An Impossibility Theorem for Price-Based Risk Constraints

Francesco Nicolai*
BI Norwegian Business School

March 2026 – Latest Version
July 2024 – First Version

Abstract

Price-based risk constraints map sampled transaction prices into measured risk and then into binding requirements. We prove a constructive local impossibility theorem: at reachable binding states with sufficiently strong amplification, no such rule can simultaneously satisfy three desirable properties. First, risk sensitivity: posted requirements increase at first order with the measured risk statistic. Second, liquidity continuity: prices and positions respond continuously to small shocks. Third, manipulation-proofness: no admissible finite-horizon round trip yields strictly positive expected profit. Unlike classic manipulation, the mechanism does not require large trades, large price moves, or making a slack constraint bind. A small trade can affect the prices used by the rule, tighten requirements, induce predictable forced selling, and then be reversed profitably. The result holds even with continuous trading between discrete measurement and reset dates, and it survives strategic responses by constrained investors. Manipulability is strongest when realized risk is low.

Keywords: price-based risk constraints; collateral design; central clearing; forced deleveraging; manipulation-proofness; portfolio margining; endogenous risk measurement.

JEL Classification: G21, G28, G12, D82, D62.

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

1 Introduction

In many financial markets, the same transaction prices that clear today's trades also help determine tomorrow's binding constraints. Margins, haircuts, internal value-at-risk (VaR) or expected-shortfall (ES) limits, leverage targets, and rebalancing rules are often reset from sample statistics computed from recent transaction prices. Once that is true, trading that moves today's price changes tomorrow's required collateral, permitted leverage, and forced rebalancing mechanically. The rule becomes therefore part of the trading environment itself.

Our main result is a constructive local impossibility theorem for such price-based rules. No designer can, in general, guarantee all three of the following properties at once: (i) local risk sensitivity, so requirements rise at first order with measured risk; (ii) liquidity continuity, so prices and aggregate positions respond continuously to small shocks; and (iii) round-trip manipulation-proofness, so no admissible finite-horizon round trip earns strictly positive expected profit by moving the prices that enter the next update. If (i) and (ii) hold, then sufficiently strong amplification implies failure of (iii) in reachable binding states (Theorem 1). If the rule reaches such a binding high-amplification state with positive probability, no design can preserve all three properties.

The result matters because each of the three properties is a standard design objective. Risk sensitivity is the basic rationale for adaptive margins and risk limits: when measured risk rises, required protection should rise as well (Biais et al., 2016; Wang et al., 2022). Liquidity continuity matters because a well-designed rule should not turn small shocks into abrupt deleveraging, liquidity dry-ups, or fire sales when constraints bind and prices are impact-sensitive (Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 2011). Manipulation-proofness is equally fundamental: a trader should not be able to profit from a round trip whose only purpose is to move prices and thereby change the next update (Allen and Gale, 1992; Huberman and Stanzl, 2004). Our theorem shows that these objectives generally cannot be achieved simultaneously. The design problem is therefore not whether all three goals can be attained together, but how the rule should trade them off. The result applies whenever a binding feasibility constraint is updated mechanically from past transaction prices in a market that is not infinitely liquid.[1]

Several points clarify the scope of the theorem and rule out incorrect interpretations. First, the result is local, but the design implication is global. It is therefore incorrect to read the theorem as a statement about a narrow or economically irrelevant corner. The relevant question is whether the rule can reach a binding state in which a marginal tightening forces deleveraging and that deleveraging moves prices. If it can, then no price-based rule can satisfy risk sensitivity, liquidity continuity, and manipulation-proofness at that state, and so no such rule can guarantee all three properties over its reachable state space. The triggering trade need not be large. An arbitrarily small trade can move the sampled price used in the next update, tighten tomorrow's requirement, induce forced sales, and then be reversed into the resulting price pressure (Theorem 1). Sections 3.6 and 6.4 show that such states are common.[2]

[1]Liquidity continuity is not redundant. It rules out designs that avoid the local slope argument only by creating liquidity cliffs. The point is therefore not just that risk sensitivity can conflict with manipulation-proofness. It is that joint feasibility already fails within the relevant class of rules that also preserve smooth market functioning.

[2]Similarly, Nicolai and Risteska (2026), in a different setting with volatility-managed portfolios, show that the same mechanism becomes relevant once linked notional is moderate relative to market depth. The mechanism is therefore not

Second, it is incorrect to read the mechanism as relying on constrained investors being myopic. They may be fully forward-looking and understand both the trader's incentive and the effect of today's price on tomorrow's requirement. In that case they adjust in advance, which can reduce the equilibrium amount of forced selling.[3] Section 4.3 provides a microfoundation for this case. But anticipation does not remove the mechanism. As long as the requirement still binds, a price move that tightens tomorrow's rule still changes constrained demand and hence feeds the sampled price into subsequent equilibrium outcomes. Forward-looking behavior can attenuate this feedback, but it cannot make risk sensitivity, liquidity continuity, and manipulation-proofness jointly feasible. The tension disappears only in limiting cases in which the constraint ceases to matter: either constrained investors withdraw so much in stress that further tightening no longer forces meaningful sales, or the rule is kept slack precisely in the states in which a risk-sensitive rule would otherwise tighten. In either case, the mechanism vanishes only because the rule is no longer binding, or no longer meaningfully risk-sensitive, where it matters most.

Third, continuous trading between reset dates does not remove the mechanism. Section 4.1 shows that the same logic survives when trading is continuous between resets. The problem comes from the reset itself, not from a lack of trading between resets. If tomorrow's binding requirement is computed from prices sampled at a reset time or over a reset window, then moving those sampled prices changes tomorrow's requirement at first order. That tighter requirement generates predictable post-reset balance-sheet pressure and order flow.

Fourth, the mechanism is not standard manipulation. The trader does not profit by exploiting mispricing, deceiving other traders, or relying on an inefficient price process. The profit opportunity is created by the rule itself. When sampled transaction prices mechanically enter next-period requirements, moving those prices changes tomorrow's binding constraint. If the constraint binds, that change induces predictable post-update forced flow and price pressure, even when the unaffected price is a martingale and the underlying impact model is manipulation-free absent rule-induced feedback in the sense of Huberman and Stanzl (2004).[4]

A further implication is that the most vulnerable states may be the ones that look safest. Even when the posted requirement is written as a VaR or ES rule, it is usually driven by an estimated volatility computed from recent returns. Under standard rolling-volatility constructions, a given distortion of the latest sampled price has a larger effect on the estimated risk input when recent realized volatility is low. After a quiet period, a small trigger trade can therefore move tomorrow's requirement the most. Tranquil states can thus be especially exposed to trigger-and-reverse manipulation, a rule-based analogue of the "paradox of financial instability" (Borio and Drehmann, 2009). This is different from the leverage-cycle mechanism (Adrian and Shin, 2010, 2014; Geanakoplos, 2010; Brunnermeier and Sannikov, 2014). There, low measured risk increases fragility by relaxing constraints

specific to margin rules, but a general feature of price-based risk rebalancing rules.

[3]Consistent with this equilibrium-response view, transaction-level evidence from cleared repo shows that higher CCP haircuts relative to OTC induce venue substitution and adverse selection, with safer borrowers shifting away from the CCP and CCP borrower composition deteriorating (Chebotarev, 2025).

[4]In this spirit, Nicolai and Risteska (2026) construct an explicit example in which the unaffected price is a martingale and the impact specification is dynamically manipulation-free for unconstrained trading in the sense of Huberman and Stanzl (2004). Profits arise only once a disclosed price-based rule maps sampled transaction prices into next-period requirements or mandated positions, creating predictable post-update forced flow.

and encouraging leverage. Here, even without prior balance-sheet expansion, low realized volatility makes the rule itself more sensitive to a single sampled price.

Theorem 1 is proved by an explicit trigger-and-reverse strategy. Profitability is governed by a simple amplification chain: how strongly the rule maps sampled prices into requirements, how strongly tighter requirements force deleveraging, and how strongly that deleveraging moves prices. Each term is, in principle, measurable from margin methodology, liquidation behavior, and market impact. The same logic also yields an upper bound on short-horizon pass-through from price-based inputs into posted requirements if the rule is to be robust to trigger-and-reverse manipulation. Section 6 studies the designer's problem under this restriction. The central implication is a tension between coverage and implementability: the more sharply margin responds to measured risk, the easier the rule is to exploit in binding states. This provides a structural rationale for limited pass-through, smoothing, and related anti-procyclicality tools.

We then extend the argument to portfolio margining with cross-impact. Portfolio systems map a vector of transaction prices into a single charge, but a binding margin call is met through a vector of asset sales that need not fall on the contracts that most effectively move the portfolio risk measure. This creates a wedge between the contracts used to trigger the tighter requirement and the contracts sold after the requirement binds. A trader can therefore move the portfolio risk measure with one set of instruments while taking positions in the instruments that will be sold when the tighter call forces deleveraging. Incentives then spill across contracts. We show that the strength of this channel is governed by an interpretable alignment term between the direction where risk increases and the direction of liquidation-induced price pressure under cross-impact. When that alignment is strong enough, profitable cross-asset round trips exist for arbitrarily small admissible triggers. The contracts that absorb trading pressure during the risk-sampling window need not be the ones that face post-update illiquidity and forced sales, so portfolio margining can shift stress across contracts.

1.1 Contribution and related literature

This paper brings together three literatures that are usually kept separate. The first studies forced sales, liquidity spirals, endogenous leverage, and, in some cases, margins or haircuts that respond to volatility or liquidity. It generally does not analyze a rule under which a short window of transaction prices mechanically determines the next posted requirement and thereby gives a trader first-order influence over that update (Kiyotaki and Moore, 1997; Geanakoplos, 2010; Brunnermeier and Pedersen, 2009; Brunnermeier and Sannikov, 2014; He and Krishnamurthy, 2013; Gromb and Vayanos, 2002; Gârleanu and Pedersen, 2011). The second studies procyclicality and risk regulation, often taking the risk input as given or statistically specified, for example through volatility dynamics. That literature emphasizes feedback from constraints into prices, but it does not isolate the incentive problem that arises when the risk input is estimated from transaction prices that traders can influence (Basak and Shapiro, 2001; Danielsson et al., 2004; Adrian and Shin, 2014; Glasserman and Wu, 2018). The third develops no-manipulation and no-dynamic-arbitrage results under market impact and exogenous constraints, where sufficiently small round trips lose money to impact and cannot earn strictly positive expected profits (Allen and Gale, 1992; Jarrow, 1992; Huberman and Stanzl, 2004; Gatheral, 2010). Our contribution is to show what changes once next-period feasibility depends on a price-based update. In binding states, an arbitrarily small distortion of the transaction prices used by the rule shifts the next requirement at first order and creates predictable rule-induced order flow after the update.

Our contribution is to make the update rule itself the object of analysis and to characterize the relevant design frontier. First, we prove a constructive local impossibility result. Near reachable binding states, no price-based, risk-sensitive update rule can simultaneously deliver risk sensitivity, liquidity continuity, and round-trip manipulation-proofness (Theorem 1). Second, the profitability of the deviation reduces to a single amplification chain, which yields a sharp diagnostic for vulnerability and, read as a design restriction, an upper bound on short-horizon pass-through from price-based inputs into posted requirements (Section 3). Third, under portfolio margining with cross-impact, the contracts that move the portfolio risk measure need not be the ones sold when tighter requirements force deleveraging. That mismatch creates cross-contract manipulation incentives and spillovers (Section 5). Fourth, Section 6 studies the implied design problem directly. It distinguishes the unconstrained first-best risk-management target from the implementable rule, derives a manipulation-proof slope cap, and shows that the cap can be quantitatively tight under benchmark calibrations. Discontinuities do not remove the incentive problem; they merely shift it to deviations that cross thresholds. Pooling and bounded adjustment regions can arise as consequences of the cap.

This paper establishes a general impossibility result for price-based constraints. It identifies the general incentive restriction created when binding feasibility depends on tradable prices in markets with price impact. Nicolai and Risteska (2026) apply the same logic in a different setting. Taking a deterministic update rule g as given, they turn the amplification condition into a viability test and a liquidity-scaled capacity bound for mechanically rebalanced, price-insensitive capital, such as volatility-managed portfolios. Our theorem therefore isolates the implementability constraint, while Nicolai and Risteska (2026) show how the same diagnostic can be used to quantify admissible scale in other environments. Examples outside CCPs include internal VaR, ES, and leverage-limit policies tied to recent prices; internal capital-allocation rules in multi-strategy hedge funds that reallocate balance-sheet capacity across strategies in response to measured risk; repo and prime-brokerage haircuts linked to volatility or liquidity measures; volatility-control indices that trigger mechanical reallocations; and risk-parity or volatility-targeting mandates that scale exposure with estimated risk. More generally, the same logic applies whenever exposure is a deterministic function of recent prices, realized volatility, drawdowns, or other price-based statistics.

The design implications speak directly to clearing and margin policy. Central clearing manages counterparty risk through margin and default resources, while netting sets and clearing architecture shape exposures and collateral demand (Duffie and Zhu, 2011; Menkveld and Vuillemey, 2021; Cont and Kokholm, 2014; Loon and Zhong, 2014; Duffie et al., 2015). Evidence from repo markets shows that haircut policy affects participation and composition: bilateral haircuts vary across counterparties and relationships, while uniform CCP haircuts can push activity toward OTC markets and worsen selection when they rise relative to OTC (Julliard et al., 2024; Chebotarev, 2025). Policy work accordingly emphasizes both coverage and stability of initial margin and documents anti-procyclicality tools,[5] while empirical and policy evidence shows that margin calls create large, state-dependent

[5] See Committee on Payments and Market Infrastructures and International Organization of Securities Commissions

liquidity demands in stress (Murphy et al., 2014, 2016; Aldasoro et al., 2023; King et al., 2023). Our mechanism adds an incentive constraint to this discussion. Once the risk input is computed from tradable prices and the requirement binds, short-horizon pass-through becomes a design choice with an incentive-based upper bound. That gives a structural rationale for slope limits, smoothing, and buffer-type overlays. The same perspective connects to work on robustness and incentives in risk measurement and margin setting, and to analyses of CCP governance and externalities in stress (Cont et al., 2010; Huang and Takáts, 2024; Cont and Ghamami, 2025; Pirrong, 2011, 2014). Kuong and Maurin (2024) provide a particularly clean contracting benchmark by endogenizing key layers of the default waterfall through a model of the CCP as an agent. Our contribution is complementary. Taking the broader collateral and default-resource architecture as given, we isolate an additional implementability constraint that arises when the posted requirement is computed from tradable prices and binds.

Roadmap. Section 2 presents a canonical cleared-market model with $M_{t+1} = g(\Gamma_t)$ to fix ideas, although the analysis applies more generally to any price-based constraint that maps sampled transaction prices into a binding requirement. For transparency, the main text works with linear impact, realized variance, and an explicit margin-elastic demand rule; Appendix A proves the corresponding results in a more general setting. Section 3 formalizes the three desiderata, constructs the trigger-and-reverse strategy, and proves the impossibility theorem. It then rewrites the profit bound in terms of the amplification chain to obtain the slope cap and a magnitude limit. Section 4 studies robustness and implementation, including continuous trading between discrete update times, alternative risk inputs such as parametric VaR and ES and scenario-based maxima and quantiles, and implementation features such as trimming, thresholds, rounding, floors, and stepwise add-ons. Section 5 extends the analysis to portfolio margining with cross-impact and forced liquidation across contracts, and derives the corresponding cross-contract impossibility result. Section 6 studies the design problem implied by the theorem, distinguishes the unconstrained risk-management target from the implementable rule, derives the manipulation-proof slope cap, and gives a quantitative interpretation of how restrictive that cap is.

2 Model

We now formalize the environment described in the introduction. A rule designer, for example a CCP, computes a risk statistic from transaction prices and updates initial margin mechanically. Near reachable binding states, a marginal increase in M_{t+1} induces forced liquidation with a nonzero local slope. Because transaction prices also enter the next update, this mapping can generate predictable liquidation-driven price pressure that a finite-horizon trigger-and-reverse strategy can exploit. The analysis does not rely on large trades, large price moves, or on making a slack constraint bind. For transparency, the main text works with three special cases: linear price impact in net order flow, realized variance over a fixed sampling window, and a target-with-cap liquidation rule that makes the

(2012); Basel Committee on Banking Supervision (2022); Basel Committee on Banking Supervision and Committee on Payments and Market Infrastructures and Board of the International Organization of Securities Commissions (2022); European Union (2013); European Securities and Markets Authority (2018); European Systemic Risk Board (2020); Bank of England (2021); European Central Bank (2023).

position elasticity with respect to margin explicit. These choices are only for exposition. Appendix A replaces realized variance with a generic price-based risk functional, allows a general local impact map, and permits broader liquidation responses.

2.1 Primitives and timing

Time is discrete, $t = 0, 1, 2, \ldots$. A single contract is traded and marked to market each period. Fundamentals $V_t \in \mathbb{R}$ are exogenous, adapted to the public filtration $\{\mathcal{F}_t\}$, and satisfy

$$\mathbb{E}[V_{t+1} \mid \mathcal{F}_t] = V_t. \tag{1}$$

Assumption (1) removes predictable drift from the strategic trader's expected profit. Fix $m \geq 2$. Let P_t be the period-t transaction price and M_t the per-unit initial margin applied in period t. The public state at the start of t is $S_t = \big(V_t,\ P_{t-m}, \ldots, P_{t-1},\ M_t\big)$. A continuum of constrained traders $i \in [0, 1]$ choose end-of-period positions $x_t(i)$ subject to M_t. A strategic trader chooses $a_t \in [-\bar{a}, \bar{a}]$, $\bar{a} > 0$, and his inventory evolves as $y_t = y_{t-1} + a_t$ with $y_{-1} = 0$. Competitive liquidity suppliers clear net order flow and P_t is realized by the price formation rule specified below. The rule designer computes a price-based input from the realized window and posts next-period margin,

$$\Gamma_t = \Gamma(P_{t-m+1}, \ldots, P_t), \qquad M_{t+1} = g(\Gamma_t).$$

At $t+1$, M_{t+1} applies. Breaching accounts adjust immediately through forced liquidation, and this order flow enters the same price formation rule in period $t+1$. Constrained traders choose $x_t(i)$ as a function of S_t and do not condition on P_t, which is realized only after orders clear. They may condition on P_t from $t+1$ onward since $P_t \in S_{t+1}$.

2.2 Price formation

Let $X_t = \int_0^1 x_t(i)\, di$ be the aggregate constrained position at the end of period t, and let $\Delta X_t = X_t - X_{t-1}$ be constrained traders' net order flow in period t. Total net order flow is

$$Q_t = \Delta X_t + a_t.$$

Transaction prices follow linear price impact with slope $\alpha > 0$:

$$P_t = V_t + \alpha Q_t = V_t + \alpha(\Delta X_t + a_t). \tag{2}$$

Thus α is the local price response to order flow, including forced-liquidation flow.[6]

[6] Appendix A.4 allows a general impact; for the main results only the local sensitivity of P_t to Q_t around the benchmark path matters.

