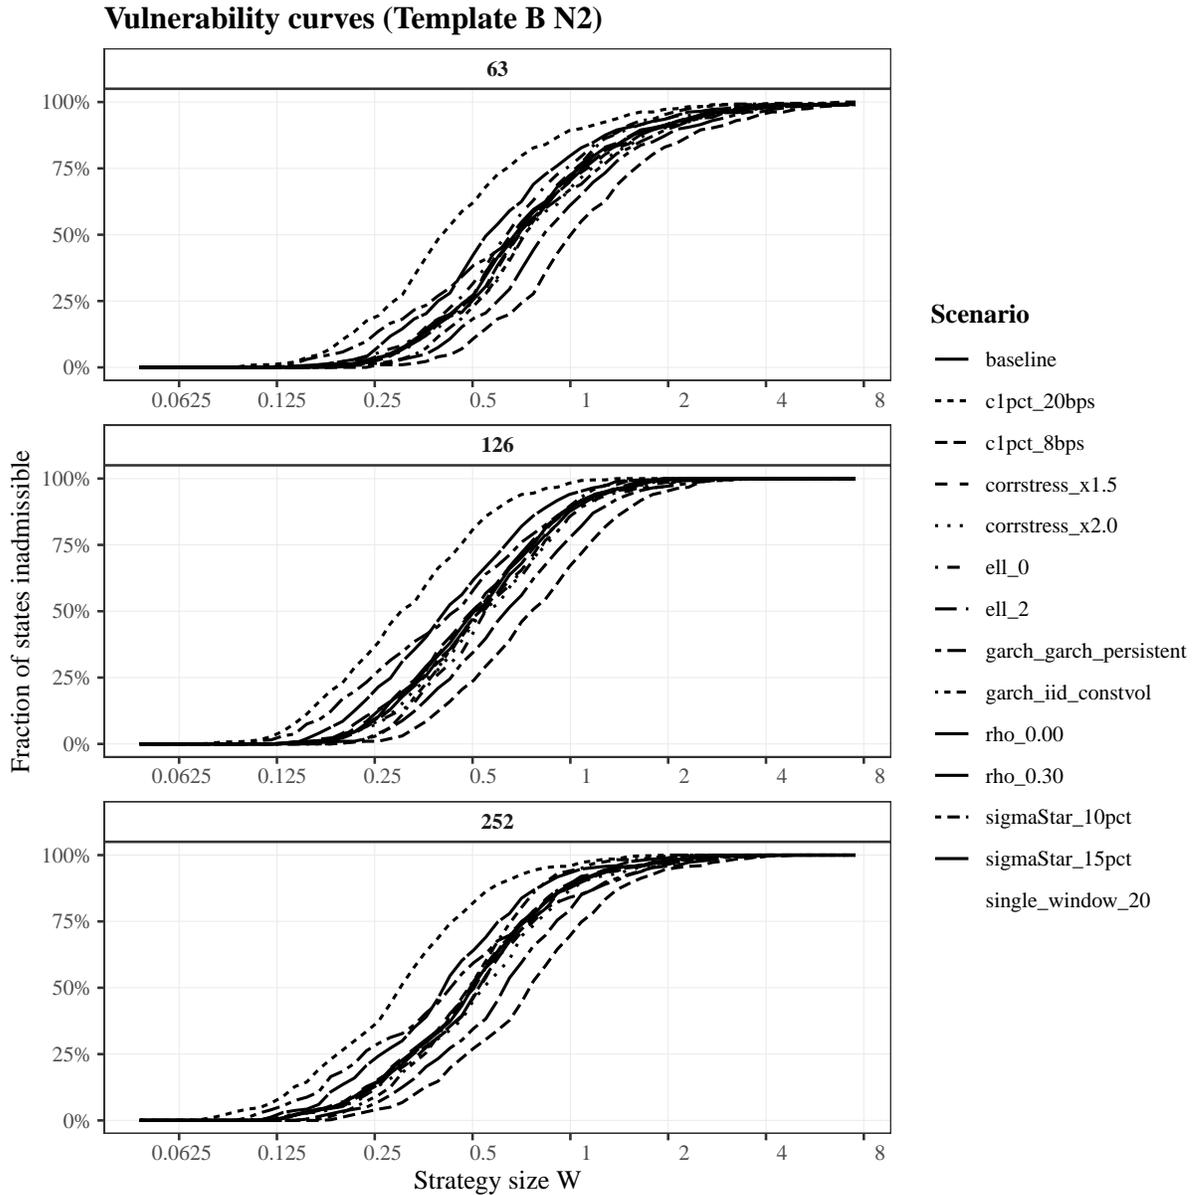
the share of stress states in which a sector of size  $W$  violates the screen. Table 1 summarizes baseline magnitudes at  $T = 126$ . For Template A (single-asset rule), the conservative capacity is  $W_{\max}(\mathcal{Z}) = 0.165$  days of ADV, while the median across stress states is 0.760. For Template 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 screen 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 increase to 94.5% to 99.5%.

Figures 2 and 3 plot the vulnerability curves for the full scenario set in Template A and in Template B with  $N = 2$ , for  $T \in \{63, 126, 252\}$ . These curves complement the conservative bound  $W_{\max}(\mathcal{Z})$  by showing how failure probability varies with sector scale. A 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 the curves upward. In Template B, changes in  $\sigma^*$  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  $\sigma^*$  (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 two primitives of the rulebook-and-liquidity environment: the target volatility  $\sigma^*$  and the market impact impact calibration written as the cost in bps of trading 1% of



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

Comparative statics for  $T = 126$ : target volatility  $\sigma^*$  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)

**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

the volume  $c_{1\%}$ . Lowering the target to  $\sigma^* = 10\%$  raises the 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 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 moves 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 results in this section show that, under the current calibrations and stress design, admissible scale is typically measured in fractions of a day of ADV, not 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 at economically large scales rises quickly: for  $W = 1$  to 2 days of ADV, the screen fails in a large share of stress states even in baseline configurations (Table 1 and

Figures 2–3). In practical terms, if a volatility-control sector is expected to operate at a multi-day-of-ADV scale, the message of these calibrations is that the screen fails in a large share of stress states.

A second lesson is that tail states drive the supervisory risk. Cross-state medians can be several times larger than worst-state bounds (Table 1), so a screen based on typical conditions would materially overstate safe capacity. The diagnostic is designed precisely for this gap: it translates disclosed rulebook sensitivities into a statewise bound  $W_{\max}(z)$ , then aggregates conservatively via  $W_{\max}(\mathcal{Z})$  and transparently via the vulnerability curve.

Third, liquidity is a first-order factor to consider. Changes in the impact calibration  $c_{1\%}$  shift capacity close to proportionally (Table 2), implying that a credible implementation must update the screen as market liquidity changes.

Finally, details on how the volatility-managed strategy is implemented matter. Target scaling interacts with caps and other nonlinear features so that capacity need not move monotonically with  $\sigma^*$  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 function as trading mechanisms. A disclosed mapping from sampled transaction prices into requirements (margin, haircuts, leverage limits, mandated exposures) makes parts of future order flow mechanically predictable. When the constrained sector is large, this creates a feedback channel from current prints to subsequent rule-driven rebalancing. We show that standard no-manipulation conditions for market impact are not sufficient once price-based constraints are operative: even when the impact model rules out price manipulation, the combination of a price-sampled risk statistic and mechanically triggered rebalancing can still generate profitable round trips because today's prints predictably move tomorrow's required trading.

We characterize a mechanism for dynamic arbitrage. An attacker can profit from trigger-and-reverse trading that nudges the sampled risk input, tightens a binding requirement, and then unwinds into the predictable, rule-driven rebalancing flow. The predictability comes from the disclosed rule, so profitable round trips can exist even without large trades or large price moves.

Our main contribution is to derive the no-arbitrage restrictions implied by price-based constraints and to turn them into an implementable test. Theorem 1 delivers a pass/fail condition that depends only on the disclosed rulebook sensitivities and a standard market-impact calibration. This gives regulators an ex ante stress test for mechanical strategies: given a proposed rule and an assumed sector scale, the screen quantifies whether and how often profitable round trips arise under stressed market states. Applying the test state by state yields a capacity bound  $W_{\max}(z)$ ; aggregating over the stress set produces the conservative bound  $W_{\max}(\mathcal{Z})$  and a vulnerability curve that maps sector scale (in ADV units) into the share of stress states admitting profitable round trips.

Beyond the admissibility condition and its implementation as a test, the paper delivers a set of results that pin down the economics of exploitation and the practical design levers. First, we characterize the optimal round trip as the combination of (i) harvesting predictable forced flow, (ii) inventory management required by the terminal round-trip constraint and any speed limits, and (iii)

a distinct rule-manipulation motive that values moving current prints because doing so reshapes future requirements and forced trading (Proposition 3 and its multi-asset analogue). The manipulation motive is disciplined by observables from the disclosed rule: it is proportional to the local sensitivity of the measured statistic to the sampled print and the local slope of the requirement schedule (Propositions 4 and 8). With lags there is a last date at which a trade can still affect within-horizon requirements, after which the problem reduces to intermediation against a predetermined forced-flow path (Section 5 and its endgame results). Second, we show that multi-asset complexity collapses to economically interpretable portfolios. The least-cost way to move the rule’s input is to trade the portfolio aligned with the statistic gradient (Lemma 6). When the rule marks a fixed portfolio direction, an optimal attack can be implemented using only two portfolios, the liquidation direction and the marked direction, and cross-impact matters only through three transparent scalars (Theorem 2). Third, we formalize why timing and within-period dependence are first-order for safety at scale. Lagged updating guarantees well-posedness by construction, while within-period dependence can create extreme amplification and, in some configurations, near-arbitrage sequences (Theorem 3).