2.3 Margin update rule and price-based risk input

The rule designer sets next-period per-unit initial margin by

$$M_{t+1} = g(\Gamma_t), \tag{3}$$

where $g : \mathbb{R}_+ \to (0, \infty)$ is public and increasing. In the baseline model the risk input is realized variance computed from transaction prices over the m-period window:

$$\Gamma_t = RV_t = \sum_{j=0}^{m-1} (P_{t-j} - P_{t-j-1})^2 . \tag{4}$$

Since transaction prices satisfy (2), trading within the window moves the sampled marks and therefore shifts Γ_t.[7]

2.4 Constrained traders

A continuum of traders $i \in [0, 1]$ face initial margin. Trader i has equity $E(i) > 0$ and must satisfy

$$|x_t(i)|\, M_t \le E(i). \tag{5}$$

When M_{t+1} is posted, any trader who violates (5) must reduce exposure at $t + 1$ until the constraint holds; this adjustment is forced liquidation. For transparency, the main text uses a target-with-cap rule. Let $\bar{x}_t \in \mathbb{R}$ be a public target chosen from the pre-trade state S_t (Section 2.1), before the strategic trade and before P_t is realized. The realized position is the target clipped by the margin cap:

$$x_t(i) = \pi(\bar{x}_t, E(i)/M_t) = \operatorname{sgn}(\bar{x}_t) \min\left\{ |\bar{x}_t|, \frac{E(i)}{M_t} \right\}. \tag{6}$$

Aggregate constrained exposure is

$$X_t = X(M_t, \bar{x}_t) = \int_0^1 \operatorname{sgn}(\bar{x}_t) \min\left\{ |\bar{x}_t|, \frac{E(i)}{M_t} \right\} di. \tag{7}$$

When M_{t+1} rises, caps $E(i)/M_{t+1}$ tighten; if targets remain high, X_{t+1} falls and the reduction is liquidation order flow at $t + 1$. The analysis uses only the state-dependent slope of forced flow with respect to the next posted margin in states where the constraint binds. Define

$$c_X(S_t) \;=\; - \left. \frac{\partial X(M, \bar{x}_{t+1})}{\partial M} \right|_{M = M_{t+1}}, \tag{8}$$

where the derivative is evaluated at the benchmark path and holds fixed the constrained sector's period-$t+1$ targets chosen from S_{t+1} before enforcement. Under the target-with-cap rule, $c_X(S_t) > 0$ whenever a positive mass of traders is margin-binding. Appendix B replaces the target rule by a

[7] Section 4.2 and Appendix A.3 allow Γ_t to be any price-based functional of the sampled prices. The proofs use only the local directional sensitivity of Γ_t to perturbations of the sampled prices.

forward-looking constrained sector that anticipates the update and chooses positions optimally. In that equilibrium the constrained block enters the mechanism only through the local liquidation slope

$$c_X^{\text{eff}}(S_t) = -\left.\frac{\partial X^*(M;S_t)}{\partial M}\right|_{M=M_{t+1}}.$$

Once positions are endogenous, the key question is whether the requirement is locally binding in the states of interest. That is exactly what the effective slope $c_X^{\text{eff}}(S_t)$ captures: the local response of aggregate forced adjustment to M_{t+1} after allowing for anticipatory repositioning or within-window re-optimization. Those are the relevant states for design. When the requirement binds, the update mechanically generates forced trading and price pressure, so implementability has to be evaluated there rather than in slack states. Anticipation can reduce this elasticity, but it does not alter the logic of the theorem. The amplification condition and the associated slope restriction are both written in terms of $c_X^{\text{eff}}(S_t)$, so they apply whenever $c_X^{\text{eff}}(S_t) > 0$. Under standard concavity assumptions, $c_X^{\text{eff}}(S_t) = 0$ only when no participating trader is locally constrained by M_{t+1}, that is, when the requirement is locally slack for the relevant set, or participation is negligible, in that state.[8]

2.5 Strategic trader and round trips

A strategic trader (or manipulator) chooses a predictable trade sequence $a_0, \ldots, a_T$ over a fixed horizon T, where $a_t > 0$ corresponds to a buy and $a_t < 0$ a sell. A round trip is a trade sequence with net-zero inventory at the end of the horizon:

$$\sum_{t=0}^{T} a_t = 0. \tag{9}$$

The resulting inventory process is $y_t = \sum_{s=0}^{t} a_s$. Let P_t denote the transaction price at time t. The trader's trading profit is

$$\Pi_T(a_{0:T}) = -\sum_{t=0}^{T} a_t P_t. \tag{10}$$

We allow an optional quadratic inventory penalty $\kappa \geq 0$. The resulting inventory-adjusted profit is

$$\Pi_T^{\kappa}(a_{0:T}) = \Pi_T(a_{0:T}) - \kappa \sum_{t=0}^{T-1} y_t^2. \tag{11}$$

Finally, we optionally impose a hard funding constraint: there exists $\bar{F} > 0$ such that the trader's margin outlay never exceeds $\bar{F}$:

$$M_t|y_t| \leq \bar{F} \qquad \text{for all } t. \tag{12}$$

[8] It is difficult to treat endogenous adjustment as a general resolution of the mechanism in the theorem. The empirical evidence overwhelmingly shows that when collateral or funding constraints bind, the result is not frictionless re-optimization that neutralizes forced flow, but sizable, state-dependent liquidity demands, forced sales, fire-sale externalities, and price spillovers. See, among others, Murphy et al. (2014, 2016); Aldasoro et al. (2023); King et al. (2023) on margin-call-induced liquidity demands, and Coval and Stafford (2007); Duarte and Eisenbach (2013); Greenwood et al. (2015) on forced sales and spillovers in binding states.

This constraint only restricts the maximum admissible trigger size and plays no role in the local gain mechanism.[9]

2.6 Information and equilibrium

A constrained trader i follows a Markov policy $\sigma_i^C : \mathcal{S} \to \mathbb{R}$, where the public state space is $\mathcal{S} = \mathbb{R} \times \mathbb{R}^m \times (0, \infty)$. The policy maps the current public state into an end-of-period position that satisfies (5). Fix a benchmark policy profile $\{\sigma_i^C\}_{i\in[0,1]}$. The strategic trader enters only as a deviator. Along the benchmark path he does not trade. Starting from a benchmark state S_t, we ask whether there exists an admissible finite-horizon round trip that earns strictly positive conditional expected profit through its effect on the next margin update, taking primitives, constrained-trader policies, and the margin rule as fixed. One such profitable deviation is enough to violate sequential optimality and therefore to reject manipulation-proofness at that state as in Definition 3.

We use $(\cdot)^{(0)}$ for benchmark, or no-deviation, objects. Fix S_t. Let $P_t^{(0)}$ denote the period-t transaction price under the benchmark and define the last return by $R_t^{(0)} = P_t^{(0)} - P_{t-1}$. Let $RV_t^{(0)}$ be realized variance computed from the sampled prices $(P_{t-m}, \ldots, P_{t-1}, P_t^{(0)})$, and let $M_{t+1}^{(0)} = g(RV_t^{(0)})$ be the posted margin. Because $\{\sigma_i^C\}$ is Markov in S_t and price formation follows (2), the objects $(P_t^{(0)}, R_t^{(0)}, RV_t^{(0)}, M_{t+1}^{(0)})$ are measurable functions of S_t. A state is reachable if it arises along the benchmark path induced by $\{\sigma_i^C\}$, g, and (2). All local assumptions are imposed only on reachable states.

Definition 1 (Markov equilibrium under the margin rule)**.** Fix g and the risk statistic (4). A Markov equilibrium is a pair $(\{\sigma_i^C\}_{i\in[0,1]}, P)$ such that, for every t, (i) constrained traders choose $x_t(i) = \sigma_i^C(S_t)$ and $X_t = \int_0^1 \sigma_i^C(S_t)\, di$; (ii) the transaction price satisfies (2) with $Q_t = \Delta X_t$ along the equilibrium; (iii) the margin designer sets $M_{t+1}^{(0)} = g(RV_t^{(0)})$, where $RV_t^{(0)}$ is computed from realized transaction prices.

3 Properties and main results

3.1 Three desirable properties

Theorem 1 is a non-joint-feasibility result: if risk sensitivity and liquidity continuity hold, we construct one admissible finite-horizon round trip with strictly positive conditional expected profit, so round-trip manipulation-proofness fails. We want to stress that liquidity continuity is a market-functioning desideratum; it is not the property contradicted in the proof.

Definition 2 (Risk sensitivity)**.** A margin rule $M_{t+1} = g(\Gamma_t)$ is risk-sensitive on an interval $I \subset \mathbb{R}_+$ if g is differentiable on I and there exists $c_g > 0$ such that $g'(r) \geq c_g$ for all $r \in I$.

Risk sensitivity requires nontrivial local pass-through from the measured risk input into posted margin, as in risk-based margining and incentive-alignment arguments (Biais et al., 2016).

[9] Appendix A.5 adds proportional execution and linear funding wedges; these only add a minimum quantity feasibility restriction.

Definition 3 (Manipulation-proofness)**.** A mechanism is round-trip manipulation-proof at a state S_t if for every admissible round trip $a_{t:t+T}$,

$$\mathbb{E}[\Pi_T^\kappa(a_{t:t+T}) \mid S_t] \leq 0.$$

This is the standard no-manipulation requirement in price-impact settings: profitable round trips correspond to manipulation/dynamic arbitrage and are ruled out by canonical no-manipulation conditions on price impact (Huberman and Stanzl, 2004).

Definition 4 (Liquidity continuity)**.** Fix a Markov equilibrium (Definition 1). The mechanism has liquidity continuity if the mapping $S_t \mapsto (P_t^{(0)}, X_t^{(0)})$ is continuous on the set of reachable states and, for each reachable S_t, the map

$$\varepsilon \mapsto \big(P_t(\varepsilon), X_t(\varepsilon), P_{t+1}(\varepsilon), X_{t+1}(\varepsilon)\big)$$

is continuous at $\varepsilon = 0$, where $P_s(\varepsilon)$ and $X_s(\varepsilon)$ are the outcomes when the strategic trader submits an additional trade ε at time t (relative to the benchmark with zero strategic trade) and all primitives are held fixed.

Liquidity continuity rules out endogenous liquidity cliffs: arbitrarily small within-window perturbations cannot trigger jumps in prices or aggregate positions. It is a continuity requirement, not a bound on steepness; discontinuities typically come from kinks, thresholds, or stepwise overlays in the composite mapping from sampled marks into posted margin (Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 2011).

3.2 Assumptions

Fix a reachable benchmark state S_t and hold primitives and other agents' policies fixed. We construct a local two-step deviation. A trade at t changes the last sampled price P_t, shifts the risk input, here RV_t, and therefore changes the posted margin M_{t+1}. When the requirement binds, the higher margin forces liquidation at $t+1$ and, through price impact, moves P_{t+1}. Reversing the initial trade can then profit from this predictable liquidation-driven price pressure. The proof bounds each step in this local transmission chain. Appendix A states the corresponding conditions for a generic price-based input and for non-smooth implementation features using one-sided slopes and directional derivatives. Nicolai and Risteska (2026) characterize multi-period trigger-and-reverse strategies for fixed deterministic update rules; here it is enough to show that one such profitable deviation exists.

Assumption 1 (Local slope of the margin rule)**.** There exists an interval $I = [\underline{r}, \overline{r}] \subset \mathbb{R}_+$ such that g is differentiable on I and

$$g'(r) \in [c_g, C_g] \quad \text{for all } r \in I,$$

with $0 < c_g \leq C_g < \infty$.

Assumption 2 (Local sensitivity of the risk input)**.** At S_t, $RV_t^{(0)} \in \operatorname{int}(I)$ and $|R_t^{(0)}| \geq \underline{R}$ for some $\underline{R} > 0$. The strategic trader can trade in either direction with capacity $\bar{a} > 0$. In the explicit construction we take $R_t^{(0)} < 0$; the case $R_t^{(0)} > 0$ is obtained by reversing signs.

Assumption 3 (Local liquidation slope under binding margin)**.** At S_t with $RV_t^{(0)} \in \mathrm{int}(I)$, there exist $c_X > 0$ and $\delta_X > 0$ such that, holding the period-t target $\bar{x}_t$ fixed, the mapping $M \mapsto X(M, \bar{x}_t)$ is continuously differentiable on $[M_t, M_t + \delta_X]$ and satisfies

$$-\frac{\partial X}{\partial M}(M, \bar{x}_t) \geq c_X \quad \text{for all } M \in [M_t, M_t + \delta_X].$$

Assumption 3 is the statement that when a positive mass of traders is margin-binding, a higher posted margin forces a first-order reduction in aggregate exposure. Under the target-with-cap policy it holds in any such binding state. Appendix B endogenizes targets in a forward-looking equilibrium; in that case the same condition holds with c_X replaced by the equilibrium slope $c_X^{\mathrm{eff}}(S_t)$, the derivative of equilibrium aggregate exposure with respect to the next posted margin.

Assumption 4 (Price impact)**.** $\alpha > 0$.

3.3 Auxiliary lemmas

We isolate four links used in the two-step deviation: (i) how RV_t moves when the last sampled mark P_t moves, (ii) how g maps this into M_{t+1}, (iii) how a higher margin forces liquidation when constraints bind, and (iv) how that forced flow moves P_{t+1} under linear impact.

Lemma 1 (Sensitivity of realized variance to the last price)**.** *Fix $(P_{t-m}, \ldots, P_{t-1})$ and a benchmark last price $P_t^{(0)}$. Let $R_t^{(0)} = P_t^{(0)} - P_{t-1}$. The map $p \mapsto RV_t(P_{t-m}, \ldots, P_{t-1}, p)$ in (4) is differentiable and*

$$\left.\frac{\partial RV_t}{\partial P_t}\right|_{P_t = P_t^{(0)}} = 2R_t^{(0)}.$$

If $|R_t^{(0)}| \geq \underline{R} > 0$, then for any ΔP with $\mathrm{sgn}(\Delta P) = \mathrm{sgn}(R_t^{(0)})$,

$$RV_t(P_t^{(0)} + \Delta P) - RV_t(P_t^{(0)}) \geq 2\underline{R}\,|\Delta P|.$$

Proof. Holding $(P_{t-m}, \ldots, P_{t-1})$ fixed, only $(P_t - P_{t-1})^2$ depends on P_t, so

$$\frac{\partial RV_t}{\partial P_t} = 2(P_t - P_{t-1}),$$

and evaluating at $P_t^{(0)}$ gives $2R_t^{(0)}$. Also,

$$RV_t(P_t^{(0)} + \Delta P) - RV_t(P_t^{(0)}) = (R_t^{(0)} + \Delta P)^2 - (R_t^{(0)})^2 = 2R_t^{(0)}\Delta P + (\Delta P)^2.$$

If $\mathrm{sgn}(\Delta P) = \mathrm{sgn}(R_t^{(0)})$, then $R_t^{(0)}\Delta P = |R_t^{(0)}||\Delta P|$ and $(\Delta P)^2 \geq 0$, so the difference is at least $2|R_t^{(0)}||\Delta P| \geq 2\underline{R}|\Delta P|$. □

Lemma 2 (Pass-through from RV_t to margin)**.** *Under Assumption 1, for any $RV, RV' \in I$ with $RV' \geq RV$,*

$$g(RV') - g(RV) \geq c_g\,(RV' - RV).$$

Proof. By the mean value theorem, $g(RV') - g(RV) = g'(\xi)(RV' - RV)$ for some ξ between RV and RV'. Assumption 1 gives $g'(\xi) \geq c_g$. □

Lemma 3 (Liquidation induced by a margin increase)**.** *Under Assumption 3, there exists $\delta_M > 0$ such that for all $\Delta M \in (0, \delta_M)$,*

$$L(\Delta M) = X(M_t, \bar{x}_t) - X(M_t + \Delta M, \bar{x}_t) \geq c_X \, \Delta M,$$

and $\Delta M \mapsto L(\Delta M)$ is continuous on $(0, \delta_M)$.

Proof. Write $X(M) = X(M, \bar{x}_t)$ and set $\delta_M = \delta_X$. For $\Delta M \in (0, \delta_M)$, the mean value theorem gives $X(M_t + \Delta M) - X(M_t) = X'(\xi)\Delta M$ for some $\xi \in (M_t, M_t + \Delta M)$. Assumption 3 implies $X'(\xi) \leq -c_X$, so $L(\Delta M) \geq c_X \Delta M$. Continuity follows from continuity of X on $[M_t, M_t + \delta_M]$. □

Lemma 4 (Liquidation-induced price change under linear impact)**.** *Fix $t+1$. Suppose the margin update induces additional forced liquidation $L > 0$ at $t+1$ relative to the benchmark. Let $\Delta X_{t+1}^{(0)}$ be benchmark constrained order flow at $t+1$, so realized constrained order flow is $\Delta X_{t+1} = \Delta X_{t+1}^{(0)} - L$. If the strategic trader buys $q \in [0, \bar{a}]$ at $t+1$, then*

$$P_{t+1} = V_{t+1} + \alpha(\Delta X_{t+1}^{(0)} - L + q).$$

Holding $\Delta X_{t+1}^{(0)}$ fixed, the liquidation component is $-\alpha L$.

Proof. Substitute $Q_{t+1} = \Delta X_{t+1} + q = \Delta X_{t+1}^{(0)} - L + q$ into (2). □

3.4 An explicit round trip

This section provides a transparent two-period deviation that exposes the feedback loop. Fix a time t and consider a trader who sells $q > 0$ at t and buys q at $t+1$. The trade sequence is $(a_t, a_{t+1}) = (-q, q)$, and the end-of-horizon inventory is zero. For the realized-variance specification (4), a sale at t changes the time-t return from $R_t^{(0)}$ to $R_t = R_t^{(0)} - \alpha q$ and hence changes realized variance by

$$\Delta RV_t = RV_t - RV_t^{(0)} = 2R_t^{(0)}(-\alpha q) + (\alpha q)^2. \tag{13}$$

When $R_t^{(0)} < 0$, a sale increases realized variance, which increases margin when g is risk-sensitive. To simplify notation, define the local sensitivity constant

$$c_{RV} = 2|R_t^{(0)}|.$$

Then, under the sign choice $R_t^{(0)} < 0$, we have $\Delta RV_t \geq c_{RV}\alpha q$. Combining this with risk sensitivity of g and the liquidation elasticity yields a lower bound on forced liquidation at $t+1$ of the form

$$L \geq c_X c_g c_{RV} \alpha q.$$

Define the corresponding feedback gain

$$G_t = \alpha c_X c_g c_{RV}. \tag{14}$$

The next lemma computes the profit from the two-period round trip.

Lemma 5 (Profit from the two-period round trip)**.** *Fix a reachable state S_t satisfying Assumptions 1–4. Consider the two-period round trip $(a_t, a_{t+1}) = (-q, q)$ with $q > 0$. Suppose q is admissible, meaning that all locality restrictions hold (including $RV_t + \Delta RV_t \in (\underline{r}, \overline{r})$ and $M_{t+1} + \Delta M_{t+1} \in I$) and that the hard funding constraint (12) holds if imposed. Then the expected inventory-adjusted profit satisfies*

$$\mathbb{E}[\Pi_1^\kappa \mid S_t] \geq \alpha q \left(G_t q - 2q\right) - \kappa q^2 = \left(\alpha(G_t - 2) - \kappa\right) q^2. \tag{15}$$

In particular, if $\alpha(G_t - 2) > \kappa$, then $\mathbb{E}[\Pi_1^\kappa \mid S_t] > 0$ for every admissible $q > 0$ (equivalently, for every $q \in (0, q_{\max}]$ in the local range).

Proof. Under the round trip, $\Pi_1^\kappa = -a_t P_t - a_{t+1} P_{t+1} - \kappa |y_t|^2$. Since $a_t = -q$, $a_{t+1} = q$, and $y_t = -q$, this becomes

$$\Pi_1^\kappa = q(P_t - P_{t+1}) - \kappa q^2.$$

Write prices under the deviation relative to the benchmark: $P_t = P_t^{(0)} - \alpha q$ and $P_{t+1} = P_{t+1}^{(0)} + \alpha(q - L)$. Hence

$$P_t - P_{t+1} = \left(P_t^{(0)} - P_{t+1}^{(0)}\right) + \alpha(L - 2q).$$

Assume the benchmark has no predictable drift from t to $t+1$, i.e., $\mathbb{E}[P_{t+1}^{(0)} \mid S_t] = P_t^{(0)}$, so $\mathbb{E}[P_t^{(0)} - P_{t+1}^{(0)} \mid S_t] = 0$. Taking conditional expectations gives

$$\mathbb{E}[\Pi_1^\kappa \mid S_t] = \alpha q(L - 2q) - \kappa q^2.$$

Substituting the lower bound $L \geq c_X c_g c_{RV} \alpha q$ yields (15). □

It is worth stressing once more that the deviation is local. Admissibility requires the trade to be small enough that (RV_t, M_{t+1}) remains in the risk-sensitive region and within the model's local elasticity bounds. In the present special case, this means $2\left|R_t^{(0)}\right| \alpha q + \alpha^2 q^2 \leq \overline{r} - RV_t^{(0)}$ when the deviation raises RV_t. Lemma 5 then shows that, whenever the local gain condition holds, expected profit is strictly positive even for arbitrarily small admissible trigger sizes.

3.5 Main Theorem

Lemma 5 gives a transparent two-period deviation and isolates a local gain condition under which the expected profit is strictly positive for arbitrarily small triggers. Call a trigger size $q > 0$ admissible at S_t if all locality restrictions required by Assumptions 1–4 (and, when imposed, the capacity and hard-funding constraints) continue to hold under the deviation. Let $q_{\max}(S_t) \in (0, \infty]$ denote the supremum of admissible trigger sizes at S_t, so that every $q \in (0, q_{\max}(S_t)]$ is admissible.[10]

[10] A sufficient admissible range can be constructed explicitly by bounding how far the deviation can move the risk input and next-period requirement; one convenient choice is $q_{\max} = \min\{q_I, q_M, \bar{a}, \bar{F}/M_t\}$ with q_I ensuring $RV_t + \Delta RV_t \leq \overline{r}$

Theorem 1 (Impossibility)**.** *Suppose Assumptions 1–4 hold at a reachable state S_t (and the benchmark drift condition used in Lemma 5 holds). If $\alpha(G_t - 2) > \kappa$ and $q_{\max}(S_t) > 0$, then the mechanism is not round-trip manipulation-proof at S_t: for every admissible $q \in (0, q_{\max}(S_t)]$, the two-period round trip has strictly positive conditional expected inventory-adjusted profit. Consequently, risk sensitivity (Definition 2), round-trip manipulation-proofness (Definition 3), and liquidity continuity (Definition 4) cannot hold simultaneously at S_t.*

Proof. Fix any admissible $q \in (0, q_{\max}(S_t)]$. By Lemma 5,

$$\mathbb{E}[\Pi_1^\kappa \mid S_t] \geq (\alpha(G_t - 2) - \kappa)\, q^2.$$

Under $\alpha(G_t - 2) > \kappa$, the right-hand side is strictly positive for every $q > 0$, so $\mathbb{E}[\Pi_1^\kappa \mid S_t] > 0$. This violates round-trip manipulation-proofness (Definition 3), hence the three properties cannot all hold at S_t. □

Three points matter for interpretation. First, Theorem 1 is local. Under the gain condition, expected inventory-adjusted profit is strictly positive for every admissible trigger size $q \in (0, q_{\max}]$, so the mechanism does not require large trades or large price moves. Second, the theorem concerns states in which the constraint already binds. The trader does not need to make a slack constraint bind; a marginal increase in M_{t+1} is enough to reduce aggregate exposure at the margin. Third, the predictable component is created by the update rule itself. The argument does not require predictable drift in fundamentals, and it applies even when the underlying impact model is manipulation-free absent this feedback channel.

The proof uses a two-period round trip because manipulation-proofness is an incentive condition: one profitable deviation is enough to violate it at the state under study. In practice, requirements reset repeatedly, so the same logic can operate across update cycles. Nicolai and Risteska (2026) develop that dynamic extension. They show that optimal finite-horizon strategies combine trading against predictable post-update forced flow with trades inside the sampling window that reshape later rebalancing by moving the measured statistic. The same two objects govern that dynamic problem: the sensitivity of the risk input to sampled transaction prices and the local slope of g in binding regions.

Theorem 1 also yields a design restriction. In binding states, a risk-sensitive rule passes sampled transaction prices into next-period requirements at first order. If that pass-through is strong enough, risk sensitivity, liquidity continuity, and manipulation-proofness cannot coexist. A price-based margin rule that seeks to avoid this problem must therefore limit short-horizon pass-through from sampled prices into posted requirements. This provides a structural rationale for slope limits, smoothing, and buffers. Section 6 shows that, in empirically interpretable units, the implied restriction can tightly bound how fast posted requirements may respond to short-run changes in measured risk.

Corollary 1 (Implied local slope cap)**.** *Fix a reachable binding state S_t satisfying Assumptions 2–4 and suppose g is differentiable at the benchmark input $RV_t^{(0)}$. If liquidity continuity holds at S_t and the mechanism*

and q_M ensuring $M_{t+1}^{(0)} + \Delta M_{t+1} \leq M_{t+1}^{(0)} + \delta_X$.

is round-trip manipulation-proof at S_t, then the local pass-through slope must satisfy

$$g'\left(RV_t^{(0)}\right) \le \frac{2\alpha+\kappa}{\alpha^2 c_X c_{RV}} = \frac{2\alpha+\kappa}{2\alpha^2 c_X \left|R_t^{(0)}\right|}, \tag{16}$$

where $c_{RV} = 2\left|R_t^{(0)}\right|$ is the local sensitivity constant of the risk-measure. In particular, since $\left|R_t^{(0)}\right| \le \sqrt{RV_t^{(0)}}$, a sufficient implementable envelope stated as a function of the variance-type input alone is

$$g'(r) \le \frac{2\alpha+\kappa}{2\alpha^2 c_X \sqrt{r}}. \tag{17}$$

Equivalently, writing $\sigma = \sqrt{r}$ and $\tilde{g}(\sigma) = g(\sigma^2)$, the envelope implies the constant volatility-scale cap

$$\tilde{g}'(\sigma) \le \frac{2\alpha+\kappa}{\alpha^2 c_X}. \tag{18}$$

Proof. Suppose, toward a contradiction, that $g'\left(RV_t^{(0)}\right) > (2\alpha+\kappa)/(\alpha^2 c_X c_{RV})$. By continuity of g', there exist an interval I containing $RV_t^{(0)}$ and a constant $\underline{c}_g > 0$ such that $g'(r) \ge \underline{c}_g$ for all $r \in I$ and $\underline{c}_g > (2\alpha+\kappa)/(\alpha^2 c_X c_{RV})$. Choose $q > 0$ small enough that the deviation keeps the perturbed input inside I and within the local elasticity region of Assumption 3. Then Lemma 5 implies

$$\mathbb{E}[\Pi_1^\kappa \mid S_t] \ge \left(\alpha\left(\alpha c_X \underline{c}_g c_{RV} - 2\right) - \kappa\right) q^2 > 0,$$

which contradicts round-trip manipulation-proofness at S_t. Hence (16) must hold. The envelope (17) follows from $\left|R_t^{(0)}\right| \le \sqrt{RV_t^{(0)}}$. Finally, for $\tilde{g}(\sigma) = g(\sigma^2)$ we have $\tilde{g}'(\sigma) = 2\sigma g'(\sigma^2)$, and (18) follows from (17). □

3.6 Back-of-the-envelope magnitudes

The magnitude implication in Theorem 1 can be summarized by two objects that map directly into data. First, the forced-liquidation multiple $G = L/q$, where q is the trigger trade at the sampling time t and L is the additional forced liquidation at $t+1$ induced by the higher posted requirement. Under linear impact, G is the number of units liquidated per unit triggered. Lemma 5 implies that the two-period trigger-and-reverse deviation is profitable when $\alpha(G-2) > \kappa$. The T-periods deviation is in general less demanding. Second, the sampled-price distortion. Let δ_{bps} denote the distortion of the sampled mark created by the trigger trade, in bps of notional. With $P_t = P_t^{(0)} - \alpha q$, the relative price displacement is $\alpha q/P_t$, so $\delta_{\text{bps}} = 100\,\delta_{\%}$ where $\delta_{\%} = 100\frac{\alpha q}{P_t}$. A conservative way to discipline δ_{bps} is to express the trigger size in participation units $f = (qP_t)/ADV$, where ADV is average daily traded value. Using a standard square-root scaling of impact in participation units, we use the benchmark mapping $\delta_{\text{bps}}(f) \approx 200\sqrt{f}$, with f measured as a fraction of full-day ADV.[11] Using full-day ADV

[11]The square-root form is a standard empirical benchmark for how implementation shortfall scales with participation (Frazzini et al., 2018; Kyle and Obizhaeva, 2016). To fix a conservative scale, start from an empirical one-way average cost curve $cost_{\text{bps}}(f) \approx k\sqrt{f}$, where f is the trade size as a fraction of daily volume. Estimates in Frazzini et al. (2018) suggest k on the order of 100 bps, so a 1% of ADV trade has average one-way cost about $100\sqrt{0.01} \approx 10$ bps. The distortion that

is conservative for margining because marks are typically computed from a much shorter sampling window, so effective participation within the window is higher than f. Lemma 5 implies the profit-per-notional lower bound in percent units,

$$\pi_{\%} = 100 \frac{\mathbb{E}[\Pi_1 \mid S_t]}{qP_t} \geq (G-2)\,\delta_{\%},$$

equivalently, in bps units,

$$\pi_{\text{bps}} = 100\,\pi_{\%} \geq (G-2)\,\delta_{\text{bps}}.$$

Normalizing by peak margin usage mqP_t, where $m = M_t/P_t$ is the margin ratio, the implied return-on-margin is

$$ROI\% = 100\frac{\mathbb{E}[\Pi_1 \mid S_t]}{m\,qP_t} \geq \frac{(G-2)\,\delta_{\text{bps}}}{100\,m}.$$

For example, if $f = 1\%$ so that $\delta_{\text{bps}} \approx 20$, and $G = 6$, then $\pi_{\text{bps}} \geq 80$ bps in one update and $ROI\% \geq 80/(100m)$; for $m = 7\%$, this is about 11% return on posted margin for a single reset.

Then, a simple accounting identity rewrites $G = L/q$ in terms of observable quantities. Let X denote the constraint-sensitive sector's position and define the position-to-margin elasticity $\varepsilon_{X,M} = -d\log X/d\log M$. For a relative requirement change $\Delta M/M$,

$$\frac{L}{X} \approx \varepsilon_{X,M}\frac{\Delta M}{M}.$$

Define the turnover ratio $\tau = (XP_t)/ADV$. Since $f = (qP_t)/ADV$, we have $q/X = f/\tau$, and therefore

$$G = \frac{L}{q} \approx \varepsilon_{X,M}\frac{\Delta M}{M}\frac{\tau}{f}.$$

This highlights why the theorem has bite in binding states: (i) short-horizon margin moves $\Delta M/M$ can be large under standard risk-based models; (ii) impact and thus δ_{bps} are larger in stress; (iii) G grows when the constrained sector is large relative to market depth (τ high) and when the trigger is a small fraction of volume (f small). Appendix C provides the empirical calibration of $\varepsilon_{X,M}$, $\Delta M/M$, and $\delta_{\text{bps}}(f)$.

Table 1 interprets the two-period trigger-and-reverse deviation in economically familiar units. Panel A reports the forced-liquidation multiple $G = L/q$: the number of units of constrained-sector liquidation induced at $t+1$ per unit of trigger trade executed at the sampling time t. Values well above 2 mean that the forced-flow response overwhelms the trader's own round-trip impact loss, which is the profitability screen in Lemma 5. Panel B translates the same cases into profitability. Each cell reports a lower bound on (i) profit in bps of trigger notional, π_{bps}, and (ii) return on posted margin, $ROI\%$, where the denominator is the post-update margin $m_1 qP_t$. For example, with $f = 1\%$ participation and a margin jump from 5% to 7%, the table implies $G = 9$ and a one-update profit bound of at least 140 bps on the trigger notional, corresponding to at least 20% return on posted

matters for a sampled transaction mark is closer to an instantaneous price displacement at the sampling time. Under the standard constant-rate execution with linear instantaneous impact, impact ramps from 0 to a terminal (peak) displacement, so the time-average displacement paid is half the peak; equivalently, peak displacement is about twice the average cost. We therefore use the conservative mapping $\delta_{\text{bps}}(f) \approx 2k\sqrt{f} \approx 200\sqrt{f}$.

Table 1 Illustrative profitability

The table reports illustrative one-update profitability for the trigger-and-reverse deviation under conservative benchmark parameters: $\epsilon_{X,M} = 0.225$ (Hedegaard, 2014), $\tau = 1$, and $\delta_{\text{bps}}(f) \approx 200\sqrt{f}$ (Frazzini et al., 2018). Columns correspond to margin jumps from $m_0 = 5\%$ to $m_1 \in \{6\%, 7\%, 8\%\}$ (so $\Delta M/M \in \{20\%, 40\%, 60\%\}$). Panel A reports the implied forced-liquidation multiple $G = L/q$. Panel B reports the corresponding lower bounds for profit in bps of trigger notional, π_{bps}, and return on posted margin, $ROI\%$, computed from $\pi_{\text{bps}} \geq (G-2)\delta_{\text{bps}}$ and $ROI\% \geq \pi_{\text{bps}}/(100\, m_1)$.

	$m_0 = 5\% \to m_1 = 6\%$	$m_0 = 5\% \to m_1 = 7\%$	$m_0 = 5\% \to m_1 = 8\%$
Panel A: liquidation multiple $G = L/q$			
$f = 0.5\%$	9.0	18.0	27.0
$f = 1\%$	4.5	9.0	13.5
$f = 2\%$	2.25	4.5	6.75
Panel B: profitability lower bounds			
	$\pi_{\text{bps}} \geq (G-2)\delta_{\text{bps}}$ and $ROI\% \geq \pi_{\text{bps}}/(100\, m_1)$		
	$\pi_{\text{bps}}/ROI\%$	$\pi_{\text{bps}}/ROI\%$	$\pi_{\text{bps}}/ROI\%$
$f = 0.5\%$	99.0/16.5	226.3/32.3	353.6/44.2
$f = 1\%$	50.0/8.3	140.0/20.0	230.0/28.8
$f = 2\%$	7.1/1.2	70.7/10.1	134.4/16.8

margin in a single update. Even in the mildest scenario ($5\% \to 6\%$) at $f = 1\%$, the bound is 50 bps and 8.3% on margin.

These numbers are deliberately conservative in three ways. First, $\delta_{\text{bps}}(f)$ is disciplined using full-day ADV rather than the typically thinner settlement or snapshot window. Second, the ROI denominator uses the higher post-update margin m_1 rather than the pre-update margin m_0. Third, the table reports only the direct one-step margin-update channel captured by Lemma 5. It therefore understates profits from dynamic attacks that repeat the mechanism across update times, exploit state dependence of impact and margin sensitivity in stress, or combine trigger and harvest legs across contracts under portfolio margining (Section 5). The calibration inputs behind Table 1 (elasticity $\epsilon_{X,M}$, stress-time margin moves $\Delta M/M$, and the impact mapping $\delta_{\text{bps}}(f)$) and the full sensitivity grids over $(\epsilon_{X,M}, \tau, f)$ are collected in Appendix C.

3.6.1 How large are one-update margin moves in practice?

In Table 1 we use 20%, 40%, and 60% increases in posted margin as benchmark cases. These magnitudes are realistic under risk-based initial margin models, even with anti-procyclicality (APC) tools. APC tools mainly raise margins in calm periods through buffers and floors. That lowers return on margin mechanically, but it does not eliminate large stress-time increases unless the floor or buffer is so large that risk sensitivity is materially weakened.

Murphy et al. (2014) report short-horizon margin calls of this order as a share of portfolio value. Interpreted as changes in the required margin ratio, and starting from a representative margin level of about 5% (Hedegaard, 2014), their results imply increases of about 20% over one day and about 40% over multi-day horizons under standard specifications. If margins start from an unusually calm state, the same absolute call implies an even larger relative increase. Evidence from March 2020 points in the same direction. Gurrola-Perez (2020) documents large short-horizon increases for S&P 500 futures under filtered historical simulation VaR and shows that materially reducing the largest multi-day jumps requires large APC floors, at the cost of substantially higher margins in normal times. This is exactly the tension in the theorem: reducing stress-time jumps requires limiting local pass-through, which moves the rule away from local risk sensitivity.

The same logic is not specific to CCPs. Any price-based rule that maps sampled prices into a binding requirement can generate large short-horizon adjustments in binding states, and those are precisely the states in which the profitability screen in Lemma 5 is relevant.

4 Extensions

For exposition, the main text uses a deliberately bare-bones specification: linear price impact, a realized-variance risk input, and a simple margin-cap demand rule. Here we record extensions useful for interpretation.

4.1 Continuous trading

The discrete-time indexing in Sections 2–3 indexes margin update times. Trading can be continuous between updates without changing the argument, because the the risk input is computed from a sampling grid of prices and applies the margin update at discrete times. Fix an update interval $h > 0$ and update times $t_n = nh$, $n = 0, 1, 2, \ldots$. Let $P_n = P_{t_n}$ and $V_n = V_{t_n}$. Let ΔX_n denote the constrained traders' net order flow aggregated over $(t_{n-1}, t_n]$, and let a_n denote the strategic trader's net order flow over the same interval. At update times, transaction prices satisfy

$$P_n = V_n + \alpha(\Delta X_n + a_n), \tag{19}$$

realized variance is computed from the last m sampled marks,

$$RV_n = \sum_{j=0}^{m-1} (P_{n-j} - P_{n-j-1})^2, \tag{20}$$

and next-period margin is

$$M_{n+1} = g(RV_n). \tag{21}$$

Constrained traders face the same margin constraint at update times and adjust mechanically when M_{n+1} is posted.