The volatility-control application is informative. These indices set next-day exposure as a deterministic function of realized volatility and then rebalance mechanically in liquid hedges, so the rulebook derivatives are known and the only remaining primitive is liquidity. Three lessons emerge. First, worst-case capacity is small even when typical capacity is not: at  $T = 126$ ,  $W_{\max}(\mathcal{Z}) = 0.165$  for Template A and 0.125 for Template B with  $N = 2$  (Table 1), and even at  $N = 8$  the conservative bound is 0.230, well below one day of ADV. Second, vulnerability becomes the relevant object at economically plausible scale. At  $W = 1$  day of ADV, failure rates are already 66.8% to 94.0%, and at  $W = 2$  they are 94.5% to 99.5% (Table 1), so a sector operating at multi-day-of-ADV scale should be expected to violate the screen in most stress states unless the rule and the liquidity environment imply otherwise. Third, the results identify what a supervisor should monitor and what a designer can change. Capacity moves almost proportionally with liquidity through  $c_1\%$  (Table 2), while rulebook features that affect marginal rebalancing, such as caps and cross-impact, materially tighten the worst-state bound (Table 3 and Appendix H). This is exactly the mapping the screen is meant to provide: from a disclosed rule and a liquidity calibration to a quantitative, stress-tested bound on admissible scale.

Price-sampled risk rules should be treated like any other large mechanical trading program: certify them before deployment, size them to market liquidity, and re-certify as conditions change. Theorem 1 makes this operational. Using only the disclosed rulebook and an impact calibration, it delivers a pass/fail screen, a conservative capacity bound, and a vulnerability curve, without relying on position data. The same quantities point to concrete fixes: add implementation lags, cap marginal pass-through in binding states, reduce sensitivity of the statistic to individual prints, and concentrate design changes on the portfolio directions that matter for exploitation. Price-based constraints can be both risk-sensitive and robust, but only if they are stress-tested as mechanisms that generate trading.

## 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.
- Allen, Franklin, Marlene Haas, Eric Nowak, Matteo Pirovano, and Angel Tenguov, "Squeezing Shorts Through Social Media Platforms," *Management Science*, 2025.
- , —, Matteo Pirovano, and Angel Tenguov, "How prevalent are short squeezes? Evidence from the US and Europe," *Journal of Banking & Finance*, 2025, 176, 107436.
- 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.
- Barroso, Pedro and Pedro Santa-Clara, "Momentum Has Its Moments," *Journal of Financial Economics*, 2015, 116 (1), 111–120.
- 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.
- Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, "Do ETFs Increase Volatility?," *The Journal of Finance*, 2018, 73 (6), 2471–2535.

- Benzaquen, Michael, Iacopo Mastromatteo, Zoltan Eisler, and Jean-Philippe Bouchaud, "Dissecting cross-impact on stock markets: An empirical analysis," *Journal of Statistical Mechanics: Theory and Experiment*, 2017, p. P07014.
- 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 Vandenbergh, *Convex Optimization*, Cambridge University Press, 2004.
- Bretscher, Lorenzo, Lukas Schmid, Ishita Sen, and Varun Sharma, "Institutional Corporate Bond Pricing," *The Review of Financial Studies*, 2025.
- 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.
- and Martin Oehmke, "Predatory Short Selling," *Review of Finance*, 2014, 18 (6), 2153–2195.
- 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.
- Da, Zhi and Sophie Shive, "Exchange traded funds and asset return correlations," *European Financial Management*, 2018, 24 (1), 136–168.

- Danielsson, Jon, Hyun Song Shin, and Jean-Pierre Zigrand, "The Impact of Risk Regulation on Price Dynamics," *Journal of Banking & Finance*, 2004, 28 (5), 1069–1087.
- 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.
- Gabaix, Xavier and Ralph S. J. Koijen, "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," Working Paper 28967, National Bureau of Economic Research June 2021.
- Gao, Pingyang, Xu Jiang, and Jinzhi Lu, "Manipulation, panic runs, and the short selling ban," *Journal of Economic Theory*, 2025, 223, 105939.
- 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.
- Glosten, Lawrence R. and Paul R. Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 1985, 14 (1), 71–100.
- Goldman Sachs, "Autocallable Goldman Sachs Momentum Builder Focus ER Index-Linked Notes due 2032 (Form 424B2)," SEC EDGAR filing December 2025.
- Goldstein, Itay and Alexander Guembel, "Manipulation and the Allocational Role of Prices," *Review of Economic Studies*, 2008, 75 (1), 133–164.
- Greenwood, Robin, Augustin Landier, and David Thesmar, "Vulnerable Banks," *Journal of Financial Economics*, 2015, 115 (3), 471–485.
- Grossman, Sanford J. and Merton H. Miller, "Liquidity and Market Structure," *Journal of Finance*, 1988, 43 (3), 617–633.
- Haddad, Valentin, Paul Huebner, and Erik Loualiche, "How Competitive Is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing," *American Economic Review*, 2025, 115 (3), 975–1018.

- Harvey, Campbell R., Edward Hoyle, Russell Korgaonkar, Sandy Rattray, Matthew Sargaison, and Otto Van Hemert, "The Impact of Volatility Targeting," *Journal of Portfolio Management*, 2018, 45 (1), 14–33.
- Hasbrouck, Joel and Duane J. Seppi, "Common Factors in Prices, Order Flows, and Liquidity," *Journal of Financial Economics*, 2001, 59 (3), 383–411.
- Hillion, Pierre and Matti Suominen, "The manipulation of closing prices," *Journal of Financial Markets*, 2004, 7 (4), 351–375.
- Ho, Thomas and Hans R. Stoll, "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial Economics*, 1981, 9 (1), 47–73.
- 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.
- Israeli, Doron, Charles MC Lee, and Suhas A Sridharan, "Is there a dark side to exchange traded funds? An information perspective," *Review of Accounting Studies*, 2017, 22 (3), 1048–1083.
- 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.
- Kirby, Chris and Barbara Ostdiek, "It's All in the Timing: Simple Active Portfolio Strategies that Outperform Naive Diversification," *Journal of Financial and Quantitative Analysis*, 2012, 47 (2), 437–467.
- Koijen, Ralph SJ and Motohiro Yogo, "A demand system approach to asset pricing," *Journal of Political Economy*, 2019, 127 (4), 1475–1515.
- 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., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, 53 (6), 1315–1335.
- 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.
- Ni, Sophie, Neil D. Pearson, and Allen M. Poteshman, "Stock Price Clustering on Option Expiration Dates," *Journal of Financial Economics*, 2005, 78 (1), 49–87.
- 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.
- 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.

Vayanos, Dimitri and Paul Woolley, “An Institutional Theory of Momentum and Reversal,” *The Review of Financial Studies*, 2013, 26 (5), 1087–1145.

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*

$$|x_t(i)| M_t \leq E(i). \quad (78)$$

*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:*

$$x_t(i) = \text{sgn}(\bar{x}_t) \min \left\{ |\bar{x}_t|, \frac{E(i)}{M_t} \right\}. \quad (79)$$

*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  $\bar{x}_t \geq 0$  and  $\bar{x}_t \geq \max_{i \in [0,1]} E(i)/M_t$ , 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*

$$X_t = \frac{W}{M_t}, \quad \frac{dX_t}{dM_t} = -\frac{W}{M_t^2}. \quad (80)$$

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

$$v_t = X_t - X_{t-1}. \quad (81)$$

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  $V : \mathbb{R} \rightarrow \mathbb{R}$  be concave. Consider the investor problem

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

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

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

In particular, if  $x^u \geq 0$ , 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  $x^u \geq 0$ , the cap  $E/M$  falls in  $M$ , so  $x(M)$  is weakly decreasing, with derivative  $-E/M^2$  wherever binding.  $\square$

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

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

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  $X : (0, \infty) \rightarrow \mathbb{R}$  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  $(M_t, M_{t+1}) \rightarrow (M_0, M_0)$ ,

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

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

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

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  $B = W/M_0^2$ .  $\square$

## B Augmented no-arbitrage condition

In this appendix we provide two complements to Section 4. First, we characterize when the linear system  $(\text{Id} + \mathcal{K})q = u$  is well posed. Second, we derive the eigenvalue-based simplification used in Section B.1. For a given attack path  $u$ , invertibility of  $\text{Id} + \mathcal{K}$  pins down a unique induced total-flow path  $q$  and hence a unique execution wedge  $P - S$  through the impact operator. If  $\text{Id} + \mathcal{K}$  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  $\text{Id} + \mathcal{K}$  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.

Let  $\mathcal{U}_T$  denote the class of admissible strategies (predictable and uniformly bounded). To ensure that the variance of the profits is well defined, assume square integrability of the benchmark: in the single-asset model,  $\mathbb{E}[S_t^2] < \infty$  for all  $t$  (in addition to Assumption 1), and in the multi-asset model,  $\mathbb{E}[\|S_t\|^2] < \infty$  for all  $t$  (in addition to Assumption 3). Under admissibility and a finite horizon, this implies  $\Pi(u) \in L^2$  and therefore  $\mathbb{V}\text{ar}(\Pi(u)) < \infty$  for all  $u \in \mathcal{U}_T$ . For  $u \in \mathcal{U}_T$ , realized profit  $\Pi(u)$  is defined as in Definition 3, using the unique execution-price path generated by the closed loop.

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

(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  $\mathbb{E}[\Pi(u^n)] \rightarrow \infty$  and*

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

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 (23)–(29). 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*

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

where  $\text{Id}$  is the  $T \times T$  identity matrix,  $\mathcal{K}$  is the loop operator in (28), and  $\mathcal{I}$  is the (linearized) impact operator. The following statements are equivalent:

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

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

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

*Proof.* (i) $\Rightarrow$ (ii). If  $\text{Id} + \mathcal{K}$  is invertible, then for each  $u$  the unique solution is  $q = (\text{Id} + \mathcal{K})^{-1}u$ . The wedge path follows uniquely from  $P - S = \mathcal{I}q$ .

(ii) $\Rightarrow$ (iii). Statement (ii) applies to every  $u \in \mathbb{R}^T$ , hence in particular to every  $u \in \mathcal{R}_T \subset \mathbb{R}^T$ .