Proposition 1 (Continuous trading with discrete margin updates)**.** Fix $h > 0$ and the update-time model (19)–(21). Suppose the local assumptions of Section 3.2 hold at some reachable update-time state, interpreted with (P_t, V_t, RV_t, M_t) replaced by (P_n, V_n, RV_n, M_n). Suppose also that the local conditions of Lemma 5 hold at that state (with t replaced by n). Then the same two-step round trip (net selling over $(t_{n-1}, t_n]$ and net buying over $(t_n, t_{n+1}]$) delivers the same profit lower bound (15). In particular, whenever the amplification condition in Lemma 5 holds, the impossibility result of Theorem 1 applies at that update-time state.

Risk measures are typically computed from settlement prices, end-of-day marks, or specified intraday snapshots. Even in a continuously traded market, the risk input is built from a discrete sample, and the margin update is a discrete event.[12]

4.2 Alternative risk inputs

We use realized variance in the main text because its local sensitivity to the last sampled price is transparent. The mechanism does not rely on squaring returns. What matters is that the CCP computes a price-based risk input from sampled marks, and that some implementable perturbation of a sampled price moves this input at first order.

4.2.1 A generic price-based risk functional

Let the CCP compute a scalar risk statistic from the sampled price vector,

$$\Gamma_t = \Gamma(P_{t-m}, \ldots, P_t), \qquad M_{t+1} = g(\Gamma_t).$$

The Appendix allows non-smooth Γ and non-smooth overlays inside g. Here we isolate the single local property that replaces the realized-variance sensitivity used in the main text.

Assumption 5 (Local one-sided sensitivity of a generic risk statistic)**.** There exists a reachable state S_t such that $\Gamma_t \in \text{int}(I_\Gamma)$ for some interval $I_\Gamma \subset \mathbb{R}_+$, and there exist a sign choice and a constant $c_\Gamma > 0$ such that, holding $(P_{t-m}, \ldots, P_{t-1})$ fixed,

$$\Gamma(P_{t-m}, \ldots, P_{t-1}, P_t + \Delta P) - \Gamma(P_{t-m}, \ldots, P_{t-1}, P_t) \;\geq\; c_\Gamma |\Delta P|$$

for all sufficiently small ΔP of that sign.

Assumption 5 says that, at some reachable state, there is an implementable one-sided perturbation of the last sampled price that raises the risk input at a nontrivial linear rate.

Theorem 2 (Impossibility for a locally sensitive price-based risk input)**.** *Replace* (4) *by* $\Gamma_t = \Gamma(P_{t-m}, \ldots, P_t)$ *and* (3) *by* $M_{t+1} = g(\Gamma_t)$*. Replace Assumption 2 by Assumption 5. Assume g is risk-sensitive on an interval I_Γ containing Γ_t in its interior, and retain Assumptions 3–4 and the local conditions used in Lemma 5. Then, for the same two-period round trip (sell q at t and buy q at $t+1$) and for q small enough that the perturbed Γ_t remains in I_Γ,*

$$\mathbb{E}[\Pi_1^\kappa \mid S_t] \;\geq\; \alpha q \left(\alpha c_X c_g c_\Gamma q - 2q\right) \;-\; \kappa q^2 \;=\; \big(\alpha(\alpha c_X c_g c_\Gamma - 2) - \kappa\big) q^2.$$

[12] This extension speaks directly to the view that continuous trading between update times eliminates the mechanism. It does not. For any fixed update interval $h > 0$, Proposition 1 applies unchanged, because the CCP still computes the risk input from a discrete set of marks and resets margin at discrete times. The only boundary case in which the first-order effect can vanish is an idealized diffusion limit that simultaneously (i) lets the sampling grid become arbitrarily fine, (ii) assumes continuous price paths, and (iii) keeps the trader's per-update trade size bounded, so that $\Delta RV_n = 2R_n \Delta P_n + (\Delta P_n)^2$ with $R_n = O(\sqrt{h})$ and the linear term disappears as $h \downarrow 0$. But that is not continuous trading in any institutional sense. It requires a smooth constraint to be recomputed and enforced essentially continuously from a continuously sampled price path, with no discrete marks, such as settlements, auctions, or snapshots, and no operational non-smoothness, such as thresholds, rounding, or stress add-ons. Actual margin systems continue to rely on discrete marks and discrete enforcement, including in deep and liquid markets. Proposition 1 therefore covers the economically relevant case.

In particular, if $\alpha c_X c_g c_\Gamma > 2 + \kappa/\alpha$, then $\mathbb{E}[\Pi_1^\kappa \mid S_t] > 0$ for every admissible q in the local range.

Proof. With $\Delta X_t = 0$, selling q at t moves the time-t transaction price by $\Delta P_t = -\alpha q$. By Assumption 5 (choosing the sign appropriately),

$$\Delta\Gamma_t \geq c_\Gamma |\Delta P_t| = c_\Gamma \alpha q.$$

Risk sensitivity of g on I_Γ yields $\Delta M_{t+1} \geq c_g \Delta\Gamma_t \geq c_g c_\Gamma \alpha q$, and Assumption 3 yields $L \geq c_X \Delta M_{t+1} \geq c_X c_g c_\Gamma \alpha q$. The profit identity in Lemma 5 gives $\mathbb{E}[\Pi_1^\kappa \mid S_t] = \alpha q(L - 2q) - \kappa q^2$, so substituting the lower bound for L yields the claim. □

4.2.2 VaR/ES built from volatility

A common implementation sets margin proportional to a VaR or ES multiple of a volatility estimate, for example $\Gamma_t = z_p \hat{\sigma}_t$ or $\Gamma_t = k_p \hat{\sigma}_t$. If $\hat{\sigma}_t$ loads on recent squared returns with positive weight, then $\hat{\sigma}_t$ inherits a local first-order sensitivity to a perturbation of a sampled price whenever the relevant sampled return is not exactly zero. For the special case $\hat{\sigma}_t = \sqrt{RV_t/(m-1)}$,

$$\frac{\partial \hat{\sigma}_t}{\partial P_t} = \frac{1}{2\sqrt{(m-1)RV_t}} \cdot \frac{\partial RV_t}{\partial P_t} = \frac{R_t}{\sqrt{(m-1)RV_t}}.$$

This expression highlights a sharp state dependence: holding the most recent return R_t fixed, the local slope scales like $1/\sqrt{RV_t}$. Hence the same small perturbation of the last sampled mark has a much larger effect on $\hat{\sigma}_t$ after a quiet spell (low RV_t) than in an already volatile state. Put differently, low past measured risk is precisely when volatility-based risk inputs are locally steep: a single nontrivial return arriving after low realized variance mechanically produces a large change in $\hat{\sigma}_t$, and therefore in Γ_t and M_{t+1}.

Existing leverage-cycle and endogenous-risk work focuses on an equilibrium channel: when measured risk is low, constraints and funding terms loosen, balance sheets expand, and fragility accumulates (Adrian and Shin, 2010; Brunnermeier and Pedersen, 2009; Geanakoplos, 2010; Brunnermeier and Sannikov, 2014; Danielsson et al., 2004). Our point is different. We isolate a purely mechanical, local object for price-based rulebooks: the directional sensitivity of the posted requirement to a manipulable sampled mark. For volatility-based inputs this sensitivity is steepest precisely in tranquil states (low recent RV_t), so the states that look safest under the rulebook are also the states that are easiest to move at first order. To the best of our knowledge, this link between low measured risk and maximal first-order manipulability of price-based constraints is not made explicit in the existing leverage-cycle and procyclicality literatures. It is central for large, liquid markets, which spend long stretches in low-volatility regimes.

4.2.3 Scenario maxima, quantiles, and kinks

Many margin rules compute a scalar risk number as a maximum over scenarios or as a high quantile of scenario losses:

$$\Gamma_t = \max_{s\in\mathcal{S}} l_s(P_{t-m:t}), \qquad \text{or} \qquad \Gamma_t = \text{Quantile}_p\big(l_1(P_{t-m:t}),\ldots,l_N(P_{t-m:t})\big).$$

These mappings are typically non-differentiable when the identity of the worst scenario changes. Away from kink points, a single active scenario governs the slope, and the analysis reduces to Theorem 2 with c_Γ equal to the one-sided slope of the active scenario. At kink points, the relevant object is a directional derivative or subgradient, which is handled in the Appendix.

4.2.4 Hard trimming and margin cliffs

To illustrate how hardening a risk input can interact with continuity, consider trimming, which discards sampled returns whose absolute size exceeds a cutoff:

$$\Gamma_t^{\text{trim}} = \sum_{j=0}^{m-1} (P_{t-j} - P_{t-j-1})^2 \cdot \mathbf{1}\{|P_{t-j} - P_{t-j-1}| \le c\}, \qquad c > 0. \tag{22}$$

A return just below c is counted (contributing close to c^2), while a return just above c is dropped (contributing 0). This creates a discrete change in the risk input, and therefore a discrete change in posted margin under any strictly increasing margin mapping.

Proposition 2 (Hard trimming creates discontinuous margin updates)**.** Fix $c > 0$ and define Γ_t^{trim} by (22). Assume g is strictly increasing and continuous. Holding $(P_{t-m},\ldots,P_{t-1})$ fixed, the mapping $P_t \mapsto M_{t+1} = g(\Gamma_t^{\text{trim}})$ is discontinuous at any state with $|R_t| = c$, where $R_t = P_t - P_{t-1}$.

Proof. Holding $(P_{t-m},\ldots,P_{t-1})$ fixed, write $\Gamma_t^{\text{trim}} = A + (P_t - P_{t-1})^2\mathbf{1}\{|P_t - P_{t-1}| \le c\}$ for a constant A. At $|P_t - P_{t-1}| = c$, perturbing P_t slightly toward the interior makes the indicator equal to 1 (adding about c^2), while perturbing slightly outward makes it equal to 0 (adding 0). Hence Γ_t^{trim} jumps by c^2 at the cutoff, and by strict monotonicity of g, so does M_{t+1}. □

Winsorizing (capping the squared return at c^2 rather than dropping it) restores continuity, but creates locally flat regions in which marginal sensitivity is zero once the cap binds. These examples illustrate the design trade-off: efforts to reduce smooth marginal manipulability tend to replace slopes with kinks, caps, or cliffs, and those features interact directly with continuity and risk sensitivity.[13]

[13] In the mechanism, profitability is governed by a local amplification: the product of (i) sensitivity of the risk input to sampled prices, (ii) pass-through from the input to posted margin, (iii) the margin elasticity of constrained demand, and (iv) price impact. Slope caps (winsorization) and hard thresholds (trimming, stepwise overlays) limit this amplification by flattening local slopes or concentrating adjustment at kinks. We use these examples only to illustrate how robustification interacts with continuity and sensitivity; the optimal choice of overlays is taken up in the design problem.

4.3 Strategic constrained traders

The baseline model takes the constrained sector's target exposure as predetermined at the start of the pricing window. Appendix B relaxes this. Constrained traders are forward-looking: they choose targets anticipating that within-window trading can move the sampled mark, raise the posted requirement $M_{t+1} = g(\Gamma_t)$, and trigger forced deleveraging with price impact. The appendix therefore studies a two-sided subgame-perfect equilibrium in which (i) constrained traders choose targets as a function of the state S_t, and (ii) the strategic trader chooses the trigger size (equivalently a_t or q) in response. The only change to the main mechanism is that equilibrium targeting alters the local forced-flow slope. In the binding states, the relevant object is an effective liquidation elasticity

$$c_X^{\text{eff}}(S_t) \;=\; -\frac{\partial X}{\partial M}\big(M_t, \bar{x}_t^*(S_t)\big),$$

which is weakly smaller than the non-strategic benchmark because forward-looking traders scale back exposures in states where the update is more susceptible to manipulation. The amplification condition is therefore modified by replacing c_X with $c_X^{\text{eff}}(S_t)$.

Crucially, forward-looking behavior can dampen the loop but cannot eliminate it whenever constraints bind locally. If the cap binds at the equilibrium target, then $c_X^{\text{eff}}(S_t) > 0$, and the same two-period trigger-and-reverse logic applies with c_X replaced by $c_X^{\text{eff}}(S_t)$. Eliminating the channel requires $c_X^{\text{eff}}(S_t) = 0$ in the relevant states, which means that endogenous participation collapses or the constraint is kept slack precisely when risk-sensitive margining is intended to be operative. The equilibrium response therefore shifts the problem into an economically meaningful trade-off: robustness to manipulation is obtained through ex ante contraction of constrained-sector exposure and liquidity.[14] From a design perspective, this strengthens the case for limiting marginal pass-through. Reducing the local sensitivity of posted requirements lowers both the direct profitability of within-window manipulation and the indirect ex-ante withdrawal of constrained-sector risk-bearing capacity highlighted in Appendix B. In the design problem of Section 6, this is exactly the tension between the first-best desire for risk-sensitive updates and the implementability requirement that the public rule remain manipulation-proof in binding states.

5 Multi-asset portfolio margining and cross-contract contagion

It is common to set initial margin at the portfolio level rather than contract by contract. The posted requirement is a scalar function of a portfolio risk statistic computed from a vector of prices. When the requirement binds, meeting a margin call can force liquidation in contracts that are not the ones that most efficiently move the portfolio statistic. This wedge does not arise in the single-contract model: a trader may be able to move the portfolio risk input relatively cheaply in one subset of contracts, while the induced deleveraging is concentrated in a different subset. The main point of this section is that portfolio margining generates a cross-contract manipulation channel and a natural notion of cross-contract contagion. The same feedback loop is at work: a trade at t moves transaction

[14] Empirically, this outside-option response is not hypothetical: in centrally cleared repo, increases in CCP haircuts push safer borrowers toward OTC trading, with a stronger effect for collateral-constrained borrowers Chebotarev (2025).

prices P_t, which shifts the portfolio risk input Γ_t and hence the posted requirement M_{t+1}. If the requirement binds, the margin call induces forced rebalancing of constrained positions X_{t+1}, and the resulting liquidation order flow feeds back into prices at $t+1$. In the multi-asset setting, P_t and X_t are vectors, while the margin call is scalar, so the induced liquidation can be concentrated in specific contracts.

5.1 Setup: vector prices, scalar margin, and cross-margining liquidation

There are $K \geq 1$ contracts. Transaction prices and fundamentals are vectors $P_t, V_t \in \mathbb{R}^K$. Constrained traders hold an aggregate vector position $X_t \in \mathbb{R}^K$, with net order flow $\Delta X_t = X_t - X_{t-1}$. A strategic trader submits a trade vector $a_t \in \mathbb{R}^K$ subject to a capacity constraint $\|a_t\|_1 \leq \bar{a}$.

5.1.1 Price formation and profits

Price formation at update times is linear with cross-impact:

$$P_t = V_t + A\big(\Delta X_t + a_t\big), \tag{23}$$

where $A \in \mathbb{R}^{K\times K}$ is an impact matrix. Trading profit over a horizon T is

$$\Pi_T(a_{t:t+T}) = -\sum_{s=0}^{T} a_{t+s}^\top P_{t+s}. \tag{24}$$

A (vector) round trip over horizon T is a trade sequence $(a_t, \ldots, a_{t+T})$ satisfying $\sum_{s=0}^{T} a_{t+s} = 0$ in $\mathbb{R}^K$. Fundamentals satisfy the same martingale restriction as in (1), componentwise:

$$\mathbb{E}[V_{t+1} \mid \mathcal{F}_t] = V_t.$$

As in the single-asset model, this removes predictable drift in fundamentals from the profit calculation; it does not remove predictability generated mechanically by margin-induced liquidation.

5.1.2 Margin rule

A scalar portfolio risk statistic is computed from the sampled vector price path and a scalar margin is posted:

$$\Gamma_t = \Gamma(P_{t-m}, \ldots, P_t), \qquad M_{t+1} = g(\Gamma_t). \tag{25}$$

Write $P_{t-m:t} = (P_{t-m}, \ldots, P_t)$ and $P_{t-m:t-1} = (P_{t-m}, \ldots, P_{t-1})$. To state local assumptions, we use the multi-asset analog of the public state in Section 2:

$$S_t = (V_t, P_{t-m}, \ldots, P_{t-1}, M_t),$$

where each $P_{t-j} \in \mathbb{R}^K$ is a vector of sampled marks.

5.1.3 Liquidation

The new object relative to the single-asset model is how a scalar margin increase maps into a vector liquidation response. We represent the deviation-induced liquidation response at $t+1$ by the linear map

$$\Delta X_{t+1} = \Delta X^{(0)}_{t+1} - B_t\,\Delta M_{t+1}, \qquad \Delta M_{t+1} = M_{t+1} - M^{(0)}_{t+1}, \tag{26}$$

where $\Delta X^{(0)}_{t+1}$ and $M^{(0)}_{t+1}$ denote the baseline outcomes absent a time-t deviation, and $B_t \in \mathbb{R}^K$ is an $\mathcal{F}_t$-measurable liquidation direction capturing cross-margining. Thus a higher scalar requirement forces an additional portfolio reduction of $B_t \Delta M_{t+1}$, and the resulting sales need not occur in the contracts that most effectively move Γ_t.

5.1.4 Assumptions

Assumption 6 (Impact matrix)**.** The matrix A in (23) has strictly positive diagonal entries and is symmetric positive definite.[15]

Assumption 7 (Directional sensitivity of the portfolio risk statistic)**.** There exist a reachable state S_t, constants $c_\Gamma > 0$ and $\delta_\Gamma > 0$, and a unit direction $d_t \in \mathbb{R}^K$ such that, holding $P_{t-m:t-1}$ fixed, for any perturbation ΔP_t with $d_t^\top \Delta P_t \in [0, \delta_\Gamma]$ and small enough that $P_t + \Delta P_t$ remains in the local neighborhood considered,

$$\Gamma(P_{t-m:t-1}, P_t + \Delta P_t) - \Gamma(P_{t-m:t-1}, P_t) \geq c_\Gamma\, d_t^\top \Delta P_t. \tag{27}$$

Assumption 8 (Local slope of the margin mapping)**.** There exists an interval $I_\Gamma \subset \mathbb{R}_+$ with $\Gamma_t \in \operatorname{int}(I_\Gamma)$ such that g is differentiable on I_Γ and

$$g'(\gamma) \geq c_g > 0 \quad \text{for all } \gamma \in I_\Gamma.$$

Assumption 9 (Cross-margining liquidation is nontrivial)**.** At the reachable state in Assumption 7, the liquidation direction in (26) satisfies $\|B_t\|_1 \geq c_X > 0$.

Assumptions 6–9 are the multi-asset counterparts of the three local links used in the single-asset argument. Assumption 6 pins down how a vector order perturbation maps into transaction prices through the cross-impact matrix A, and imposes the standard no-dynamic-arbitrage restriction on the impact matrix (so round-trip profits do not arise absent the margin-feedback loop). Assumption 7 is the local one-sided bound on how the scalar portfolio risk statistic reacts to an admissible perturbation of the last sampled mark. Assumption 9 is the local statement that a scalar margin increase induces a nontrivial vector liquidation response, summarized by B_t.