(iii) $\Rightarrow$ (i). Since  $0 \in \mathcal{R}_T$ , statement (iii) applied to  $u = 0$  yields uniqueness of the solution to  $(\text{Id} + \mathcal{K})q = 0$ , hence  $\text{null}(\text{Id} + \mathcal{K}) = \{0\}$ . For a square matrix,  $\text{null}(A) = \{0\}$  is equivalent to  $A$  being nonsingular (hence invertible); see (Stachurski and Sargent, 2022, Theorem 6.1.14).  $\square$

**Lemma 11** (Invertibility of  $\text{Id} + \mathcal{K}$  under lagged updating). *Fix a horizon  $T$  and stack time-indexed paths in chronological order as column vectors  $x = (x_0, \dots, x_{T-1})^\top$ . 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)  $\Gamma_t$  depends only on prices up to date  $t$ , so its Jacobian  $J_{t,s} = 0$  for all  $s > t$ . Then the map  $q \mapsto \mathcal{K}q$  is such that its  $T \times T$  matrix representation,  $\mathcal{K}_{t,s} = 0$  for all  $s \geq t$ . In particular,  $\mathcal{K}$  is strictly lower triangular,  $\text{Id} + \mathcal{K}$  is invertible on  $\mathbb{R}^T$ , and*

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

*Proof.* Given  $x = (x_0, \dots, x_{T-1})^\top$ , the one-step lag operator  $L$  satisfies  $(Lx)_0 = 0$  and  $(Lx)_t = x_{t-1}$  for  $t = 1, \dots, T-1$ , 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  $t = 1, \dots, T-1$ , 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

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

so  $\text{Id} + \mathcal{K}$  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).  $\square$

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 (23)–(29).*

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

(i) **Lagged updating implies automatic well-posedness.** *Under the timing  $M_{t+1} = g(\Gamma_t)$ ,  $\text{Id} + \mathcal{K}$  is invertible on  $\mathbb{R}^T$  (Lemma 11). Hence for each  $u \in \mathbb{R}^T$  there is a unique induced flow path  $q$  and a unique wedge path  $P - S = \mathcal{I}q$ .*

(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 loop operator, so the fixed point becomes*

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

Invertibility of  $\text{Id} + \mathcal{K}^0$  is no longer automatic. If  $\text{Id} + \mathcal{K}^0$  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  $\text{Id} + \mathcal{K}$  is invertible, then

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

and the same identity holds with  $\mathcal{K}$  replaced by  $\mathcal{K}^0$  whenever  $\text{Id} + \mathcal{K}^0$  is invertible. Hence if  $\sigma_{\min}(\text{Id} + \mathcal{K}) \downarrow 0$  (or  $\sigma_{\min}(\text{Id} + \mathcal{K}^0) \downarrow 0$ ) along a sequence of configurations while invertibility is maintained, arbitrarily small inputs can generate arbitrarily large induced flows, and therefore arbitrarily large wedge-price responses.

Moreover, at any well-posed configuration where the sign restriction in Theorem 1 fails (equivalently,  $\widehat{H}$  has a negative direction on  $\mathcal{R}_T$ ), there exists a margin-feedback quasi-arbitrage in the sense of Definition 4(ii).

*Proof.* (i) See Lemma 11.

(ii) If  $\text{Id} + \mathcal{K}^0$  is singular, there exists  $q_0 \neq 0$  with

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

For the round-trip input  $u = 0 \in \mathcal{R}_T$ , both  $q = 0$  and  $q = q_0$  solve  $(\text{Id} + \mathcal{K}^0)q = u$ , so the induced flow path is not unique. The induced wedge paths differ as well. If  $\mathcal{I}q_0 = 0$ , then  $\mathcal{K}^0q_0 = 0$  because the loop operator acts only through the wedge channel, hence  $(\text{Id} + \mathcal{K}^0)q_0 = q_0$ , contradicting  $q_0 \neq 0$ . Thus  $\mathcal{I}q_0 \neq 0$ , so  $P - S = \mathcal{I}q$  is not uniquely determined at  $u = 0$ . Since realized profit is computed from the execution-price path, it is undefined.

(iii) By definition of the  $\ell_2$  norm,

$$\sup_{\|u\|_2=1} \|(\text{Id} + \mathcal{K})^{-1}u\|_2 = \|(\text{Id} + \mathcal{K})^{-1}\|_2,$$

and likewise with  $\mathcal{K}$  replaced by  $\mathcal{K}^0$  whenever  $\text{Id} + \mathcal{K}^0$  is invertible. To relate this norm to singular values, let  $A$  be any invertible matrix and write

$$A = U\Sigma V^\top, \quad \Sigma = \text{diag}(\sigma_1, \dots, \sigma_T), \quad \sigma_1 \geq \dots \geq \sigma_T > 0.$$

Then

$$A^{-1} = V\Sigma^{-1}U^\top,$$

so the singular values of  $A^{-1}$  are  $(1/\sigma_T, \dots, 1/\sigma_1)$ . Since  $\|B\|_2 = \sigma_{\max}(B)$  for any matrix  $B$ ,

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

Applying this with  $A = \text{Id} + \mathcal{K}$  (or  $A = \text{Id} + \mathcal{K}^0$ ) yields the displayed identity; see (Boyd and

Vandenberghe, 2004, Section A.5.4). The amplification statement follows immediately: as  $\sigma_{\min}(A) \downarrow 0$ , the worst-case gain  $\|A^{-1}\|_2$  diverges. For the quasi-arbitrage claim, failure of the sign restriction in Theorem 1 means that  $\widehat{H}$  has a negative direction on  $\mathcal{R}_T$ . Thus there exists a deterministic vector  $u_0 \in \mathcal{R}_T$  with  $u_0^\top \widehat{H} u_0 < 0$ . View  $u_0$  as a deterministic strategy; choosing the admissibility bound large enough makes  $u_0 \in \mathcal{U}_T \cap \mathcal{R}_T$ , and (32) gives  $\mathbb{E}[\Pi(u_0)] = -\frac{1}{2} u_0^\top \widehat{H} u_0 > 0$ . For  $n \geq 1$ , define  $u^n = n u_0$ . Each  $u^n$  is again a deterministic admissible round trip. Under the linearized loop, the induced wedge path satisfies  $P - S = \mathcal{I}(\text{Id} + \mathcal{K})^{-1} u$  (or  $P - S = \mathcal{I}(\text{Id} + \mathcal{K}^0)^{-1} u$ ), hence scales linearly with  $u$  and the wedge-profit term is a deterministic quadratic form in  $u$ . Using  $\mathbb{E}[u^\top S] = 0$  on round trips (Lemma 1), we therefore have

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

Moreover, because  $u_0$  is deterministic and the wedge path is a deterministic linear function of  $u$ , the wedge component of  $\Pi(u^n)$  is nonrandom; all randomness comes from the unaffected-price term  $u^n^\top S$ . Under square-integrability,  $u_0^\top S \in L^2$ , so the variance terms below are finite. Hence

$$\text{Var}(\Pi(u^n)) = \text{Var}(u^n^\top S) = n^2 \text{Var}(u_0^\top S),$$

and therefore

$$\frac{\mathbb{E}[\Pi(u^n)]}{\sqrt{\text{Var}(\Pi(u^n))}} = n \frac{\mathbb{E}[\Pi(u_0)]}{\sqrt{\text{Var}(u_0^\top S)}} \rightarrow \infty.$$

This is the quasi-arbitrage logic in Huberman and Stanzl (2004). □

Part (iii) uses the standard quasi-arbitrage notion of Huberman and Stanzl (2004). The result implies that once the sign condition fails at a well-posed configuration, there exists a round-trip direction  $u_0$  whose expected profit is positive under the linearized closed loop, and scaling that same trading pattern up produces a sequence of admissible round trips  $u^n = n u_0$  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 feedback generates a predictable execution wedge that scales deterministically with trade size, whereas fundamental price risk enters only through the benchmark martingale and scales linearly. The result is that the Sharpe ratio of the scaled round trip can be made arbitrarily large, which is exactly the sense in which Definition 4(ii) 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 (31), 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  $\hat{H}$  is positive semidefinite on  $\mathcal{R}_T$  if and only if  $\lambda_{\min}^{\mathcal{R}_T}(\hat{H}) \geq 0$ . Equivalently, with  $Q = -\frac{1}{2}\hat{H}$ ,

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

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

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

## C Robustness to nonlinearity

Theorem 1 is obtained by linearizing the rulebook components  $(\Gamma, g, X)$  around a binding configuration. This appendix shows that the screen is not a linearization artifact: under local smoothness and away from kinks, if the screen fails then the exact nonlinear feedback system admits a genuinely profitable sufficiently small round trip. With linear price formation in total flow, the closed-loop map from strategic order flow to wedge prices is differentiable at the reference configuration, so expected profit has a second-order expansion whose quadratic term is precisely the matrix  $\widehat{H}$  used in the test.

Fix a reference configuration  $z$  where the requirement is binding, and fix the sector scale  $W$ . Keep the same linear impact model for prices, and keep the rulebook and constrained-demand maps in their exact form. Focus on a small neighborhood of  $z$  where no caps, floors, or other piecewise components switch regime, so the mapping is locally smooth. At  $z$ , the induced feedback is summarized by the loop operator  $\mathcal{K}(z; W)$ , and the second-order term in expected profit is the closed-loop cost matrix  $\widehat{H}(z; W)$  from Theorem 1.

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

- (i) *The rulebook map from transaction prices to the statistic  $\Gamma$ , 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) *The linearized closed loop is well posed:  $\text{Id} + \mathcal{K}(z; W)$  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*

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

*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 (17)–(18) 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  $v = \Phi(q)$ . The closed loop is the fixed point

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

Under (i) and (ii),  $\Phi$  is continuously differentiable near the reference point and satisfies  $\Phi(0) = 0$ . Its derivative at the reference point is the negative of the linearized loop operator,  $D\Phi(0) = -\mathcal{K}(z; W)$ , by the same chain-rule construction used in Theorem 1. Define

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

Then  $\partial_q F(0, 0) = \text{Id} + \mathcal{K}(z; W)$ , 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  $Dq(0) = (\text{Id} + \mathcal{K}(z; W))^{-1}$ . □

**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 screen 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  $\varepsilon$  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  $Dq(0) = (\text{Id} + \mathcal{K}(z; W))^{-1}$ . Hence, for  $\varepsilon$  small,

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

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

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

where the last equality uses (31). 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  $\varepsilon^2$ . The remainder is negligible relative to  $\varepsilon^2$  for sufficiently small  $\varepsilon$ , hence there exists  $\bar{\varepsilon} > 0$  such that  $\mathbb{E}[\Pi(U^\varepsilon)] > 0$  for all  $\varepsilon \in (0, \bar{\varepsilon})$ . Admissibility holds for  $\varepsilon$  small because  $U^\varepsilon$  is deterministic (hence predictable) and its per-date magnitudes scale linearly in  $\varepsilon$  (Definition 3 and its multi-asset analogue).  $\square$

A negative direction of  $\widehat{H}(z; W)$  means that, at the binding configuration, the feedback dominates impact. The trader nudges the sampled statistic, which moves the next requirement; because the constraint is binding, the sector must rebalance in a predictable direction, and that induced flow improves the trader's execution when the position is unwound. Thus failure of the test identifies a real local arbitrage at  $(z, W)$ ; passing the test rules out these small trigger-and-reverse deviations in that state.