[15] Schneider and Lillo (2019) characterize the no-dynamic-arbitrage restriction for cross-impact. If $I = (A + A^\top)/2$ is not positive semidefinite, pick v with $v^\top I v < 0$; under linear impact $P_t = P_t^{(0)} + A a_t$ and zero expected drift, the round trip $a_t = qv$, $a_{t+1} = -qv$ has expected execution cost $q^2 v^\top A v = q^2 v^\top I v < 0$, hence yields arbitrage. We therefore impose a standard no-manipulation restriction on the impact matrix, so the round trips we identify arise despite, not because of, an arbitrageable impact model.

5.2 Local bounds in the multi-asset setting

Fix a reachable state satisfying Assumptions 6–9. We consider local deviations so that (27) applies and so that g is evaluated inside I_Γ. The multi-asset argument uses the same slope chain as in the single-asset model, with vectors and cross-impact.

Lemma 6 (A deviation moves the transaction-price vector)**.** *Holding V_t and ΔX_t fixed, perturbing a_t by Δa_t perturbs the time-t transaction price by*

$$\Delta P_t = A\,\Delta a_t.$$

Proof. Immediate from (23): $P_t = V_t + A(\Delta X_t + a_t)$. □

Lemma 7 (A price perturbation moves the portfolio risk statistic at first order)**.** *Suppose Assumption 7 holds at S_t. Holding $P_{t-m:t-1}$ fixed, any perturbation ΔP_t with $d_t^\top \Delta P_t \in [0, \delta_\Gamma]$ satisfies*

$$\Delta\Gamma_t = \Gamma(P_{t-m:t-1}, P_t + \Delta P_t) - \Gamma(P_{t-m:t-1}, P_t) \;\geq\; c_\Gamma\, d_t^\top \Delta P_t.$$

Proof. This is exactly (27). □

Lemma 8 (Risk sensitivity of g passes through to margin)**.** *Under Assumption 8, for any $\Delta\Gamma_t \geq 0$ such that $\Gamma_t + \Delta\Gamma_t \in I_\Gamma$,*

$$\Delta M_{t+1} = g(\Gamma_t + \Delta\Gamma_t) - g(\Gamma_t) \;\geq\; c_g\,\Delta\Gamma_t.$$

Proof. By the mean value theorem, $g(\Gamma_t + \Delta\Gamma_t) - g(\Gamma_t) = g'(\xi)\Delta\Gamma_t$ for some $\xi \in (\Gamma_t, \Gamma_t + \Delta\Gamma_t) \subset I_\Gamma$. Assumption 8 gives $g'(\xi) \geq c_g$. □

Lemma 9 (Cross-margining liquidation and cross-impact)**.** *Under (26), the deviation-induced additional constrained order flow at $t+1$ equals $-B_t\Delta M_{t+1}$. Holding V_{t+1}, $\Delta X^{(0)}_{t+1}$, and a_{t+1} fixed, the liquidation-driven component of the time-$(t+1)$ transaction price is*

$$\Delta P^{\text{liq}}_{t+1} = -AB_t\,\Delta M_{t+1}.$$

Proof. The first claim is (26). Substituting $\Delta X_{t+1} = \Delta X^{(0)}_{t+1} - B_t\Delta M_{t+1}$ into (23) at $t+1$ yields

$$P_{t+1} = V_{t+1} + A\big(\Delta X^{(0)}_{t+1} - B_t\Delta M_{t+1} + a_{t+1}\big),$$

so, holding V_{t+1}, $\Delta X^{(0)}_{t+1}$, and a_{t+1} fixed, the deviation contributes $-AB_t\Delta M_{t+1}$. □

5.3 A two-period cross-asset round trip

We focus on the same object as in the single-contract proof: a two-period round trip that isolates profits generated mechanically by the margin-update channel. Consider a trade vector $a_t \in \mathbb{R}^K$ at t and its reversal $a_{t+1} = -a_t$ at $t+1$, with $\|a_t\|_1 \leq \bar{a}$. As in the single-asset case, we impose a local condition that removes predictable drift in fundamentals and predictable baseline constrained flow, so that any predictability in prices comes from deviation-induced liquidation.

Assumption 10. At S_t, (i) $\Delta X_t = 0$; (ii) letting $\Delta X_{t+1}^{(0)}$ denote constrained order flow at $t+1$ absent a deviation at t, $\mathbb{E}[\Delta X_{t+1}^{(0)} \mid S_t] = 0$; (iii) $\mathbb{E}[V_{t+1} \mid S_t] = V_t$.

To keep the algebra readable, we summarize a direction $a \in \mathbb{R}^K$ by two scalars:

$$m_t(a) = d_t^\top A a, \qquad l_t(a) = a^\top A B_t.$$

The first quantity, $m_t(a)$, measures how strongly the time-t price move loads on the risk-increasing direction d_t because $\Delta P_t = Aa$ and $d_t^\top \Delta P_t = m_t(a)$. The second quantity, $l_t(a)$, measures exposure to the liquidation-driven price component at $t+1$, since liquidation moves prices in the direction AB_t.

Lemma 10 (A profitable two-period round trip)**.** *Suppose Assumptions 6–9 and 10 hold, and consider the two-period round trip $a_{t+1} = -a_t$ with $\|a_t\|_1 \leq \bar{a}$. Assume the deviation is local so that $m_t(a_t) \in [0, \delta_\Gamma]$ and $\Gamma_t + \Delta\Gamma_t \in I_\Gamma$. If $l_t(a_t) \leq 0$, then*

$$\mathbb{E}[\Pi_1 \mid S_t] \;\geq\; -(c_g c_\Gamma)\, l_t(a_t)\, m_t(a_t) \;-\; 2\, a_t^\top A a_t. \tag{28}$$

Proof. Under $a_{t+1} = -a_t$, two-period trading profit is $\Pi_1 = a_t^\top (P_{t+1} - P_t)$. Using (23) at t and $t+1$, together with $\Delta X_t = 0$ and $\Delta X_{t+1} = \Delta X_{t+1}^{(0)} - B_t \Delta M_{t+1}$ from (26), we obtain

$$P_t = V_t + A a_t, \qquad P_{t+1} = V_{t+1} + A\big(\Delta X_{t+1}^{(0)} - B_t \Delta M_{t+1} - a_t\big),$$

hence

$$P_{t+1} - P_t = (V_{t+1} - V_t) + A \Delta X_{t+1}^{(0)} - A B_t \Delta M_{t+1} - 2 A a_t.$$

Taking $\mathbb{E}[\cdot \mid S_t]$ and applying Assumption 10 yields

$$\mathbb{E}[P_{t+1} - P_t \mid S_t] = -A B_t \Delta M_{t+1} - 2 A a_t,$$

so

$$\mathbb{E}[\Pi_1 \mid S_t] = -\big(a_t^\top A B_t\big) \Delta M_{t+1} - 2 a_t^\top A a_t = -l_t(a_t) \Delta M_{t+1} - 2 a_t^\top A a_t.$$

It remains to lower bound ΔM_{t+1}. By Lemma 6, $\Delta P_t = A a_t$, so $d_t^\top \Delta P_t = m_t(a_t)$. If $m_t(a_t) \in [0, \delta_\Gamma]$, Assumption 7 gives $\Delta \Gamma_t \geq c_\Gamma m_t(a_t)$, and Lemma 8 gives $\Delta M_{t+1} \geq c_g \Delta \Gamma_t \geq c_g c_\Gamma m_t(a_t)$. Substituting into the expression for $\mathbb{E}[\Pi_1 \mid S_t]$ yields (28). □

Lemma 10 isolates what is genuinely new under portfolio margining. The time-t leg must raise the scalar portfolio risk input, captured by $m_t(a_t) = d_t^\top A a_t > 0$. But the predictable price component that generates profit at $t+1$ is liquidation-driven and points in the cross-impact direction AB_t, so the relevant exposure of the round trip is $l_t(a_t) = a_t^\top A B_t$. Under contract-by-contract margining these objects collapse to the same sign choice; under portfolio margining they generally decouple because the risk-increasing direction d_t and the forced-deleveraging direction B_t need not align. This creates a trigger-versus-harvest wedge: a trader can move the scalar charge efficiently using one set of contracts while taking offsetting exposure in the contracts that are unwound when the call binds, with the strength governed by the alignment term $-d_t^\top A B_t$. This is distinct from standard predatory-trading/fire-sale mechanisms.

5.4 A sufficient amplification condition

Lemma 10 gives a direction-dependent lower bound. For interpretation, it is useful to have a simple sufficient condition that guarantees existence of at least one admissible two-period round trip with strictly positive conditional expected profit. The next result does this by selecting a concrete direction, namely the cross-margining liquidation direction B_t, and checking when the margin channel dominates the trader's own round-trip impact loss.

Theorem 3 (A sufficient condition for a profitable two-period round trip)**.** *Suppose Assumptions 6–10 hold and that there exists a liquidation direction $B_t \in \mathbb{R}^K$ with $\|B_t\|_1 = 1$ and $d_t^\top A B_t < 0$. Define*

$$H_t = (B_t^\top A B_t)\Big(-(c_g c_\Gamma)\, d_t^\top A B_t - 2\Big). \tag{29}$$

If $H_t > 0$ and $q > 0$ is admissible (so that the locality bounds apply), then the two-period round trip $a_t = -qB_t$, $a_{t+1} = qB_t$ satisfies

$$\mathbb{E}[\Pi_1 \mid S_t] \geq q^2 H_t > 0.$$

In particular, whenever locality permits arbitrarily small $q > 0$, profitable round trips exist for arbitrarily small admissible trigger sizes.

Proof. Let $a_t = -qB_t$ and $a_{t+1} = qB_t$. Then $l_t(a_t) = a_t^\top A B_t = -qB_t^\top A B_t \leq 0$ and $m_t(a_t) = d_t^\top A a_t = -q\, d_t^\top A B_t > 0$. For admissible q, we have $m_t(a_t) \in [0, \delta_\Gamma]$ and $\Gamma_t + \Delta\Gamma_t \in I_\Gamma$, so Lemma 10 applies:

$$\mathbb{E}[\Pi_1 \mid S_t] \geq -(c_g c_\Gamma)\, l_t(a_t)\, m_t(a_t) - 2a_t^\top A a_t = q^2 H_t.$$

If $H_t > 0$, this bound is strictly positive for every admissible $q > 0$. When $K = 1$ and $A = \alpha > 0$, taking $d_t = -1$ and $B_t = c_X$ reduces $H_t > 0$ to the scalar amplification condition. □

Theorem 3 is only a sufficient condition. We choose B_t because it has a clear meaning: it is the direction in which the constrained sector is forced to sell when a portfolio margin call arrives. Portfolio margining separates two things that coincide in the single-contract model. A trader can trigger the margin call using the contracts that move the portfolio risk number most cheaply, but the forced selling that follows can be concentrated in different contracts. By loading on those contracts, the trader can earn predictable profits from the margin-induced flow at $t + 1$.

5.5 Multi-asset impossibility

Theorem 4 (Impossibility under portfolio margining and cross-contract contagion)**.** *Fix a reachable state S_t satisfying the local assumptions in Section 5. Let $\widetilde{q}_{\max} > 0$ denote the maximal trigger size allowed by locality and capacity in the multi-asset setting. If $H_t > 0$ as defined in (29), then for every admissible $q \in (0, \widetilde{q}_{\max}]$ the two-period round trip $a_t = -qB_t$, $a_{t+1} = qB_t$ has strictly positive conditional expected profit:*

$$\mathbb{E}[\Pi_1 \mid S_t] \geq q^2 H_t > 0.$$

In particular, portfolio margining can generate profitable round trips for arbitrarily small admissible trigger sizes. Consequently, risk sensitivity (Definition 2), round-trip manipulation-proofness (Definition 3), and *liquidity continuity (Definition 4) cannot hold simultaneously at S_t.*

Proof. Fix any admissible $q \in (0, \widetilde{q}_{\max}]$ and take $a_t = -qB_t$, $a_{t+1} = qB_t$. Under the local conditions defining H_t, the bound in Theorem 3 applies and gives $\mathbb{E}[\Pi_1 \mid S_t] \geq q^2 H_t$. If $H_t > 0$, expected profit is strictly positive for every admissible $q > 0$, which violates round-trip manipulation-proofness (Definition 3) at S_t. The incompatibility with risk sensitivity and liquidity continuity follows by the same contrapositive logic as in Theorem 1. □

Portfolio margining turns the mechanism into a genuinely cross-contract channel. A vector of prices is mapped into a single portfolio charge, but a binding charge forces liquidation in a particular direction B_t, which need not coincide with the contracts that most efficiently move the portfolio statistic. A trader can therefore trigger a higher portfolio margin using one set of contracts and profit from the predictable forced flow in another set of contracts, with cross-impact transmitting that flow into prices across the cleared complex even when fundamentals are martingales.[16]

5.6 What is new relative to the single-contract case

Portfolio margining creates a wedge absent in the single-contract case. With one contract, the same instrument both determines the risk input for next-period margin and absorbs the forced liquidation when margin tightens. Under portfolio margining, the CCP maps a vector of marks into one scalar charge, but a binding charge can force liquidation in a different set of contracts. This separation creates a distinct cross-contract channel. The trade at t moves marks by $\Delta P_t = Aa_t$, so its first-order effect on the portfolio risk input is

$$d_t^\top \Delta P_t = d_t^\top A a_t,$$

which identifies the trigger direction. The resulting margin increase induces additional constrained selling in direction B_t, which moves next-period prices through cross-impact in direction AB_t. The trader's exposure to this liquidation-driven price component is

$$a_t^\top A B_t,$$

which is the harvest. When $K = 1$, trigger and harvest reduce to the same sign choice. When $K > 1$, they need not coincide: it can be feasible to choose a_t such that the trade at t raises the scalar portfolio charge, $d_t^\top A a_t > 0$, while the position is short the contracts sold to meet the margin call, $a_t^\top A B_t < 0$. The trader triggers the call in one set of contracts and harvests the forced-flow price pressure in another. This trigger-harvest separation is not an artifact of the two-period setup. In a dynamic treatment of price-based constraints, Nicolai and Risteska (2026) show that optimal finite-horizon round trips can be confined to the span of two portfolios: a trigger portfolio that most efficiently moves the marked statistic and a liquidation portfolio that loads on the forced-flow direction. In our notation, these are the directions behind $d_t^\top A a_t$ and AB_t. The portfolio case can therefore make manipulation easier. Manipulation-proofness must rule out all admissible vector round trips, and as K grows, the set of directions that both raise the scalar risk input and load negatively on liquidation expands.

[16] As an example, this is operational in SPAN-like systems, where combined-commodity offsets and portfolio netting link contracts through a common scalar charge (CME Group, 2019a,b).

It is enough that one such direction is profitable at a reachable state for manipulation-proofness to fail. The strength of the effect is governed by an interpretable alignment term: how closely the risk-increasing direction aligns with the liquidation-driven price direction. In the sufficient condition, the key scalar is $-d_t^\top AB_t$.

Portfolio margining also implies an empirical signature absent under contract-by-contract margining. The contracts with trading pressure at t need not be the contracts with illiquidity and forced sales at $t+1$. Churn and subsequent stress can arise in different legs of the market. Empirically, collateral calls on derivatives portfolios have at times translated into concentrated sales and liquidity dislocations in related cash markets, including the UK gilt episode (Sep–Oct 2022) (Pinter, 2023) and the European energy derivatives episode in 2022 (Avalos et al., 2023).

6 Optimal Margin Design

In this section we study the designer's problem of choosing the public margin rule

$$M_{t+1} = g(r_t), \qquad r_t = \Gamma_t, \qquad \sigma_t = \sqrt{r_t}.$$

In the baseline environment, r_t is the scalar price-based risk input constructed from recent transaction prices. The rule may also depend on background state variables z_t, such as current posted margin, price impact, and the local liquidation elasticity, but we suppress that dependence unless it matters for the argument. Absent manipulation concerns, the designer wants margin to rise with measured risk, because a higher value of r_t corresponds to a larger next-period loss distribution. This defines a first-best target rule. The impossibility theorem adds an implementability constraint. If g is too steep, a trader can move the sampled price used in the update, raise next-period margin, induce predictable forced liquidation, and profit from the resulting price pressure.

6.1 The design problem

Fix a state z and let $r \in \mathcal{I} = [\underline{r}, \overline{r}]$. Let $M_-(z)$ be the margin currently in place before the reset. We assume next-period close-out losses scale with measured risk:

$$D_{t+1} = \sqrt{r}\, Y_{t+1}, \tag{30}$$

where Y_{t+1} is a scale-free shock with conditional distribution $F_Y(\cdot \mid z)$. This covers the usual VaR and ES cases: absent manipulation concerns, the desired margin is increasing in r, typically with square-root scaling. For a candidate margin level M, define the objective

$$\Psi(r, M; z) = \mathbb{E}[\ell(D_{t+1} - M) \mid r, z] + \kappa_M M \tag{31}$$

$$+ \lambda_L C\Big(c_X^{\text{eff}}(z)\,(M - M_-(z))_+\Big) + \lambda_P H(1 - N(M, M_-(z); z))\,. \tag{32}$$

The first term is the coverage objective. It takes the expected value of a loss function applied to the shortfall $D_{t+1} - M$. When posted margin is low relative to the next-period close-out loss, this term is large; when margin is high enough to absorb the loss, it falls. The function ℓ is increasing and convex, so larger uncovered losses are penalized more than proportionally. This term captures the CCP's core objective: choosing margin to control expected exposure to default losses.

The second term is the carrying cost of margin. Higher margin protects the CCP, but it also ties up collateral, raises funding needs for clearing members, and makes central clearing more expensive. The coefficient κ_M measures the shadow cost of one additional unit of posted margin. Without such a term, the designer would have too strong an incentive to raise margin everywhere.

The third term captures the market-functioning cost created by an upward reset in margin. Only increases relative to the currently posted level $M_-(z)$ matter, which is why the term uses $(M - M_-(z))_+$. A higher reset forces constrained traders to post collateral or reduce positions. The factor $c_X^{\text{eff}}(z)$ translates the margin increase into predictable net selling pressure, and the function $C(\cdot)$ maps that forced flow into a cost borne by the clearing system, for example through price dislocation, impaired liquidity, or destabilizing fire-sale dynamics. The coefficient λ_L governs how much weight the designer places on these market-functioning losses.

The fourth term captures participation and migration effects. The object $N(M, M_-(z); z) \in [0, 1]$ is the fraction of activity that remains centrally cleared after the reset. If the new requirement is too high, too unstable, or too difficult to fund, some activity may leave the CCP, shrink, or migrate to less transparent venues. The term $1 - N(M, M_-(z); z)$ therefore measures the fraction of business lost, and $H(\cdot)$ converts that loss of participation into a welfare cost. The coefficient λ_P determines how strongly the designer values preserving cleared activity.

The designer then chooses a monotone rule g to minimize expected cost:

$$J(g; z) = \int_{\mathcal{I}} f(r \mid z)\, \Psi\big(r, g(r); z\big)\, dr. \tag{33}$$

6.2 The unconstrained first-best target

For each risk level r, define the first-best target as the margin level that minimizes the per-state objective:

$$M^{\text{fb}}(r; z) \in \underset{M \geq 0}{\arg\min}\ \Psi(r, M; z). \tag{34}$$

This is the rule the designer would choose if manipulation were not a concern. Under the maintained assumptions, $M^{\text{fb}}(r; z)$ is increasing in r. Higher measured risk shifts the loss distribution to the right, so the coverage benefit of additional margin is larger. In the simple benchmark with $\ell(x) = x_+$ and $C(x) = x^2/2$, any interior optimum satisfies

$$\Pr\Big(D_{t+1} > M^{\text{fb}}(r; z) \mid r, z\Big) = \kappa_M + \lambda_L \big(c_X^{\text{eff}}(z)\big)^2 \big(M^{\text{fb}}(r; z) - M_-(z)\big) + \lambda_P p\Big(M^{\text{fb}}(r; z); z\Big), \tag{35}$$

where

$$p(M; z) = -H'(1 - N(M, M_-(z); z))\ \partial_1 N(M, M_-(z); z) > 0. \tag{36}$$

Equation (35) has a simple interpretation. The left-hand side is the marginal benefit of raising margin: the probability that one more unit of margin prevents an uncovered loss. The right-hand side is the full marginal cost of doing so. The term κ_M is the direct burden of higher collateral. The term proportional to λ_L captures the extra liquidation pressure created by a tighter upward reset. The term proportional to λ_P captures the marginal loss of cleared activity when higher margin drives participation away.

When $\lambda_L = \lambda_P = 0$, the target collapses to the standard coverage rule:

$$M^{\text{fb}}(r; z) = \sqrt{r}\, Q_{1-\kappa_M}(Y_{t+1} \mid z)\,. \tag{37}$$

Thus, absent manipulation concerns, the desired rule is simply a conventional quantile-based margin schedule, increasing with volatility. The same square-root scaling arises under expected shortfall. More broadly, the first-best target is steeper when the designer puts more weight on coverage and flatter when higher margin creates larger liquidation costs or larger participation losses. This is the benchmark the designer would like to implement before imposing the no-manipulation constraint.

6.3 The implementable rule

The first-best target is not obviously implementable. Corollary 1 implies that if the rule passes measured risk into next-period margin too aggressively, a trader can profit from a trigger-and-reverse deviation. Under $\kappa = 0$, the local no-manipulation condition at a binding state is

$$g'(r_t) \leq \frac{1}{\alpha_t c^{\text{eff}}_{X,t} |u_t|}, \tag{38}$$

where u_t is the last sampled return entering the risk statistic. Since $|u_t| \leq \sigma_t = \sqrt{r_t}$, a convenient sufficient condition on the slope is

$$0 \leq g'(r) \leq \bar{s}(r; z), \qquad \bar{s}(r; z) = \frac{1}{\alpha(z) c^{\text{eff}}_X(z) \sqrt{r}}. \tag{39}$$

Writing $\widetilde{g}(\sigma; z) = g(\sigma^2; z)$, the same restriction becomes

$$0 \leq \widetilde{g}'(\sigma; z) \leq \bar{s}_\sigma(z), \qquad \bar{s}_\sigma(z) = \frac{2}{\alpha(z) c^{\text{eff}}_X(z)}. \tag{40}$$

In volatility units, the maximal safe pass-through is a constant. In variance units, the cap is tighter when r is low. Intuitively, in tranquil states a given price move produces a larger proportional change in the sampled risk measure, so an aggressive update rule is easier to manipulate.

The designer therefore solves the original problem choosing an implementable margin rule that satisfies the slope cap:

$$\min_{g \in \mathcal{G}_{\text{IC}}(z)} J(g; z). \tag{41}$$

with $\mathcal{G}_{\text{IC}}(z) = \{g \text{ is weakly increasing and } 0 \leq g'(r) \leq \bar{s}(r; z)\}$. When the first-best target is already feasible, the implementable rule coincides with it. Where the first-best is too steep, the rule rises at the maximal manipulation-proof slope. Where even that capped response would be too aggressive over a range of states, the optimal policy pools and keeps margin constant across those states. Thus the implementable rule is the closest monotone approximation to the first-best target that respects the no-manipulation bound.[17]

On any interval over which the slope cap binds, the rule takes the form

$$M^*(r;z) = a(z) + \frac{2}{\alpha(z) c_X^{\text{eff}}(z)} \sqrt{r}, \tag{42}$$

with the constant $a(z)$ chosen to match the adjoining parts of the schedule. On any pooling interval $[a, b]$, the common level k satisfies

$$\int_a^b f(r \mid z)\, \partial_M \Psi(r, k; z)\, dr = 0. \tag{43}$$

The designer would like margin to respond strongly to higher measured risk, but cannot allow that response to be arbitrarily steep. The no-manipulation constraint therefore turns a standard quantile-style target into a capped rule: it tracks the target when possible, flattens when necessary, and may pool nearby risk states when the cap is especially tight.

6.4 Quantitative interpretation of the slope cap

For interpretation, the most useful version of the cap is in margin-ratio and annualized-volatility units. Let $m = M/P$ denote the margin ratio and let $\sigma^{ann} = \sqrt{252}\,\sigma$. Appendix D shows that the local no-manipulation restriction implies the bound

$$\frac{\Delta m}{m} \le \frac{2f}{\sqrt{252}\,\epsilon_{X,M} \tau \delta(f)}\, \Delta\sigma^{ann}, \tag{44}$$

where $\epsilon_{X,M}$ is the elasticity of the constrained sector's position with respect to margin, $\tau = XP/ADV$ measures the size of the constrained sector relative to market depth, $f = qP/ADV$ is the trigger trade as a fraction of average daily traded dollar volume (ADV), and $\delta(f)$ is the sampled-price distortion generated by such a trade. Equation (44) shows that the admissible responsiveness of the rule is smaller when the constrained sector is larger relative to market depth, when positions are more margin-sensitive, and when small trades move the sampled price more sharply. Using the benchmark calibration from Section 3.6, i.e., with $\epsilon_{X,M} = 0.225$, $f = 1\%$, $\delta_{bps}(1\%) \approx 20$, the bound becomes

$$\frac{\Delta m}{m} \le \frac{2.80}{\tau}\, \Delta\sigma^{ann}, \tag{45}$$

with $\Delta\sigma^{ann}$ measured in decimal annualized-volatility units. When $\tau = 1$, so that the constrained sector's notional is equal to one day of average daily traded value, even a large ten-point increase in annualized volatility, for example from 10% to 20%, permits at most a 28% increase in the margin ratio. Starting from 5%, the largest manipulation-proof increase is therefore only to about 6.4%. When $\tau = 2$, the same volatility shock permits only a 14% increase, so a 5% margin can rise only to about 5.7%. Thus, once the constrained sector is of the same order as market depth, or larger, even a substantial increase in measured risk supports only a modest increase in posted margin. The converse implication is just as important. A rule that appears only moderately reactive can already

[17] All the derivations are in Appendix D. Intuitively these are just standard KKT conditions.

violate the cap. For example, raising margin from 5% to 6% is a 20% relative increase. Under the benchmark calibration, that increase is manipulation-proof only if annualized volatility rises by at least about 7 percentage points when $\tau = 1$, and by at least about 14 percentage points when $\tau = 2$. If the same jump from 5% to 6% is triggered by a smaller volatility increase, the rule is too steep relative to the no-manipulation bound and admits a profitable trigger-and-reverse trade.

Margin can still respond to risk, but only within a tight quantitative window. Once τ is moderate or large, even sizeable increases in measured volatility justify only modest upward resets in margin. This creates a genuine and economically important design trade-off: the very feature that makes a rule prudent in the standard risk-management sense, namely strong short-run responsiveness to measured risk, is precisely what makes it manipulable. Theorem 1 therefore identifies a novel, empirically relevant, tension between coverage and implementability. A designer cannot, in general, have both a sharply reactive margin rule and local incentive compatibility in binding states.

6.5 Portfolio margining

The same design logic extends to portfolio margining, with two caveats. First, an upward reset in margin generates a vector of forced sales rather than a single-asset liquidation. Let B_t denote the liquidation vector induced by a one-unit increase in posted margin, and let A_t be the cross-impact matrix from Section 5. The resulting liquidation-driven price pressure is $A_t B_t$, so the relevant market-functioning scale is $\|A_t B_t\|_2$. A larger value means that any given increase in margin causes more disruptive price pressure, which makes the unconstrained target rule flatter. Second, the manipulation constraint depends on the relation between the portfolio used to move the scalar risk input and the portfolio that constrained traders later liquidate. Let d_t denote the direction that increases the risk input r_t, and let $c_{r,t}$ be the sensitivity of r_t in that direction. Define $\psi_t = \max\{0, -d_t^\top A_t B_t\}$. The portfolio analogue of the slope cap is then

$$g'(r_t) \leq \frac{2}{c_{r,t}\psi_t}. \tag{46}$$

If the contracts that move the risk input are well aligned with the contracts that will later be liquidated, then a small trigger trade can create a large predictable harvest, and the admissible slope must be low. Netting may lower the desired level of margin, but cross-impact and trigger-liquidation alignment can make the public rule less responsive. If the contracts that are most effective at moving the scalar risk input are also well aligned, through cross-impact, with the contracts that constrained traders later liquidate, then the admissible slope can be tighter than in the single-asset benchmark. Appendix D.6 gives a simple two-asset example, calibrated to the single-asset benchmark, showing that portfolio margining can either loosen or tighten the cap depending on the strength of cross-impact and trigger-liquidation alignment.

6.6 Jumps, cliffs, and non-smooth rules

Replacing a steep slope with a jump does not remove the incentive problem. It changes smooth local manipulation into threshold-crossing manipulation. Suppose g has an upward jump of size

$\Delta M > 0$ at threshold $r^\dagger$, and suppose the current state lies just below that threshold. If a sufficiently small admissible trade can push the sampled input above $r^\dagger$, the trader triggers a discrete increase in next-period margin and therefore a discrete increase in forced liquidation. In the single-asset case, a threshold-crossing trade of size q generates an expected harvesting gain approximately equal to

$$q\,\alpha_t c^{\text{eff}}_{X,t}\Delta M,$$

while the trader's own round-trip impact cost is proportional to q^2. In the portfolio case, the corresponding harvesting gain is approximately

$$q\,u^\top A_t B_t\,\Delta M$$

for any direction u that loads positively on the liquidation price-pressure vector, while the direct round-trip cost is again proportional to q^2. Hence, for sufficiently small threshold-crossing trades, the gain from triggering the jump dominates the trader's own impact loss.

7 Conclusions

This paper studies price-based risk constraints: public rules that map sampled transaction prices into a risk statistic and then into a binding requirement, $M_{t+1} = g(\Gamma_t)$. Once the rule binds in an impact-sensitive market, it becomes part of the trading environment.

Our main result is a constructive local impossibility theorem. At reachable binding states, a price-based rule cannot in general deliver local risk sensitivity, liquidity continuity, and round-trip manipulation-proofness at the same time (Theorem 1). The mechanism does not rely on large trades, on making a slack constraint bind, on predictable fundamentals, or on standard forms of manipulation. The profit opportunity is created by the rule itself. A small trigger trade can shift the next requirement at first order and then be reversed into the forced flow generated by that tighter requirement. Profitability is governed by a simple amplification chain from sampled prices to requirements, from requirements to forced sales, and from forced sales to prices.

Read this way, the theorem is also a design result. Any binding price-based rule that seeks to remain manipulation-proof must limit short-horizon pass-through from sampled prices into posted requirements. This creates a first-order trade-off between coverage and implementability. A rule that responds more sharply to measured risk provides more protection against loss, but it is also easier to exploit in precisely the states where it binds. The implied slope cap therefore gives a structural rationale for bounded pass-through, smoothing, buffers, and related anti-procyclicality tools. These are not only statistical devices. They are also incentive-compatibility devices. With volatility-based inputs, the restriction can be tightest after quiet periods, so the states that look safest under the rule can be the most locally vulnerable.

The same logic becomes stronger under portfolio margining. When a scalar requirement is computed from a vector of prices, the contracts that move the portfolio risk measure need not be the contracts sold when the tighter call binds. This trigger-harvest wedge creates a genuinely cross-contract channel: a trader can move the portfolio statistic in one set of instruments and profit from forced liquidation in another. Portfolio margining can therefore shift stress across contracts and markets, so trading pressure during the sampling window and post-update illiquidity need not appear in the same place.

More broadly, the paper identifies a general implementability constraint for any institution that ties binding decisions mechanically to risk measures built from recent prices. CCP margining is a central application, but the same issue arises whenever recent prices determine binding leverage, exposure, or capital-allocation rules. Once a public rulebook makes tradable prices part of a binding constraint, prudent risk management and manipulation-proof design are no longer aligned. The rule must be designed jointly with the market impact it creates.

References

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

_ and _ , "Procyclical Leverage and Value-at-Risk," *The Review of Financial Studies*, 2014, *27* (2), 373–403.

Aldasoro, Inaki, Torsten Ehlers, and Egemen Eren, "Liquid assets at CCPs and systemic liquidity risks," *BIS Quarterly Review*, December 2023.

Allen, Franklin and Douglas Gale, "Stock-price manipulation," *Review of Financial Studies*, 1992, *5* (3), 503–529.

Avalos, Fernando, Wenqian Huang, and Kevin Tracol, "Margins and Liquidity in European Energy Markets in 2022," BIS Bulletin 77, Bank for International Settlements September 2023.

Bank of England, "A CBA of APC: analysing approaches to procyclicality reduction in CCP initial margin models," Staff Working Paper 950, Bank of England 2021.

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.

Basel Committee on Banking Supervision, "Transparency and responsiveness of initial margin in centrally cleared markets: review and policy proposals," Technical Report d568, Bank for International Settlements 2022.

Basel Committee on Banking Supervision and Committee on Payments and Market Infrastructures and Board of the International Organization of Securities Commissions, "Review of Margining Practices," Technical Report, Bank for International Settlements and IOSCO September 2022. Available at: https://www.iosco.org/library/pubdocs/pdf/IOSCOPD714.pdf.

Biais, Bruno, Florian Heider, and Marie Hoerova, "Risk-sharing or risk-taking? Counterparty risk, incentives, and margins," *Journal of Finance*, 2016, *71* (4), 1669–1698.

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.

Brunnermeier, Markus K. and Lasse Heje Pedersen, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 2009, *22* (6), 2201–2238.

_ and Yuliy Sannikov, "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 2014, *104* (2), 379–421.

Chebotarev, Dmitry, "Adverse Selection Effect of CCP Haircuts," March 2025. Working paper, Indiana University Kelley School of Business.

CME Group, "CME SPAN Methodology Overview," 2019.

_ , "SPAN Methodology," Technical Report, CME Group March 2019.

_ , "SPAN Reference Documents," 2026.

Committee on Payments and Market Infrastructures and International Organization of Securities Commissions, "Principles for Financial Market Infrastructures," Technical Report, Bank for International Settlements and IOSCO April 2012.

Cont, Rama and Seyyedmostafa Ghamami, "Skin in the game: risk analysis of central counterparties," *Journal of Financial Market Infrastructures*, 2025.

_ and Thomas Kokholm, "Central clearing of OTC derivatives: Bilateral vs multilateral netting," *Statistics and Risk Modeling*, 2014, *31* (1), 3–22.

_ , Romain Deguest, and Giacomo Scandolo, "Robustness and sensitivity of risk measures," *Quantitative Finance*, 2010, *10* (6), 593–606.

Coval, Joshua and Erik Stafford, "Asset fire sales (and purchases) in equity markets," *Journal of Financial Economics*, 2007, *86* (2), 479–512.

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

Duarte, Fernando and Thomas M. Eisenbach, "Fire-Sale Spillovers and Systemic Risk," Staff Reports 645, Federal Reserve Bank of New York October 2013.

Duffie, Darrell and Haoxiang Zhu, "Does a central clearing counterparty reduce counterparty risk?," *Review of Asset Pricing Studies*, 2011, *1* (1), 74–95.

_ , Martin Scheicher, and Guillaume Vuillemey, "Central Clearing and Collateral Demand," *Journal of Financial Economics*, 2015, *116* (2), 237–256.

European Central Bank, "CCP initial margin models in Europe," Occasional Paper Series 314, European Central Bank April 2023.

European Securities and Markets Authority, "Final report on guidelines on CCP anti-procyclicality margin measures (ESMA70-151-1293)," Technical Report, ESMA 2018.

_ , "Consultation paper on review of EMIR RTS with respect to procyclicality of CCP margin," Technical Report, ESMA 2022.

_ , "Final report: review of RTS No 153/2013 with respect to procyclicality of CCP margin," Technical Report, ESMA 2023.

European Systemic Risk Board, "Mitigating the procyclicality of margins and haircuts in derivatives markets and securities financing transactions," Technical Report, ESRB January 2020.

European Union, "Commission Delegated Regulation (EU) No 153/2013, Article 28 (Procyclicality)," 2013.

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz, "Trading Costs," Technical Report, AQR Capital Management 2018. Working paper, August 23, 2018.

Gârleanu, Nicolae and Lasse Heje Pedersen, "Margin-Based Asset Pricing and Deviations from the Law of One Price," *The Review of Financial Studies*, 2011, *24* (6), 1980–2022.

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, pp. 1–65.

Glasserman, Paul and Qi Wu, "Persistence and Procyclicality in Margin Requirements," *Management Science*, 2018, *64* (12), 5705–5724.

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

Gromb, Denis and Dimitri Vayanos, "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs," *Journal of Financial Economics*, 2002, *66* (2-3), 361–407.

Gurrola-Perez, Pedro, "Procyclicality of CCP margin models: systemic problems need systemic approaches," WFE Working Paper, World Federation of Exchanges December 2020. Available at SSRN: 3779896.

He, Zhiguo and Arvind Krishnamurthy, "Intermediary Asset Pricing," *American Economic Review*, 2013, *103* (2), 732–770.

Hedegaard, Esben, "Causes and Consequences of Margin Levels in Futures Markets," Working Paper, Arizona State University 2014.

Huang, Runhua and Elod Takáts, "Model risk at central counterparties: Is skin in the game a game changer?," *International Journal of Central Banking*, 2024, *20* (3).

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

Jarrow, Robert A., "Market manipulation, bubbles, corners, and short squeezes," *Journal of Financial and Quantitative Analysis*, 1992, *27* (3), 311–336.

Julliard, Christian, Gabor Pinter, Karamfil Todorov, Jean-Charles Wijnandts, and Kathy Yuan, "What Drives Repo Haircuts? Evidence from the UK Market," BIS Working Papers 1027, Bank for International Settlements November 2024. July 2022, revised November 2024.

King, Thomas B., Travis D. Nesmith, Anna Paulson, and Todd Prono, "Central clearing and systemic liquidity risk," *International Journal of Central Banking*, 2023, *19* (4).