## 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 (39) (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  $(\varepsilon_{t+1})$ .

The stochastic value function is

$$\tilde{V}_t(\tilde{s}) = \sup_{(u_\tau)_{\tau=t}^{T-1}} \mathbb{E} \left[ \sum_{\tau=t}^{T-1} \pi_\tau(\tilde{s}_\tau, u_\tau) \mid \tilde{s}_t = \tilde{s} \right], \quad (82)$$

and the Bellman recursion is

$$\tilde{V}_t(\tilde{s}) = \max_{u \in \mathcal{U}_t(x)} \left\{ \pi_t(\tilde{s}, u) + \mathbb{E} \left[ \tilde{V}_{t+1}(\tilde{s}_{t+1}) \mid \tilde{s}_t = \tilde{s}, u_t = u \right] \right\}, \quad (83)$$

with the terminal condition  $\tilde{V}_T(\tilde{s}) = 0$  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

$$\mu_{t+1} = \mathbb{E} \left[ \partial_x \tilde{V}_{t+1}(\tilde{s}_{t+1}) \mid \tilde{s}_t, u_t^* \right], \quad \nu_{t+1} = \mathbb{E} \left[ \frac{d}{dP_t} \tilde{V}_{t+1}(\tilde{s}_{t+1}) \mid \tilde{s}_t, u_t^* \right],$$

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

$$\mu_{t+1} = \mathbb{E} \left[ \nabla_x \tilde{V}_{t+1}(\tilde{s}_{t+1}) \mid \tilde{s}_t, u_t^* \right], \quad \nu_{t+1} = \mathbb{E} \left[ \nabla_{P_t} \tilde{V}_{t+1}(\tilde{s}_{t+1}) \mid \tilde{s}_t, u_t^* \right],$$

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

### 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  $u = u_t^*$  and denote the resulting values by  $\mu_{t+1}$  and  $\nu_{t+1}$ . The

Bellman objective at date  $t$  is  $\pi_t(s_t, u) + V_{t+1}(\mathcal{T}(s_t, u))$ . By (38),  $\partial_u \pi_t(s_t, u) = -\eta(2u + v_t)$ . 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 (39) 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  $u = u_t^*$ . At an interior optimum,

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

which gives (45) and hence (46). Finally, (47) follows from Lemma 3 and the identity

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

which is obtained by rearranging (45). □

### D.3 Proof of Corollary 1

*Proof.* Since  $0 \in \mathcal{U}_t(x_t)$  by assumption,  $u = 0$  is an admissible deviation in (42). If  $v_t = 0$ , then Lemma 3 gives  $\pi_t(s_t, u) \leq 0 = \pi_t(s_t, 0)$  for all  $u$ . Optimality of  $u_t^*$  in (42) implies

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

so if  $u_t^* \neq 0$  then the continuation term must strictly outweigh the contemporaneous loss  $\pi_t(s_t, u_t^*) \leq 0$ . Under the conditions of Proposition 3, setting  $v_t = 0$  in (46) yields the stated reduction. Finally, if  $\mu_{t+1} = \nu_{t+1} = 0$  at an interior optimum, then (46) gives  $u_t^* = -v_t/2$ . □

### D.4 Proof of Corollary 2

*Proof.* From (46),  $u_t^* = -v_t/2 + \chi_{t+1}/(2\eta)$ . If  $v_t \neq 0$ , then  $u_t^* v_t > 0$  is equivalent to

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

which holds if and only if  $\text{sign}(\chi_{t+1}) = \text{sign}(v_t)$  and  $|\chi_{t+1}| > \eta|v_t|$ . For (ii), the payoff identity (47) implies

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

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

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

By the chain rule,

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

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  $\nu_{t+1}$  and grouping the  $(\partial_M, \partial_X)$  terms yields

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

which is (49) after inserting the definition (48).  $\square$

## D.6 Proof of Corollary 3

*Proof.* The decomposition (49) is a sum of three terms. The only term that depends on  $(g, \Gamma)$  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.  $\square$

## D.7 Proof of Corollary 4

*Proof.* Fix  $t \geq T - 1 - \ell$ . For any  $\tau \in \{t + 1, \dots, T - 1\}$  we have  $M_\tau = g(\Gamma_{\tau-1-\ell})$  with  $\tau - 1 - \ell \leq T - 2 - \ell \leq t - 1$ , so  $M_\tau$  (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  $(v_{t+1}, \dots, v_{T-1})$  are predetermined at time  $t$  and cannot be influenced by actions  $(u_t, \dots, u_{T-1})$ . 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  $t \leq T - 2 - \ell$  then  $t + \ell + 1 \leq T - 1$ , so  $\Gamma_t$  enters  $M_{t+\ell+1} = g(\Gamma_t)$  within the horizon and can affect forced flow. If  $t \geq T - 1 - \ell$ , the first part shows that no within-horizon requirement  $(M_{t+1}, \dots, M_{T-1})$  depends on  $\Gamma_t$ .  $\square$

## D.8 Proof of Proposition 5

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

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

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  $\lambda$  on the inventory constraint. The Lagrangian is

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

The first-order condition for each  $\tau$  is

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

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

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

which gives (50).

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

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

Under (50),

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

which is constant across  $\tau$ . 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 (51).  $\square$

## D.9 Proof of Corollary 5

*Proof.* The first display is exactly (51). If  $v_\tau = c$ , then  $\bar{v}_t = c$  and  $\sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2 = 0$ , so (51) gives  $V_t^{\text{harv}} = \eta x_t c - \frac{\eta}{n} x_t^2$ . Equation (50) yields  $u_\tau^* = -\frac{1}{2}(c - c) - \frac{x_t}{n} = -\frac{x_t}{n}$  for all  $\tau$ . Finally, if  $x_t = 0$  then  $u_\tau^* = 0$  and  $V_t^{\text{harv}} = 0$ .  $\square$

## D.10 Proof of Corollary 6

*Proof.* Under the stated forced-flow path,  $\bar{v}_t = \frac{1}{n} \sum_{\tau=t}^{T-1} v_\tau = v/n$ . Applying (50) with  $x_t = 0$  gives, for  $\tau = t, \dots, T-2$ ,

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

and for  $\tau = T-1$ ,

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

Since  $v < 0$ , the terminal trade is positive. Finally, with  $x_t = 0$ , Corollary 5 gives  $V_t^{\text{harv}} = \frac{\eta}{4} \sum_{\tau=t}^{T-1} (v_\tau - \bar{v}_t)^2$ . 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

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

which yields the stated value.  $\square$

## 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  $\lambda$  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  $\tau$ ,

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

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

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

Setting  $c_t = \lambda/(2\eta)$  yields (52).

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

nondecreasing function

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

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

Finally, if no bound binds at the optimizer, then  $\text{clip}_{[-\bar{u}, \bar{u}]}$  is the identity and stationarity gives  $u_{\tau}^* = -v_{\tau}/2 + c_t$  with a common offset  $c_t$ . Imposing  $\sum_{\tau=t}^{T-1} u_{\tau}^* = -x_t$  yields  $c_t = \bar{v}_t/2 - x_t/(T-t)$ , which recovers (50).  $\square$

## E Proofs and benchmark material 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  $\text{Range}(A)$ ; see [Boyd and Vandenberghe \(2004, Appendix A.5.4\)](#).

**Lemma 13.** *For symmetric and p.s.d.  $A$  and  $\Pi_A$  as above:*

1. For all  $z \in \mathbb{R}^N$ ,  $z^\top Az = (\Pi_A z)^\top A(\Pi_A z)$ .
2. For any  $w \in \text{Range}(A)$ , 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  $u_t^*$  be an interior maximizer in (65). 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 (61) 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,

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

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,  $\nabla_u \Psi(u_t^*) = 0$ , hence

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

which is (68).

**Step 2: decomposition.** Since  $A$  may be singular, (68) identifies  $u_t^*$  only up to additions in  $\ker(A)$ . Choose the canonical representative with  $u_t^* \in \text{Range}(A)$ , so  $\Pi_A u_t^* = u_t^*$ . Multiplying (68) by  $A^\dagger$  gives

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

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

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

hence

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

which is (69).

**Step 3: payoff identity.** By Lemma 5,

$$\pi_t(s_t, u_t^*) = \frac{1}{4} v_t^\top A v_t - \left(u_t^* + \frac{v_t}{2}\right)^\top A \left(u_t^* + \frac{v_t}{2}\right) = \frac{1}{4} v_t^\top A v_t - \frac{1}{4} (2u_t^* + v_t)^\top A (2u_t^* + v_t).$$

Using (68),  $A(2u_t^* + v_t) = \mu_{t+1} + A\nu_{t+1} \in \text{Range}(A)$ . By Lemma 13(2) (see Boyd and Vandenberghe (2004, Example 4.5)),

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

Substituting yields (70).

### 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,

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

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

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

so grouping the  $g'(\Gamma_t) \partial_P \Gamma$  terms yields (72) with  $\Lambda_{t+1}$  defined in (71).

## E.4 Proof of Theorem 2

Assume the rule is portfolio marked in direction  $d$  as in the theorem statement. Let  $\mathcal{W} = \text{span}\{b, d\}$  and write  $B = [b \ d] \in \mathbb{R}^{N \times 2}$ . Define the  $2 \times 2$  Gram matrix  $G = B^\top A B$  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

$$c^* = G^\dagger B^\top A u, \quad \Pi_{\mathcal{W}}u = Bc^* = B G^\dagger B^\top A u,$$

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.

$$w^\top A (u - \Pi_{\mathcal{W}}u) = 0 \quad \text{for all } w \in \mathcal{W}. \quad (84)$$

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

$$b^\top A \Pi_{\mathcal{W}}u = b^\top A u, \quad d^\top A \Pi_{\mathcal{W}}u = d^\top A u. \quad (85)$$

Moreover, (84) implies the Pythagorean identity

$$u^\top A u = (\Pi_{\mathcal{W}}u)^\top A (\Pi_{\mathcal{W}}u) + (u - \Pi_{\mathcal{W}}u)^\top A (u - \Pi_{\mathcal{W}}u), \quad (86)$$

so the projection weakly decreases the  $A$ -norm:

$$(\Pi_{\mathcal{W}}u)^\top A (\Pi_{\mathcal{W}}u) \leq u^\top A u. \quad (87)$$

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 (60),

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

Using (85) with  $u = u_t$  gives  $d^\top A \tilde{u}_t = d^\top A u_t$ , 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  $v_t \in \text{span}\{b\}$ , the interaction term depends on  $u$  only through  $b^\top Au$ . By (85) with  $w = b$ , this interaction term is preserved when replacing  $u_t$  by  $\tilde{u}_t$ . By (87), the direct impact loss weakly falls:  $\tilde{u}_t^\top A\tilde{u}_t \leq u_t^\top Au_t$ . Therefore  $\pi_t(\tilde{u}_t) \geq \pi_t(u_t)$  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  $\mathcal{W} = \text{span}\{b, d\}$  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  $v_t \in \text{span}\{b\}$ , 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 $\ell_1$ per-date trade bound

This subsection records the first-order conditions for the constrained one-period maximization in (65) when the  $\ell_1$  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

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

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  $\ell_1$  balls,

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

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

$$0 \in \nabla_u G_t(u_t^*) + \lambda_t \zeta_t + \kappa_t \xi_T, \tag{88}$$

$$\lambda_t (\|u_t^*\|_1 - \bar{u}) = 0, \quad \kappa_t (\|x_t + u_t^*\|_1 - (T - 1 - t)\bar{u}) = 0. \tag{89}$$

Here  $\partial \|z\|_1$  is the standard  $\ell_1$  subdifferential: componentwise,  $\zeta_i = \text{sign}(z_i)$  if  $z_i \neq 0$  and  $\zeta_i \in [-1, 1]$  if  $z_i = 0$ . Equation (88) reduces to the interior Euler equation (68) 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 rulebook object (American Academy of Actuaries, 2026b; Securities and Exchange Commission, 2024).
3. *Dealer replication and hedge instruments.* Replication is implemented in liquid instruments (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; the empirical inputs are the linked notional and the liquidity and impact of the hedge instruments.

## F.1.2 Single-underlying versus multi-underlying templates

Most marketed designs share a common architecture: a reference portfolio and a cash account, with a volatility-control overlay 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.

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

For the admissibility computations below we focus on Templates A and B. Template C can be accommodated by adding the derivatives of the internal portfolio rule, but the volatility-control overlay that drives the feedback loop is the same object in all three templates.

## F.2 Formulae and timing conventions

This section records the objects that are fixed by disclosure: (i) the volatility estimator, (ii) the exposure (or leverage) rule, and (iii) the timing with which exposures are implemented. We use these documents only to pin down the deterministic map from sampled prices to next-day holdings. The admissibility test then adds market liquidity primitives for the hedge instruments.

### F.2.1 Methodology and disclosure

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.2.2 Template A: single-underlying daily clipped target-vol

The single-underlying template allocates between a Parent 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. Rulebooks typically construct an annualized realized volatility

estimate as a function of recent returns and then take a conservative aggregation across horizons. 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  $\sigma^*$ , the exposure to the Parent Index is set as a ratio of target to estimated volatility, subject to a maximum leverage cap  $L_{\max}$ ,

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

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  $\ell$  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).

We will work with this template in a local neighborhood of a given state. Kinks induced by caps, floors, and turnover buffers are handled with one-sided slopes in the derivative calculations below.

### F.2.3 Template B: multi-underlying risk-control

The multi-underlying template generalizes the same architecture. 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 (e.g., equity and rates sleeves). 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  $\hat{\sigma}_t$  from the history of  $r^{\text{ref}}$  (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.2.4 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 (e.g., MSCI) and in underlier disclosures (MSCI, 2024; Morgan Stanley, 2025).
3. Caps, floors, turnover buffers, rounding, and exchange-holiday conventions create kinks and discrete updates. These are addressed by working locally away from kinks when possible, and using one-sided derivatives and conservative bounds when kinks bind (MSCI, 2024; BlackRock, 2023).

### F.3 Derivatives and construction of the admissibility test

This subsection maps a published rulebook into the objects used in Theorem 1. An index methodology gives a deterministic rule from recent returns to next-day risky weight, including caps, floors, and any implementation lag. Fix a stress-test path  $z$  and horizon  $T$ . The rule then pins down the local derivatives we need: (i) the slope of the exposure schedule with respect to its risk statistic, and (ii) the Jacobian of that statistic with respect to the underlying price inputs. Combining these derivatives with an impact model for the hedge instruments delivers the loop operator  $\mathcal{K}$  and the closed-loop quadratic form  $\hat{H}$  that determines admissibility on round trips.

#### F.3.1 Step 1: Identify $(\Gamma, g)$ and the local slope $s = g'(\Gamma)$

A rulebook specifies (i) a sampled risk statistic  $\Gamma_t$  from recent returns and (ii) a deterministic map from that statistic into next-period risky exposure. In our notation,

$$M_{t+1} = g(\Gamma_{t-\ell}), \quad (90)$$

where  $\ell \in \{0, 1, 2\}$  is the disclosed implementation lag. For a standard target-volatility rule,

$$w_{t+1} = \text{clip}\left(w_{\min}, w_{\max}, \frac{\sigma^*}{\hat{\sigma}_{t-\ell}}\right), \quad M_{t+1} = \frac{1}{w_{t+1}}, \quad (91)$$

with  $\hat{\sigma}_t = \sqrt{\Gamma_t}$ . Many methodologies add a turnover buffer: the weight is held fixed unless the raw update exceeds a threshold  $\tau$ ,

$$w_{t+1}^{\text{raw}} = \text{clip}\left(w_{\min}, w_{\max}, \frac{\sigma^*}{\hat{\sigma}_{t-\ell}}\right), \quad w_{t+1} = \begin{cases} w_t, & \text{if } |w_{t+1}^{\text{raw}} - w_t| \leq \tau, \\ w_{t+1}^{\text{raw}}, & \text{if } |w_{t+1}^{\text{raw}} - w_t| > \tau, \end{cases} \quad M_{t+1} = \frac{1}{w_{t+1}}.$$

When the buffer binds (first case), exposure is locally flat in  $\hat{\sigma}_{t-\ell}$ , so  $s_t = 0$ . On the interior region (cap/floor inactive and buffer inactive),

$$M_{t+1} = \frac{\hat{\sigma}_{t-\ell}}{\sigma^*} = \frac{\sqrt{\Gamma_{t-\ell}}}{\sigma^*}, \quad s_t = \frac{\partial M_{t+1}}{\partial \Gamma_{t-\ell}} = \frac{1}{2\sigma^* \hat{\sigma}_{t-\ell}}. \quad (92)$$

At caps, floors, and buffer kinks we use the appropriate one-sided slope; in particular  $s_t = 0$  whenever exposure is locally flat in  $\hat{\sigma}_{t-\ell}$ .

### F.3.2 Step 2: Compute the statistic Jacobian $J$ for Template A (single-underlying)

Template A uses a single price series  $P_t$  and a rolling realized-variance statistic. Define log returns and the statistic by

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right), \quad \Gamma_t = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} r_{t-i}^2, \quad (93)$$

with window length  $m$  and annualization factor  $a_{\text{ann}} > 0$ . Since  $P_t = S_t + \Delta P_t$ , a wedge-price perturbation  $\delta\Delta P$  enters  $\Gamma$  exactly as a perturbation of  $P$  at the linearization point. Some rulebooks compute realized variance at multiple horizons and set  $\Gamma_t$  equal to the maximum across horizons. For a finite set  $\mathcal{M}$ , define

$$\Gamma_t^{(m)} = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} r_{t-i}^2, \quad m \in \mathcal{M}, \quad \Gamma_t = \max_{m \in \mathcal{M}} \Gamma_t^{(m)}.$$

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, the max is kinked; for implementation we break ties deterministically in favor of the smallest window,  $m_t^* = \min \arg \max_{m \in \mathcal{M}} \Gamma_t^{(m)}$ . Differentiating (93) gives, for any date  $j$ ,

$$\frac{\partial \Gamma_t}{\partial P_j} = \frac{2a_{\text{ann}}}{m P_j} \left( \mathbf{1}\{t - m + 1 \leq j \leq t\} r_j - \mathbf{1}\{t - m \leq j \leq t - 1\} r_{j+1} \right). \quad (94)$$

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 (94) evaluated at the reference state. Only prices inside the rolling window affect  $\Gamma_t$ .

### F.3.3 Step 3: Compute the statistic Jacobian $J$ for Template B (multi-underlying, Option A)

Template B (Option A) scales exposure to a fixed risky basket against cash. Let  $P_t \in \mathbb{R}^N$  collect the constituent prints used by the methodology, and define constituent log returns

$$r_t^k = \log\left(\frac{P_t^k}{P_{t-1}^k}\right), \quad k = 1, \dots, N. \quad (95)$$

Fix basket weights  $d \in \mathbb{R}^N$  and write the basket return as  $\bar{r}_t = d^\top r_t$ . A standard realized-variance statistic is

$$\Gamma_t = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} \bar{r}_{t-i}^2. \quad (96)$$

As in Template A, some rulebooks take the maximum across horizons. For  $\mathcal{M} \subset \mathbb{N}$ ,

$$\Gamma_t^{(m)} = \frac{a_{\text{ann}}}{m} \sum_{i=0}^{m-1} \bar{r}_{t-i}^2, \quad m \in \mathcal{M}, \quad \Gamma_t = \max_{m \in \mathcal{M}} \Gamma_t^{(m)}.$$

Away from ties, let  $m_t^*$  be the active window and take the Jacobian row at date  $t$  from that window. At ties, the max is kinked; for implementation we break ties in favor of the smallest window,  $m_t^* = \min \arg \max_{m \in \mathcal{M}} \Gamma_t^{(m)}$ . Differentiating (96) with respect to a constituent print  $P_j^k$  yields

$$\frac{\partial \Gamma_t}{\partial P_j^k} = \frac{2a_{\text{ann}}}{m P_j^k} d_k \left( \mathbf{1}\{t - m + 1 \leq j \leq t\} \bar{r}_j - \mathbf{1}\{t - m \leq j \leq t - 1\} \bar{r}_{j+1} \right). \quad (97)$$

Stack wedge-price deviations over dates and assets as

$$\delta \Delta P = \left( (\delta \Delta P_0)^\top, \dots, (\delta \Delta P_{T-1})^\top \right)^\top \in \mathbb{R}^{NT}.$$

Then (97) defines the block Jacobian  $J$  such that  $\delta \Gamma = J \delta \Delta P$ , as in (24), with all time operators in Theorem 1 interpreted blockwise in the stacked system.