Kiyotaki, Nobuhiro and John Moore, "Credit Cycles," *Journal of Political Economy*, 1997, *105* (2), 211–248.

Kuong, John Chi-Fong and Vincent Maurin, "The design of a central counterparty," *Journal of Financial and Quantitative Analysis*, 2024, *59* (3), 1257–1299.

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

LCH, "LCH Ltd margin methodology," 2026.

LCH Limited, "Procedures Section 2C: SwapClear Clearing Service," Rulebook / Procedures December 2025.

LCH Ltd, "CPMI-IOSCO Self-Qualitative Assessment of Observance of the PFMI of LCH Ltd (Q3 2022)," 2022.

LCH.Clearnet Limited, "4d(f) Request (CFTC) (describes PAIRS and volatility scaling via EWMA)," Technical Report, U.S. Commodity Futures Trading Commission November 2015.

London Metal Exchange, "LME Clear SPAN Methodology (Version 1.1)," Technical Report, LME Clear August 2015.

_ , "LME Clear margin parameter files," 2026.

Loon, Yee Cheng and Zhaodong (Ken) K. Zhong, "The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market," *Journal of Financial Economics*, 2014, *112* (1), 91–115.

Menkveld, Albert J. and Guillaume Vuillemey, "The economics of central clearing," *Annual Review of Financial Economics*, 2021, *13*, 153–178.

Murphy, David, Michalis Vasios, and Nicholas Vause, "An investigation into the procyclicality of risk-based initial margin models," Financial Stability Paper 29, Bank of England 2014.

_ , _ , and _ , "A comparative analysis of tools to limit the procyclicality of initial margin requirements," Working Paper 597, Bank of England 2016.

Nicolai, Francesco and Simona Risteska, "Dynamic Arbitrage from Price-Based Risk Constraints," Technical Report, Working paper 2026. Working paper (February 2026).

Pinter, Gabor, "An Anatomy of the 2022 Gilt Market Crisis," Staff Working Paper 1019, Bank of England March 2023.

Pirrong, Craig, "The Economics of Central Clearing: Theory and Practice," ISDA Discussion Papers Series, International Swaps and Derivatives Association 2011.

—, "A Bill of Goods: CCPs and Systemic Risk," *Journal of Financial Market Infrastructures*, 2014, *2* (3), 55–89.

Reserve Bank of Australia, "Assessment of LCH Limited's SwapClear Service: Appendix B (PAIRS description and procyclicality floor)," Technical Report, Reserve Bank of Australia 2019.

Schneider, Michael and Fabrizio Lillo, "Cross-Impact and No-Dynamic-Arbitrage," *Quantitative Finance*, 2019, *19* (1), 137–154.

Shleifer, Andrei and Robert W. Vishny, "Fire Sales in Finance and Macroeconomics," *Journal of Economic Perspectives*, 2011, *25* (1), 29–48.

The Options Clearing Corporation, "Margin Methodology," 2026.

U.S. Securities and Exchange Commission, "Release No. 34-99393: Notice of Filing of Proposed Rule Change by The Options Clearing Corporation," January 2024.

Wang, Jessie Jiaxu, Agostino Capponi, and Hongzhong Zhang, "A Theory of Collateral Requirements for Central Counterparties," *Management Science*, 2022, *68* (9), 6993–7017.

Internet Appendix

for

An Impossibility Theorem for Price-Based Risk Constraints

A Solving the general case

For simplicity and clarity, the main text specializes to realized variance and linear impact. The mechanism itself is local and does not require differentiability. It only requires that an admissible perturbation of the sampled transaction price at t that raises the risk input also raises posted margin by a nontrivial amount; that, in states where margins bind, this tightening induces a nontrivial adjustment in positions; and that the induced forced order flow moves transaction prices at first order. This Appendix states these local conditions without derivatives and then reruns the same two-step construction around a benchmark path with no deviation at time t. We write $(\cdot)^{(0)}$ for benchmark objects (prices, inputs, and margins) and compare them to outcomes under a small two-period round trip.

A.1 One-sided slopes

Definition 5 (One-sided strong increase)**.** A function $f : \mathbb{R} \to \mathbb{R}$ is one-sided strongly increasing at x_0 with constant $c_f > 0$ if there exists $\delta_f > 0$ such that for all $h \in [0, \delta_f]$,

$$f(x_0 + h) - f(x_0) \geq c_f h. \tag{47}$$

Definition 6 (Directional sensitivity of the risk input in the last price)**.** Let $\Gamma : \mathbb{R}^{m+1} \to \mathbb{R}$ and $\Gamma_t = \Gamma(P_{t-m}, \dots, P_t)$. Fix $(P_{t-m}, \dots, P_{t-1})$ and define $\Gamma^{\text{last}}(p) = \Gamma(P_{t-m}, \dots, P_{t-1}, p)$. We say that Γ is directionally sensitive at p_0 if there exist $s \in \{-1, +1\}$ and constants $c_\Gamma > 0$ and $\delta_\Gamma > 0$ such that for all $h \in [0, \delta_\Gamma]$,

$$\Gamma^{\text{last}}(p_0 + sh) - \Gamma^{\text{last}}(p_0) \geq c_\Gamma h. \tag{48}$$

Definition 5 replaces a right-derivative bounded away from zero with a one-sided linear lower bound. Definition 6 is the only step where the realized-variance structure in the main text mattered: it ensures that an admissible perturbation of the last sampled mark can raise the risk input at first order; conditional on that, the remainder of the construction is unchanged.

A.2 Local assumptions

Fix a reachable benchmark configuration at time t and interpret all statements below locally at that benchmark. Let $P_t^{(0)}$ denote the benchmark transaction price at t absent a deviation, and let $\Gamma_t^{(0)}$ denote the corresponding risk input.

Assumption 11 (Risk input is directionally sensitive)**.** The risk input Γ is directionally sensitive at the benchmark last mark $P_t^{(0)}$ in the sense of Definition 6. Moreover, $\Gamma_t^{(0)}$ lies in the interior of an interval I_Γ on which the posted margin rule is evaluated.

Assumption 12 (Posted margin has a one-sided slope bound)**.** The posted margin rule is $M_{t+1} = g(\Gamma_t)$, where $g : \mathbb{R} \to (0, \infty)$ is weakly increasing. There exists $c_g > 0$ such that for all $x, y \in I_\Gamma$ with $y \geq x$,

$$g(y) - g(x) \geq c_g(y - x). \tag{49}$$

Assumption 13 (Forced liquidation has a local lower bound)**.** There exist constants $c_X > 0$ and $\delta_M > 0$ such that for any margin increase $\Delta M \in [0, \delta_M]$, the induced reduction in aggregate constrained position satisfies

$$L(\Delta M) = X(M_t) - X(M_t + \Delta M) \geq c_X \Delta M. \tag{50}$$

Assumption 14 (Generic price impact)**.** Transaction prices satisfy

$$P_u = V_u + \Psi_u(Q_u), \qquad Q_u = \Delta X_u + a_u, \tag{51}$$

where $\Psi_u(0) = 0$ and Ψ_u is odd and weakly increasing. There exist constants $0 < \underline{\alpha}_u \leq \overline{\alpha}_u < \infty$ and $\delta_\Psi > 0$ such that for all $q \in [0, \delta_\Psi]$,

$$\underline{\alpha}_u q \leq \Psi_u(q) \leq \overline{\alpha}_u q. \tag{52}$$

By oddness, (52) also implies $-\overline{\alpha}_u|q| \leq \Psi_u(q) \leq -\underline{\alpha}_u|q|$ for $q \in [-\delta_\Psi, 0]$. Assumption 14 contains linear impact as a special case and covers any monotone impact schedule with local slopes bounded away from zero and infinity on the relevant flow range.

A.2.1 Broader constrained demand schedules and margin elasticity

This subsubsection gives a sufficient condition for Assumption 13. The main text uses a target-with-cap rule to make margin elasticity explicit; the same local slope obtains for a broad class of truncated demands whenever the margin cap binds for a positive mass of traders with nontrivial total equity.

Fix time t and a benchmark state. Let $\tilde{x}_t(i)$ denote trader i's unconstrained desired position and let $E(i)$ denote equity. The implemented position under margin level $M > 0$ is

$$x_t(i; M) = \operatorname{sgn}(\tilde{x}_t(i)) \min\left\{|\tilde{x}_t(i)|, \frac{E(i)}{M}\right\}. \tag{53}$$

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

Lemma 11 (Sufficient condition for Assumption 13)**.** *Fix a benchmark margin $M_t > 0$ and let $B \subseteq [0, 1]$ be measurable such that*

$$|\tilde{x}_t(i)| > \frac{E(i)}{M_t} \text{ for all } i \in B, \qquad \int_B E(i)\, di > 0, \tag{54}$$

and $\operatorname{sgn}(\tilde{x}_t(i))$ is constant on B. Then for any $\delta_M > 0$ and all $\Delta M \in [0, \delta_M]$,

$$X(M_t) - X(M_t + \Delta M) \geq \left(\frac{\int_B E(i)\, di}{(M_t + \delta_M)^2}\right) \Delta M. \tag{55}$$

In particular, Assumption 13 holds with $c_X = \dfrac{\int_B E(i)\, di}{(M_t + \delta_M)^2}$.

Proof. Fix $\Delta M \in [0, \delta_M]$ and write $M' = M_t + \Delta M$. For $i \in B$, the cap binds at M_t by (54), and since $M' \geq M_t$ it also binds at M'. Hence

$$x_t(i; M_t) = \operatorname{sgn}(\tilde{x}_t(i)) \frac{E(i)}{M_t}, \qquad x_t(i; M') = \operatorname{sgn}(\tilde{x}_t(i)) \frac{E(i)}{M'}.$$

With constant sign on B,

$$x_t(i; M_t) - x_t(i; M') = \frac{E(i)\Delta M}{M_t M'}.$$

Integrating over B and using $M_t M' \le (M_t + \delta_M)^2$ yields (55). □

Demand (53) arises from any concave problem with a unique unconstrained maximizer $\tilde{x}_t(i)$ and a margin constraint $|x|M \le E(i)$. Lemma 11 shows that if the cap binds for a positive-mass set with nontrivial total equity, then aggregate positions respond to margin with a strictly positive local slope. The bound increases with $\int_B E(i)\,di$ and decreases with the relevant margin level M_t.

A.2.2 A two-step round trip in the general case

To keep the profit bound clean under a fully nonparametric Ψ_{t+1}, we impose the same local benchmark configuration as in the main text.

Assumption 15. At the benchmark S_t: (i) $\Delta X_t = 0$; (ii) letting ΔX^0_{t+1} denote constrained order flow at $t+1$ absent a deviation at t, we have $\Delta X^0_{t+1} = 0$; (iii) $\mathbb{E}[V_{t+1} \mid S_t] = V_t$.

Lemma 12 (A profitable two-period round trip under non-smooth primitives)**.** *Suppose Assumptions 11–15 hold at a reachable benchmark state S_t. For concreteness, assume the risk input increases when the last sampled mark is pushed down (in Definition 6, $s = -1$).*[18]

Let $\bar{q}_t > 0$ be such that for every $q \in (0, \bar{q}_t]$, the round trip $(a_t, a_{t+1}) = (-q, q)$ keeps all local restrictions in force: (i) $\Psi_t(q) \le \delta_\Gamma$; (ii) the perturbed input remains in I_Γ; (iii) the induced margin change $\Delta M_{t+1} = g(\Gamma_t) - g(\Gamma_t^{(0)})$ lies in $[0, \delta_M]$; (iv) the induced liquidation $L = X(M_t) - X(M_t + \Delta M_{t+1})$ satisfies $L - q \in [0, \delta_\Psi]$; and (v) when imposed, the hard funding constraint (12) holds (equivalently, $q \le \bar{F}/M_t$). Define

$$q_{\max,t} = \min\{\bar{a}, \delta_\Psi, \bar{q}_t\}.$$

Fix any $q \in (0, q_{\max,t}]$ and consider $a_t = -q$, $a_{t+1} = q$. Then

$$\Delta\Gamma_t = \Gamma_t - \Gamma_t^{(0)} \ge c_\Gamma \underline{\alpha}_t q, \qquad \Delta M_{t+1} \ge c_g c_\Gamma \underline{\alpha}_t q, \qquad L \ge c_X c_g c_\Gamma \underline{\alpha}_t q.$$

If, in addition, the gain condition

$$\underline{\alpha}_{t+1}\, c_X c_g c_\Gamma\, \underline{\alpha}_t > \overline{\alpha}_t + \underline{\alpha}_{t+1} \tag{56}$$

holds, then the conditional expected profit satisfies

$$\mathbb{E}[\Pi_1 \mid S_t] \ge \Big(\underline{\alpha}_{t+1} c_X c_g c_\Gamma \underline{\alpha}_t - (\overline{\alpha}_t + \underline{\alpha}_{t+1})\Big) q^2 - 2\tau q. \tag{57}$$

Let

$$A = \underline{\alpha}_{t+1} c_X c_g c_\Gamma \underline{\alpha}_t - (\overline{\alpha}_t + \underline{\alpha}_{t+1}), \qquad q_{\min} = \frac{2\tau}{A} \text{ when } A > 0.$$

Under (56), $A > 0$, and if $q_{\max,t} > q_{\min}$ there exists an admissible $q \in (q_{\min}, q_{\max,t}]$ such that $\mathbb{E}[\Pi_1 \mid S_t] > 0$.

[18] If $s = +1$, reverse the signs of the round trip and use the buy-in analogue of Assumption 13; see the sign-reversal discussion following Theorem 1.

Proof. Under Assumption 15(i), $Q_t = a_t = -q$, so

$$P_t = V_t + \Psi_t(-q) = V_t - \Psi_t(q), \qquad P_t^{(0)} = V_t,$$

and by (52), $\Psi_t(q) \geq \underline{\alpha}_t q$. With $s = -1$, Assumption 11 yields $\Delta\Gamma_t \geq c_\Gamma \Psi_t(q) \geq c_\Gamma \underline{\alpha}_t q$. Since $\Gamma_t, \Gamma_t^{(0)} \in I_\Gamma$ and $\Gamma_t \geq \Gamma_t^{(0)}$, Assumption 12 implies $\Delta M_{t+1} \geq c_g \Delta\Gamma_t \geq c_g c_\Gamma \underline{\alpha}_t q$, and then Assumption 13 gives $L \geq c_X \Delta M_{t+1} \geq c_X c_g c_\Gamma \underline{\alpha}_t q$. Under Assumption 15(ii), the constrained sector has zero net flow at $t+1$ absent the deviation; with the deviation it sells L while the deviator buys q, so $Q_{t+1} = q - L$ and $P_{t+1} = V_{t+1} + \Psi_{t+1}(q - L)$. Condition (56) implies $c_X c_g c_\Gamma \underline{\alpha}_t > 1$, hence $L > q$ and $q - L < 0$. By the definition of $\bar{q}_t$, $L - q \in [0, \delta_\Psi]$, so oddness and (52) give

$$P_{t+1} = V_{t+1} - \Psi_{t+1}(L - q) \leq V_{t+1} - \underline{\alpha}_{t+1}(L - q), \qquad P_t \geq V_t - \overline{\alpha}_t q.$$

Therefore

$$P_t - P_{t+1} \geq (V_t - V_{t+1}) - \overline{\alpha}_t q + \underline{\alpha}_{t+1}(L - q).$$

Under the round trip, $\Pi_1 = q(P_t - P_{t+1}) - 2\tau q$, so taking conditional expectations and using Assumption 15(iii) yields

$$\mathbb{E}[\Pi_1 \mid S_t] \geq q\Big(- \overline{\alpha}_t q + \underline{\alpha}_{t+1}(L - q)\Big) - 2\tau q.$$

Substituting $L \geq c_X c_g c_\Gamma \underline{\alpha}_t q$ gives (57). Under (56), $A > 0$ and the lower bound equals $Aq^2 - 2\tau q = q(Aq - 2\tau)$, which is positive whenever $q > 2\tau/A$. If $q_{\max,t} > q_{\min} = 2\tau/A$, pick any $q \in (q_{\min}, q_{\max,t}]$. □

Assumption 15(ii) is a convenience: it avoids imposing structure on Ψ_{t+1} beyond local slope bounds. With mild local regularity of Ψ_{t+1} on the support of baseline flow (for example local affinity on that range), a weaker mean-zero condition such as $\mathbb{E}[\Delta X_{t+1}^0 \mid S_t] = 0$ suffices, as in the main text.

A.2.3 The general impossibility statement

Lemma 12 gives a lower bound on the conditional expected profit of a two-period trigger-and-reverse deviation under local one-sided conditions. If the gain condition (56) holds and the admissible local range contains some q above the execution-cost threshold $q_{\min}$, the bound is strictly positive and yields a profitable round trip.

Theorem 5 (General impossibility theorem)**.** *Fix a Markov equilibrium with price formation* (51) *and posted margin rule $M_{t+1} = g(\Gamma_t)$. Suppose there exists a reachable benchmark state S_t at which liquidity continuity holds (Definition 4) and Assumptions 11–15 hold. Let $q_{\max,t}$ and $q_{\min}$ be defined as in Lemma 12. If* (56) *holds and $q_{\max,t} > q_{\min}$, then the mechanism is not round-trip manipulation-proof at S_t in the sense of Definition 3 with $\rho = \kappa = 0$: there exists an admissible two-period round trip with $\mathbb{E}[\Pi_1 \mid S_t] > 0$.*

Consequently, at S_t the following cannot all hold simultaneously: one-sided risk sensitivity in the sense of Assumption 12, liquidity continuity (Definition 4), and round-trip manipulation-proofness net of execution costs (Definition 3 with $\rho = \kappa = 0$). Equivalently, at any reachable S_t with liquidity continuity, round-trip manipulation-proofness implies that at least one of Assumptions 11–15, the gain inequality (56)*, or the feasibility condition $q_{\max,t} > q_{\min}$ fails.*

Proof. Under the stated hypotheses, Lemma 12 implies that any $q \in (q_{\min}, q_{\max,t}]$ yields an admissible two-period round trip $(a_t, a_{t+1}) = (-q, q)$ with $\mathbb{E}[\Pi_1 \mid S_t] > 0$, contradicting round-trip manipulation-proofness at S_t (net of execution costs). □

A.3 Verifying directional sensitivity for standard CCP-style inputs

Assumption 11 is intentionally weak. We do not model full CCP implementation details; we only give simple sufficient conditions ensuring that a risk input computed from recent marks is directionally sensitive in the last sampled mark.

A.3.1 Smooth functionals of recent returns

Let $R_j = P_j - P_{j-1}$ and suppose the input is a smooth function of the most recent m returns,

$$\Gamma_t = \phi(R_{t-m+1}, \ldots, R_t),$$

for a continuously differentiable map ϕ defined on an open neighborhood of the realized return vector.