### F.3.4 Step 4: Assemble the loop operator $\mathcal{K}$ and the augmented matrix $\hat{H}$

Given the statistic Jacobian  $J$  and the local schedule slope  $s$  from Step 1, requirement deviations satisfy

$$\delta M = s L^{\ell+1} \delta \Gamma = s L^{\ell+1} J \delta \Delta P, \quad (98)$$

where  $L$  is the one-step lag operator and  $\ell$  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 (25). Stacking over dates,

$$\delta X = -B (\text{Id}_T \otimes b) \delta M, \quad v = (D \otimes \text{Id}_N) \delta X = -Bs (D \otimes \text{Id}_N) (\text{Id}_T \otimes b) L^{\ell+1} J \delta \Delta P, \quad (99)$$

where  $D$  is the first-difference operator on the  $T$ -date horizon (applied blockwise via  $D \otimes \text{Id}_N$ ). Wedge prices are generated by total flow:  $\delta \Delta P = \mathcal{I}q$ , with  $q = u + v$ . Substituting (99) yields the closed-loop fixed point

$$q = u - \mathcal{K}q, \quad \mathcal{K} = Bs (D \otimes \text{Id}_N) (\text{Id}_T \otimes b) L^{\ell+1} J \mathcal{I}. \quad (100)$$

If  $\text{Id} + \mathcal{K}$  is invertible, then

$$q = (\text{Id} + \mathcal{K})^{-1} u, \quad \delta \Delta P = \mathcal{I}(\text{Id} + \mathcal{K})^{-1} u. \quad (101)$$

Define the effective impact operator  $\tilde{\mathcal{I}} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1}$ . Theorem 1 applies to the symmetric part of  $\tilde{\mathcal{I}}$ ,

$$\hat{H} = \tilde{\mathcal{I}} + \tilde{\mathcal{I}}^\top = \mathcal{I}(\text{Id} + \mathcal{K})^{-1} + \left( (\text{Id} + \mathcal{K})^{-1} \right)^\top \mathcal{I}^\top, \quad (102)$$

and the linearized manipulation-free condition is that  $\hat{H}$  is positive semidefinite on the relevant round-trip subspace (scalar  $\mathcal{R}_T$  in the single-underlying case, stacked  $\mathbf{R}_T$  in the multi-underlying case).

### F.3.5 A one-step closed-form diagnostic

This subsection gives a one-step version of the admissibility condition for the daily target-vol rule. Assume a single asset and temporary total-flow impact, so  $\delta\Delta P_t = \eta q_t$  and  $\mathcal{I} = \eta \text{Id}_2$ . Let

$$\Gamma_t = \frac{a_{\text{ann}}}{m} \sum_{j=0}^{m-1} r_{t-j}^2, \quad r_t = \log(P_t/P_{t-1}),$$

and assume the exposure rule is locally interior (cap/floor and any turnover buffer not binding), so  $M_{t+1} = \hat{\sigma}_t/\sigma^*$  with  $\hat{\sigma}_t = \sqrt{\Gamma_t}$ . Then

$$s_t = \frac{\partial M_{t+1}}{\partial \Gamma_t} = \frac{1}{2\sigma^* \hat{\sigma}_t}, \quad \frac{\partial \Gamma_t}{\partial P_t} = \frac{2a_{\text{ann}}}{m} \frac{r_t}{P_t}.$$

In the fully constrained benchmark  $X = W/M$ , one has  $B = W/M_0^2$ . Evaluating at  $M_0 = M_{t+1} = \hat{\sigma}_t/\sigma^*$  gives the scalar one-step loop coefficient

$$k_t = B s_t \eta \frac{\partial \Gamma_t}{\partial P_t} = \eta \frac{W}{P_t} \frac{\sigma^* a_{\text{ann}}}{m \hat{\sigma}_t^3} r_t, \quad (103)$$

with  $k_t = 0$  whenever the rule is locally flat (cap/floor/buffer binding).

**Corollary 8** (One-step inequality for the daily target-vol rule). *Consider  $T = 2$  trading dates and the immediate  $t \rightarrow t + 1$  channel under temporary impact  $\mathcal{I} = \eta \text{Id}_2$ . In this one-step specialization, admissibility on round trips is equivalent to*

$$2 + k_t \geq 0, \quad (104)$$

where  $k_t$  is given by (103). Equivalently, writing  $r_t^- = \max\{-r_t, 0\}$ ,

$$\eta \frac{W}{P_t} \frac{\sigma^* a_{\text{ann}}}{m \hat{\sigma}_t^3} r_t^- \leq 2. \quad (105)$$

*Proof.* In the one-step specialization,  $\mathcal{K}$  is  $2 \times 2$  strictly lower triangular with single entry  $k_t$  in the  $(2, 1)$  position. With  $\mathcal{I} = \eta \text{Id}_2$ ,

$$\hat{H} = \mathcal{I}(\text{Id} + \mathcal{K})^{-1} + ((\text{Id} + \mathcal{K})^{-1})^\top \mathcal{I}^\top.$$

A direct calculation gives  $\hat{H} = \eta \begin{pmatrix} 2 & -k_t \\ -k_t & 2 \end{pmatrix}$ . 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 (104). Rearranging yields (105).  $\square$

Rearranging (105) gives the one-step capacity implied by the immediate channel:

$$W_{\max}^{(1\text{-step})}(t) = \frac{2m P_t \hat{\sigma}_t^3}{\eta \sigma^* a_{\text{ann}} r_t^-}, \quad W_{\max}^{(1\text{-step})}(t) = +\infty \text{ if } r_t^- = 0. \quad (106)$$

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 rulebook pins down the feedback map from recent returns to next-day exposure (including caps, floors, buffers, and lags), hence  $(\Gamma, g)$  and the local derivatives  $(s, J)$ . The execution-cost model for the hedge instruments pins down the impact operator  $\mathcal{I}$ . Given  $(s, J, \mathcal{I})$ , the formulas in Appendix E produce  $\mathcal{K}$  and  $\widehat{H}$ , hence  $W_{\max}(z)$  for each state  $z$  and the aggregate bound over a stress-test set  $\mathcal{Z}$ .

A stress state  $z$  collects exactly what the rule and execution-cost model condition on at an evaluation date: (i) the recent return history that determines the realized-vol statistic and which regime is active (interior, cap/floor, buffer), and (ii) a local liquidity environment summarized by  $\mathcal{I}$ . Given  $z$ , the rulebook fixes  $(s, J)$  and  $\mathcal{I}$  fixes the wedge response to flow; together they determine  $\mathcal{K}(z; W)$  and  $\widehat{H}(z; W)$ .

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):

$$r_t^k = \sqrt{h_t^k} \varepsilon_t^k, \quad h_{t+1}^k = \omega_k + \alpha (r_t^k)^2 + \beta h_t^k, \quad k = 1, \dots, N,$$

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  $\nu$  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:

$$\begin{aligned} N = 2 : \quad & \sigma_{\text{ann}} = (0.20, 0.10), \quad d = b = (0.6, 0.4), \\ N = 4 : \quad & \sigma_{\text{ann}} = (0.20, 0.15, 0.12, 0.10), \quad d = b = \frac{1}{4} \mathbf{1}, \\ N = 8 : \quad & \sigma_{\text{ann}} = (0.22, 0.20, 0.18, 0.16, 0.14, 0.12, 0.10, 0.08), \quad d = b = \frac{1}{8} \mathbf{1}. \end{aligned}$$

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 wedge cost of  $c_{1\%} \in \{8, 12, 20\}$  bps. In the multi-asset case we use a common  $c_{1\%}$  across constituents (homogeneous self-impact).

Cross-impact is added as a separate sensitivity. Empirically, cross terms can account for a non-trivial share of measured price response (Benzaquen et al., 2017), and commonality in order flow supports a cross-impact channel (Hasbrouck and Seppi, 2001). Our baseline sets  $\mathcal{I}$  diagonal (most conservative and most transparent) and then considers a symmetric cross-impact parameterization

$$\mathcal{I}_{ij} = \rho \sqrt{\mathcal{I}_{ii}\mathcal{I}_{jj}} \text{corr}_{ij} \quad (i \neq j),$$

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 (73) and the rolling realized-volatility statistic (75). Section 7.2.2 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 (77),  $J_{t,j} \neq 0$  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  $\Gamma_t$  only for  $t \in \{j, j + 1, \dots, j + m\}$  (through  $r_j$  and  $r_{j+1}$ ). With lag  $\ell$ , 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  $\Gamma$  is built from squared returns, a spike-and-revert pattern can raise  $\Gamma$  even if the price level ends near its starting point.
- Setting  $j = t$  in (77) gives

$$\frac{\partial \Gamma_t}{\partial P_t} = \frac{2a_{\text{ann}}}{m P_t} r_t.$$

To raise  $\Gamma_t$  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 (76). 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  $X_{t+1} = W/M_{t+1}$ :

$$\delta X_{t+1} \approx -B_{t+1} \delta M_{t+1}, \quad B_{t+1} = \frac{W}{M_{t+1}^2}.$$

If  $\delta M_{t+1}$  is block-constant over the affected dates, then  $\delta X_{t+1}$  is also block-constant, and forced flow is its first difference,  $\delta v_{t+1} = \delta X_{t+1} - \delta X_t$ . Hence  $\delta v_{t+1}$  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 7.2.2,  $\delta \Gamma_t$  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 (73) 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  $\ell$ , 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

$$j + \ell + m + 1 \leq T - 1, \quad \text{equivalently} \quad j \leq T - 2 - \ell - m.$$

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  $\ell$ , the earliest forced flow induced by a date- $t$  manipulation arrives at  $t + \ell + 1$ . A larger  $\ell$  both moves the last feasible trigger date earlier and delays any harvesting. In the terminal region  $t \geq T - 1 - \ell$ , 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 (74) 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

$$\partial_{P_t} \Gamma_t = \frac{2a_{\text{ann}}}{mP_t} r_t,$$

raising  $\Gamma_t$  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  $\Gamma$  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  $\sigma^*$  and falls with  $\Gamma_{t-\ell}$ . Holding the impact environment fixed, the loop operator and the admissibility test in Section 7.2 are therefore most stringent in low-volatility states: when  $\hat{\sigma}_{t-\ell}$  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  $\sigma^*$  scales both the position level and the gain through  $B_{t+1}s_t$ , so higher  $\sigma^*$  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  $M \propto \hat{\sigma}$  and the hard rolling-window design.

## H Additional results for the volatility-control stress tests

This appendix complements Section 7.2.2 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

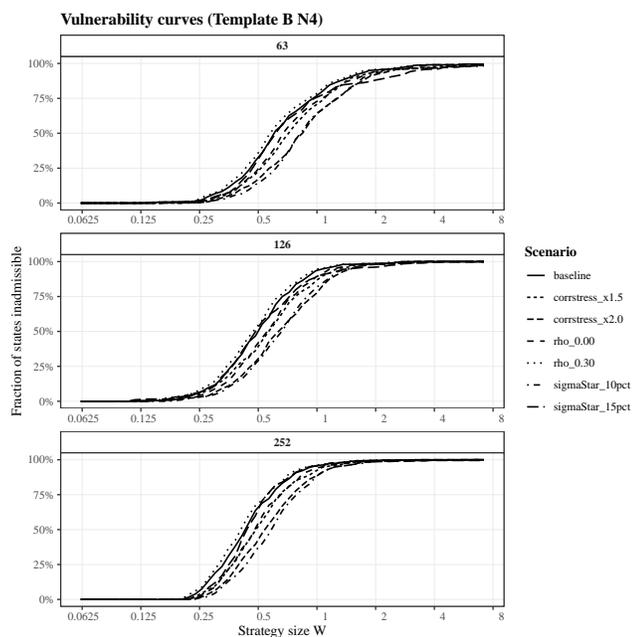
**Table 6** Template B,  $N = 4$ : conservative admissible scale  $W_{\max}(\mathcal{Z})$  by scenario and horizon. Entries are  $W_{\max}(\mathcal{Z}) = \min_{z \in \mathcal{Z}} W_{\max}(z)$  in ADV-normalized units.

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

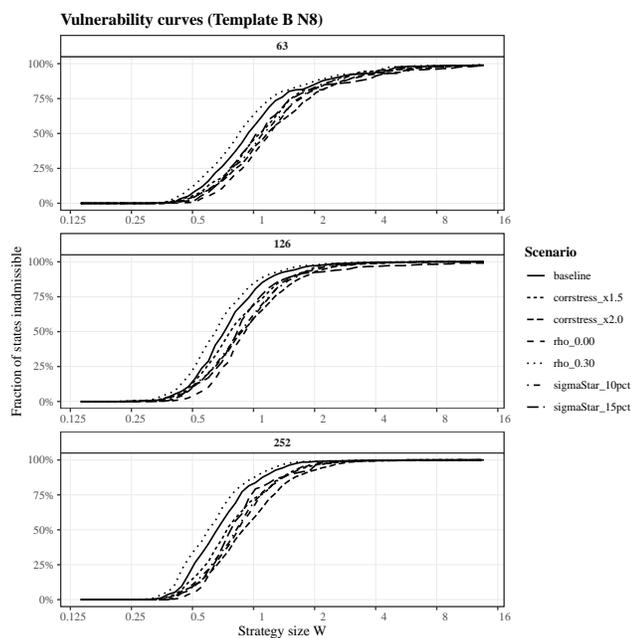
**Table 7** Template B,  $N = 8$ : conservative admissible scale  $W_{\max}(\mathcal{Z})$  by scenario and horizon. Entries are  $W_{\max}(\mathcal{Z}) = \min_{z \in \mathcal{Z}} W_{\max}(z)$  in ADV-normalized units.

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

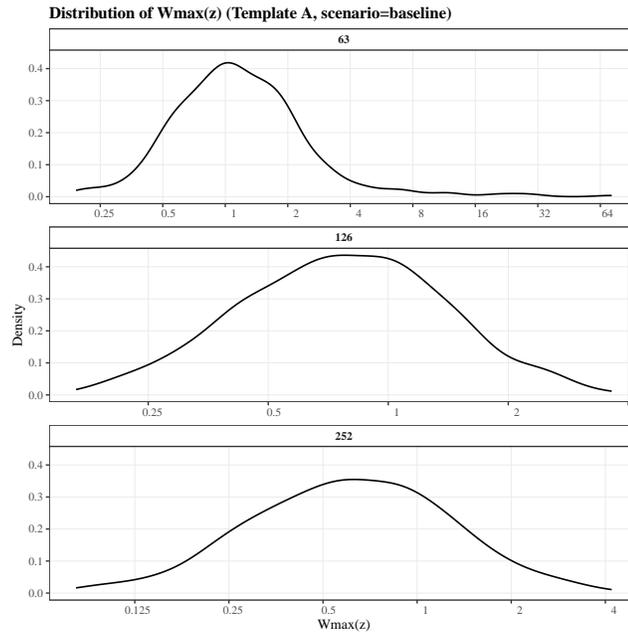


**Figure 4** Vulnerability curves for Template B with  $N = 4$ . Panels correspond to horizons  $T \in \{63, 126, 252\}$ ; lines correspond to the scenarios implemented for  $N = 4$ .

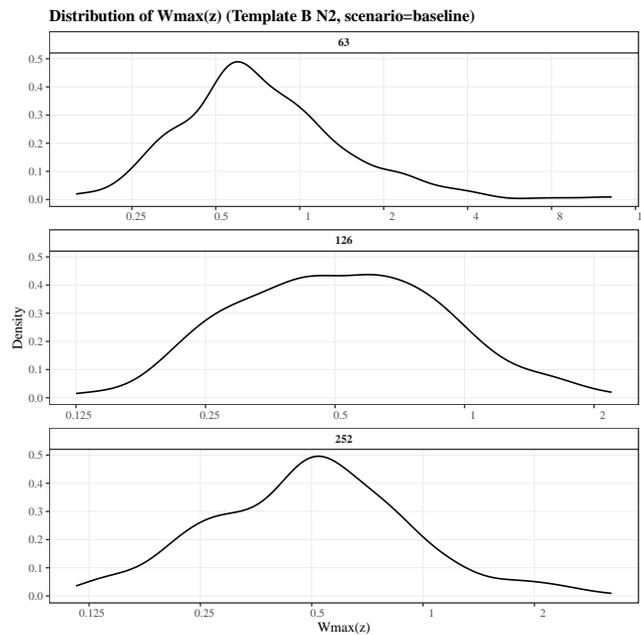


**Figure 5** Vulnerability curves for Template B with  $N = 8$ . Panels correspond to horizons  $T \in \{63, 126, 252\}$ ; lines correspond to the scenarios implemented for  $N = 8$ .

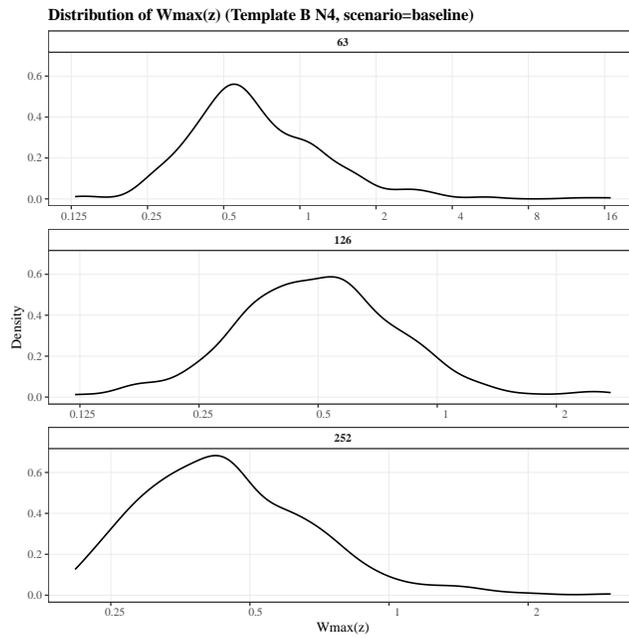
## H.2 Distribution of statewise capacities $W_{\max}(z)$



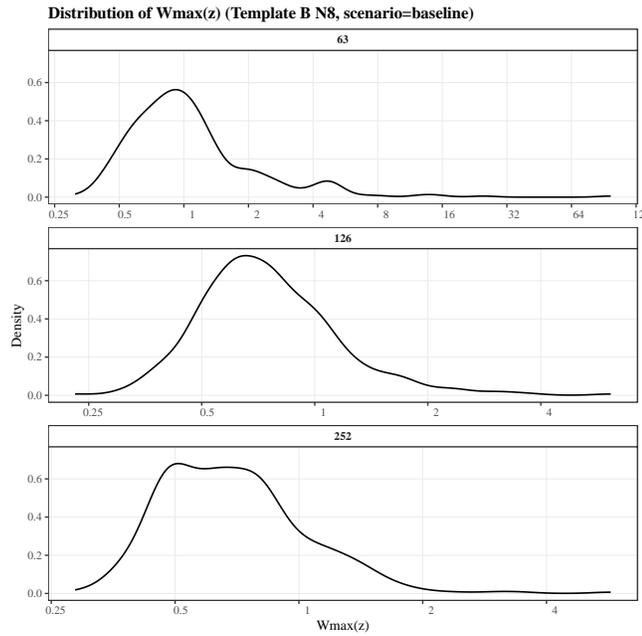
**Figure 6** Kernel density estimate of  $W_{\max}(z)$  across stress states for Template A (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .



**Figure 7** Kernel density estimate of  $W_{\max}(z)$  across stress states for Template B with  $N = 2$  (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .

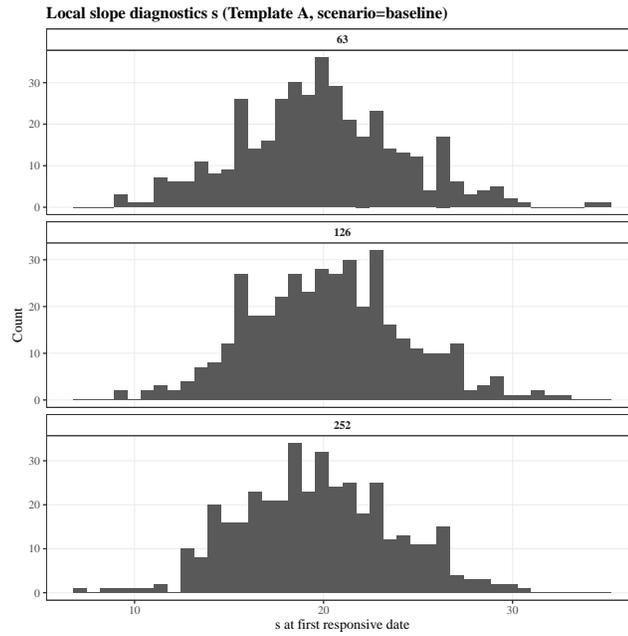


**Figure 8** Kernel density estimate of  $W_{\max}(z)$  across stress states for Template B with  $N = 4$  (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .

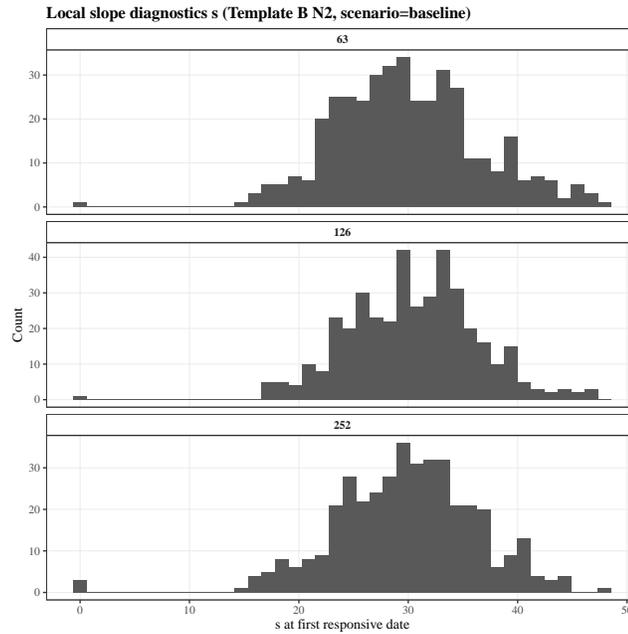


**Figure 9** Kernel density estimate of  $W_{\max}(z)$  across stress states for Template B with  $N = 8$  (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .

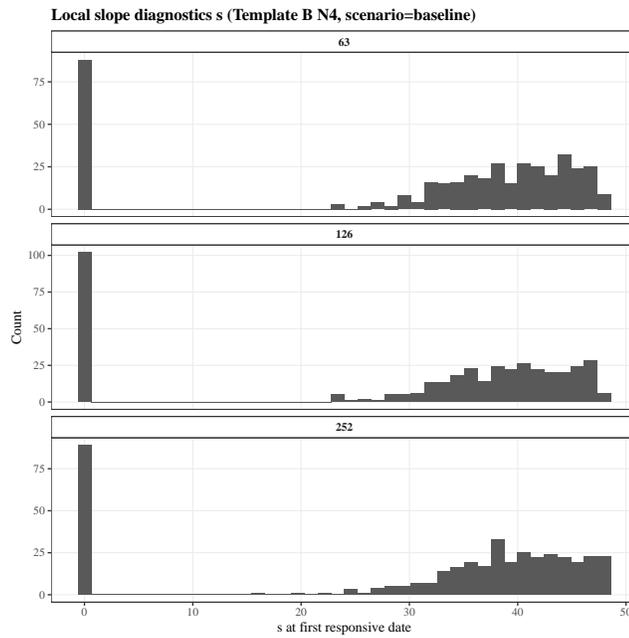
### H.3 Local slope diagnostics



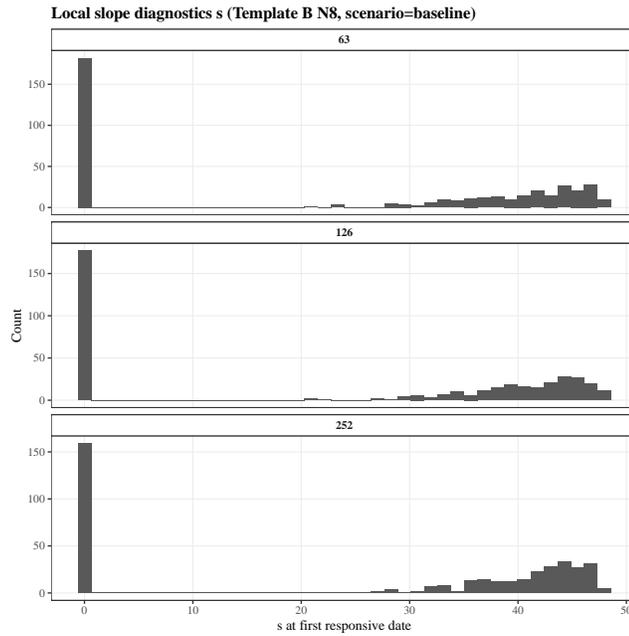
**Figure 10** Histogram of the local slope diagnostic  $s$  at the first responsive date for Template A (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .



**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  $T \in \{63, 126, 252\}$ .

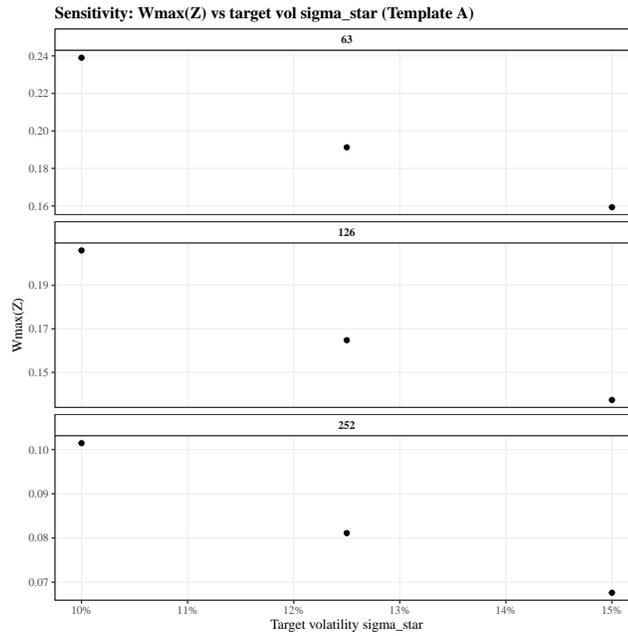


**Figure 12** Histogram of the local slope diagnostic  $s$  at the first responsive date for Template B with  $N = 4$  (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .

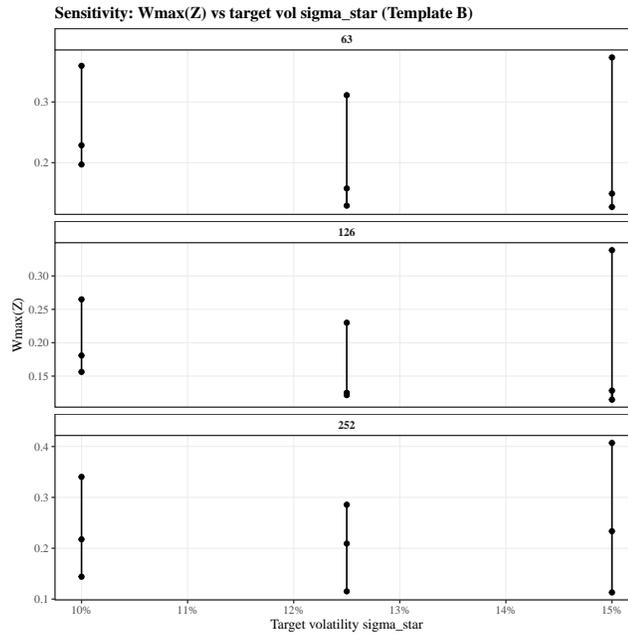


**Figure 13** Histogram of the local slope diagnostic  $s$  at the first responsive date for Template B with  $N = 8$  (baseline scenario), shown separately for  $T \in \{63, 126, 252\}$ .

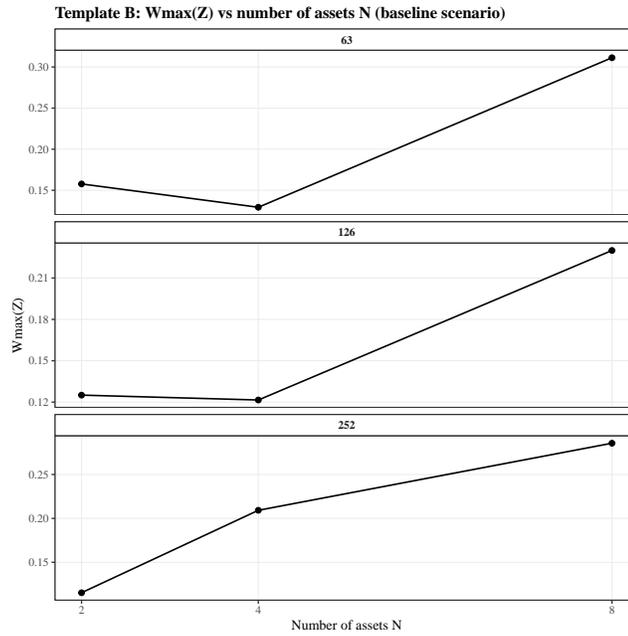
## H.4 Sensitivity summaries



**Figure 14** Sensitivity of  $W_{\max}(\mathcal{Z})$  to the target volatility  $\sigma^*$  for Template A, shown across horizons  $T \in \{63, 126, 252\}$ .



**Figure 15** Sensitivity of  $W_{\max}(\mathcal{Z})$  to the target volatility  $\sigma^*$  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  $(N, \sigma^*, T)$ , the table reports the number of stress states with regime  $\text{cap}$  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[fail   cap] at $W_{\text{med}}$
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%