Lemma 13. *If ϕ is continuously differentiable near the realized return vector and $\partial\phi/\partial R_t \neq 0$ at the realized point, then Γ is directionally sensitive at the benchmark last mark in the sense of Definition 6.*

Proof. Hold $(P_{t-m}, \ldots, P_{t-1})$ fixed and perturb only the last mark: $P_t = p_0 \mapsto p_0 + h$. Then only the last return changes, $R_t \mapsto R_t + h$. Define $u(h) = \phi(R_{t-m+1}, \ldots, R_{t-1}, R_t + h)$, so $\Gamma^{\text{last}}(p_0 + h) = u(h)$. By differentiability, $u'(0) = \partial\phi/\partial R_t \neq 0$. Let $s = \operatorname{sgn}(u'(0))$. Continuity of u' implies that for some $\delta > 0$, $u'(sh)\, s \geq |u'(0)|/2$ for all $h \in [0, \delta]$. The mean value theorem then gives $u(sh) - u(0) = u'(\xi)\, sh \geq (|u'(0)|/2)h$ for some $\xi \in (0, sh)$. This is Definition 6 with $c_\Gamma = |u'(0)|/2$ and $\delta_\Gamma = \delta$. □

Realized variance is the special case $\phi(r) = \sum_j r_j^2$.

A.3.2 Parametric VaR/ES built from volatility estimates

A common parametric form is $\Gamma_t = k_p \hat{\sigma}_t$ where $\hat{\sigma}_t = h(R_{t-m+1}, \ldots, R_t)$ and $k_p > 0$ is a fixed multiplier.

Lemma 14. *If h is continuously differentiable near the realized return vector and $\partial h/\partial R_t \neq 0$ at that point, then $\Gamma_t = k_p \hat{\sigma}_t$ is directionally sensitive at the benchmark last mark in the sense of Definition 6.*

Proof. Apply Lemma 13 with $\phi = k_p h$. Since $k_p > 0$, $\partial\phi/\partial R_t = k_p(\partial h/\partial R_t) \neq 0$, so directional sensitivity follows. □

A.4 Transient impact and longer horizons

The two-period round trip used in the main argument is the shortest horizon that can (i) perturb the sampled marks that enter the next update and then (ii) trade against the forced liquidation induced by the tighter requirement. Two extensions clarify what the proof does and does not rely on.

First, the mechanism does not require permanent impact. Allow P_{t+1} to depend on the history of total order flow $(Q_{t+1}, Q_t, \ldots)$, but assume that the contemporaneous mapping $Q_{t+1} \mapsto P_{t+1}$ has a strictly positive local one-sided slope around the realized flow at $t+1$, in the sense of Assumption 14. Then the proof of Lemma 12 goes through after replacing $\underline{\alpha}_{t+1}$ by this instantaneous local slope. Additional dependence of P_{t+1} on earlier flows enters as extra terms; these do not remove the rule-induced predictable component created by forced liquidation, and any persistence of the initial leg typically makes the reversal leg cheaper conditional on the same liquidation response. The argument therefore needs only that liquidation flow moves the transaction price at first order when it arrives.

Second, longer-horizon round trips are not needed for the impossibility result, but they can strengthen it. If a profitable two-period round trip exists at a reachable state, it can be embedded into any longer horizon by setting $a_{t+2} = \cdots = a_{t+T} = 0$, which preserves admissibility and preserves the same expected profit. Longer horizons can also generate additional profitable deviations when successive updates propagate the initial perturbation forward (for example, because prices affected at $t+1$ enter subsequent inputs and the relevant one-sided slopes remain active). A full finite-horizon characterization of such multi-update trigger-and-reverse strategies, and the corresponding stronger stress tests for a fixed disclosed rulebook, is developed in Nicolai and Risteska (2026). For the purposes of this paper, a single profitable two-period deviation is sufficient to violate manipulation-proofness at the evaluated state, so the slope cap used in Section 6 and Appendix D should be read as a tractable sufficient restriction based on the minimal deviation. Enforcing the full finite-horizon requirement in Definition 3 against multi-update strategies can only tighten the admissible local pass-through bound, and thus weakly expands the pooling regions in the optimal schedule.

A.5 Linear execution and funding costs

Let $\tau \geq 0$ denote proportional execution costs, so each trade a_t incurs cost $\tau|a_t|$. Let $\rho \geq 0$ denote a linear funding wedge, so carrying inventory y_t over $(t, t+1]$ incurs cost $\rho M_t |y_t|$. For a trade sequence $a_{0:T}$, define

$$\Pi_T^{\tau,\rho,\kappa}(a_{0:T}) = -\sum_{t=0}^{T} a_t P_t - \tau \sum_{t=0}^{T} |a_t| - \rho \sum_{t=0}^{T-1} M_t |y_t| - \kappa \sum_{t=0}^{T-1} |y_t|^2.$$

When $\tau = \rho = 0$, this reduces to Π_T^{κ} in (11). The optional hard funding constraint (12) continues to impose an upper bound on admissible trigger sizes.

Lemma 15. *Fix a reachable state S_t satisfying Assumptions 5, 1, 3, and 4. Suppose the maximum admissible trigger size is $q_{\max} > 0$, i.e., every $q \in (0, q_{\max}]$ satisfies the locality restrictions (and, if imposed, the hard funding constraint). Let*

$$\underline{G} = \alpha\, c_X\, c_g\, c_\Gamma,$$

where c_Γ is the local sensitivity constant in Assumption 5. For the two-period round trip $a_t = -q$, $a_{t+1} = q$,

$$\mathbb{E}[\Pi_1^{\tau,\rho,\kappa} \mid S_t] \;\geq\; Aq^2 - Bq, \qquad A = \alpha(\underline{G} - 2) - \kappa, \quad B = 2\tau + \rho M_t. \tag{58}$$

If $A > 0$ and $q_{\min} = B/A$, then $\mathbb{E}[\Pi_1^{\tau,\rho,\kappa} \mid S_t] > 0$ for every admissible $q \in (q_{\min}, q_{\max}]$. Thus linear wedges *add the feasibility condition* $q_{\max} > q_{\min}$.

Proof. As in Lemma 5, the induced forced liquidation satisfies $L \geq c_X c_g c_\Gamma \alpha q$. Under $\mathbb{E}[V_{t+1} \mid S_t] = V_t$ and $\mathbb{E}[\Delta X^0_{t+1} \mid S_t] = 0$,

$$\mathbb{E}[\Pi_1^{\tau,\rho,\kappa} \mid S_t] = \alpha q(L - 2q) - (2\tau + \rho M_t)q - \kappa q^2.$$

Substituting the lower bound for L yields (58). If $A > 0$, then $Aq^2 - Bq > 0$ is equivalent to $q > q_{\min} = B/A$. □

B Strategic constrained traders and endogenous participation

Sections 2.4–2.5 treat the constrained sector's target exposure $\bar{x}_t$ as fixed at the start of the pricing window. This makes the forced-liquidation channel transparent, but it raises a natural question: if constrained traders understand that within-window trading can move Γ_t and hence $M_{t+1} = g(\Gamma_t)$, will they scale back ex ante and eliminate amplification? The answer is no. Endogenous target choice can reduce the local liquidation slope by shrinking the mass of traders near the binding cap, but the mechanism remains whenever a positive mass of traders is still locally constrained. The main theorem then applies with the baseline slope c_X replaced by an equilibrium object c_X^{eff}.

B.1 A two-period game

Fix time t and consider the two-period horizon $(t, t+1)$ in the environment of Section 2. The designer commits to the posted rule $M_{t+1} = g(\Gamma_t)$, where Γ_t is computed from sampled transaction prices inside period t (Section 2.3). The sampling protocol is common knowledge and, because it depends on tradable marks, can be influenced by trades within the window.

1. At the start of t, the public state S_t is observed. Each constrained trader $i \in [0,1]$ chooses a target exposure $\bar{x}_t(i)$, and the implemented position is

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

   as in (6). The difference from the main text is that $\bar{x}_t(i)$ is now chosen optimally as a function of S_t and the known rule g.

2. Inside the pricing window, the strategic trader chooses a_t. Prices satisfy (2), so a_t affects sampled marks and hence Γ_t.

3. The new requirement $M_{t+1} = g(\Gamma_t)$ is posted. Constrained traders then adjust at $t+1$ to satisfy the margin constraint under M_{t+1}. When $M_{t+1} > M_t$ and the constraint binds, this adjustment is forced liquidation and generates price-impactful order flow.

4. The strategic trader chooses a_{t+1} and completes the round trip, $a_t + a_{t+1} = 0$.

B.2 Objective

We use a parsimonious objective with three elements: a genuine exposure motive, a convex penalty for large positions, and costs from posted margin, including both funding costs and anticipated losses from manipulation-driven margin changes. Focus on the long side in a neighborhood where forced adjustment is selling.[19] Fix a state S_t. Constrained trader i chooses $\bar{x}_t(i) \geq 0$ to maximize

$$\max_{\bar{x} \geq 0} \left\{ b_t(i)\, x(\bar{x}) - \frac{\chi}{2}\, x(\bar{x})^2 - \rho\, M_t\, x(\bar{x}) - \frac{\eta}{2}\, \sigma_{M,t}^2(S_t)\, x(\bar{x})^2 \right\}, \tag{59}$$

[19] With sign reversals, the same analysis applies to forced buy-ins when short constraints tighten.

where $x(\bar{x}) = \min\{\bar{x}, E(i)/M_t\}$ is the end-of-t position, $b_t(i) \geq 0$ is the marginal value of exposure, $\chi > 0$ is a baseline quadratic penalty, and $\rho \geq 0$ is the per-period funding spread for posted margin (Appendix A.5). The final term is a reduced-form expected loss from exposure to margin jumps. The key equilibrium object is

$$\sigma^2_{M,t}(S_t) = \mathbb{E}\big[(M_{t+1} - M_t)^2 \mid S_t\big]\,, \tag{60}$$

the conditional second moment of the next margin change under equilibrium behavior.

To close the manipulator side explicitly, suppose the strategic trader chooses the trigger size $q \geq 0$ in the two-leg round trip $(a_t, a_{t+1}) = (-q, q)$ to maximize conditional expected funding-adjusted profit, net of execution and funding costs, subject to admissibility. Given S_t and the constrained-sector targets chosen at step 1, the trigger trade moves the sampled mark through impact, which shifts the sampled input and hence the posted margin. A chain-rule approximation is

$$\Delta M_{t+1} \approx g'(\Gamma_t) \cdot \frac{\partial \Gamma_t}{\partial P_t} \cdot \frac{\partial P_t}{\partial q} \cdot q, \qquad \text{with } \frac{\partial P_t}{\partial q} = -\alpha \text{ under (2) for } a_t = -q.$$

If the strategic trader plays a pure best response $q^*(S_t)$, then conditional on S_t the next margin is pinned down by $M_{t+1} = g(\Gamma_t)$ evaluated at $q^*(S_t)$, so $\sigma^2_{M,t}(S_t) = (M_{t+1} - M_t)^2$. If he mixes over q or over the trigger sign, then $\sigma^2_{M,t}(S_t)$ is the corresponding conditional second moment.

States with greater local pass-through in g and greater susceptibility of the sampling window to within-window trading tend to have larger $\sigma^2_{M,t}(S_t)$. This raises the ex ante cost of large constrained positions. If the cap does not bind for trader i, so $x(\bar{x}) = \bar{x}$, then (59) has the interior solution

$$\bar{x}^*_t(i) = \frac{\big(b_t(i) - \rho M_t\big)_+}{\chi + \eta\, \sigma^2_{M,t}(S_t)}. \tag{61}$$

Thus $\bar{x}^*_t(i)$ falls with M_t whenever $b_t(i) > \rho M_t$, and holding M_t fixed it also falls with $\sigma^2_{M,t}(S_t)$. If $b_t(i) \leq \rho M_t$, then $\bar{x}^*_t(i) = 0$. Anticipated manipulation risk therefore reduces participation and shrinks the mass of agents near the binding cap.

B.3 Aggregate demand and an effective liquidation elasticity

Let $\bar{x}^*_t(i)$ denote the target chosen at the start of period t in (61). For any posted requirement M, define realized exposure for trader i by

$$x(i; M, \bar{x}^*_t(i)) = \min\left\{\bar{x}^*_t(i), \frac{E(i)}{M}\right\},$$

and aggregate constrained exposure by

$$X(M, \bar{x}^*_t) = \int_0^1 x(i; M, \bar{x}^*_t(i))\, di.$$

In period t, the strategic trader can move Γ_t and hence M_{t+1}, but the constrained sector's targets are already chosen. The relevant local object is therefore the response of next-period realized constrained exposure to the next posted requirement, holding those prechosen targets fixed. Define the effective liquidation elasticity at state S_t by

$$c_X^{\text{eff}}(S_t) = -\frac{\partial}{\partial M} X(M, \bar{x}_t^*)\Big|_{M=M_{t+1}^{(0)}}. \tag{62}$$

Let

$$m_t^{(0)}(S_t) = \int_0^1 \mathbf{1}\left\{ \bar{x}_t^*(i) \geq \frac{E(i)}{M_{t+1}^{(0)}} \right\} di$$

denote the mass of traders whose cap binds at the benchmark next requirement $M_{t+1}^{(0)}$. If $m_t^{(0)}(S_t) > 0$, then

$$c_X^{\text{eff}}(S_t) = \frac{1}{(M_{t+1}^{(0)})^2} \int_0^1 E(i)\, \mathbf{1}\left\{ \bar{x}_t^*(i) \geq \frac{E(i)}{M_{t+1}^{(0)}} \right\} di > 0. \tag{63}$$

Since (61) implies that $\bar{x}_t^*(i)$ is decreasing in $\sigma_{M,t}^2(S_t)$, higher anticipated manipulation risk lowers the mass of locally binding traders and therefore lowers $c_X^{\text{eff}}(S_t)$.

B.4 Anticipation dampens but does not eliminate the channel

Lemma 16. *Fix a state S_t. If $m_t^{(0)}(S_t) > 0$, then there exist $\delta_X > 0$ and $\underline{c}_X^{\text{eff}}(S_t) > 0$ such that, holding the equilibrium targets fixed, $M \mapsto X(M, \bar{x}_t^*)$ is continuously differentiable on*

$$[M_{t+1}^{(0)},\, M_{t+1}^{(0)} + \delta_X]$$

and

$$\frac{\partial X}{\partial M}(M, \bar{x}_t^*) \leq -\underline{c}_X^{\text{eff}}(S_t) \quad \text{for all } M \in [M_{t+1}^{(0)},\, M_{t+1}^{(0)} + \delta_X].$$

If instead $m_t^{(0)}(S_t) = 0$, then

$$\frac{\partial X}{\partial M}(M_{t+1}^{(0)}, \bar{x}_t^*) = 0.$$

Proof. For long positions,

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

On any interval over which the binding set is unchanged,

$$\frac{\partial X}{\partial M}(M, \bar{x}_t^*) = -\frac{1}{M^2} \int_0^1 E(i)\, \mathbf{1}\left\{ \bar{x}_t^*(i) \geq \frac{E(i)}{M} \right\} di.$$

If $m_t^{(0)}(S_t) > 0$, the binding set is non-empty at $M_{t+1}^{(0)}$, hence on some neighborhood

$$[M_{t+1}^{(0)},\, M_{t+1}^{(0)} + \delta_X],$$

and the derivative is bounded above by $-\underline{c}_X^{\text{eff}}(S_t)$ for some $\underline{c}_X^{\text{eff}}(S_t) > 0$. If $m_t^{(0)}(S_t) = 0$, the indicator set is empty at $M_{t+1}^{(0)}$, so the derivative there is zero. □

Proposition 3 (Anticipation dampens but does not eliminate the manipulation channel)**.** Fix a margin schedule g and a state S_t satisfying the local risk-input sensitivity conditions of the main theorem.

If $m_t^{(0)}(S_t) > 0$, then the two-period deviation in Lemma 5 remains strictly profitable whenever the amplification condition holds with c_X replaced by $c_X^{\text{eff}}(S_t)$. Anticipation can reduce $c_X^{\text{eff}}(S_t)$, but it does not eliminate the round-trip channel while constraints remain locally binding.

If instead $m_t^{(0)}(S_t) = 0$, then $c_X^{\text{eff}}(S_t) = 0$ and the forced-liquidation leg is locally absent.

Proof. Lemma 16 supplies the forced-liquidation step used in Lemma 5, with $c_X^{\text{eff}}(S_t)$ replacing c_X. The profit calculation is otherwise unchanged. If $m_t^{(0)}(S_t) = 0$, there is no local forced-liquidation response. □

B.5 Design implications and the participation trade-off

This extension highlights an additional cost of steep pass-through in g. A larger local slope increases the direct profitability of within-window manipulation, but it also raises the constrained sector's exposure to manipulation-driven margin jumps, increases $\sigma_{M,t}^2(S_t)$, and lowers equilibrium targets in (61). The resulting decline in c_X^{eff} is stabilizing, but it operates through reduced leveraged participation and hence through weaker liquidity provision and less effective use of the cleared market. This reinforces the design logic in Section 6: slope caps and pooling regions reduce manipulation incentives while also limiting ex ante contractions in constrained participation.

C Calibration details for Section 3.6

This appendix documents the empirical inputs used to interpret the loop-gain magnitude objects in Section 3.6.

C.1 Objects and identities

Fix the one-update trigger-and-reverse deviation from Lemma 5: sell $q > 0$ at the sampling time t and buy q at $t+1$. Under linear impact, selling q distorts the sampled mark by $\Delta P_t = \alpha q$. Define

$$\delta_{\text{bps}} = 10^4 \frac{\alpha q}{P_t}, \qquad G = \frac{L}{q}, \qquad m = \frac{M_t}{P_t}.$$

Ignoring proportional costs and inventory penalties ($\kappa = \tau = \rho = 0$), Lemma 5 implies the bps profit and ROI lower bounds

$$\pi_{\text{bps}} = 10^4 \frac{\mathbb{E}[\Pi_1 \mid S_t]}{qP_t} \;\geq\; (G-2)\,\delta_{\text{bps}}, \qquad ROI\% = 100 \frac{\mathbb{E}[\Pi_1 \mid S_t]}{m\,qP_t} \;\geq\; \frac{(G-2)\,\delta_{\text{bps}}}{100\,m}. \tag{64}$$

The liquidation-driven price pressure satisfies $|\Delta P^{\text{liq}}_{t+1}| = \alpha L = G\alpha q$, so in bps the induced liquidation pressure is $G\,\delta_{\text{bps}}$.

C.2 Impact

Let ADV denote average daily traded value (dollars) and let $N_q = qP_t$ denote the trigger notional. Define the participation rate

$$f = \frac{N_q}{ADV} \in (0, 1].$$

Empirical microstructure evidence often supports square-root scaling of costs in participation units. A convenient benchmark is

$$\delta_{\text{bps}}(f) \;\approx\; 200\sqrt{f}, \tag{65}$$

where f is expressed as a fraction of full-day ADV. The constant 200 corresponds to a one-way cost scale of roughly $100\sqrt{f}$ bps combined with the standard triangle heuristic that maps average execution cost into a peak mark distortion at the sampling time. Using full-day ADV is conservative for settlement and margining since the relevant sampling-window volume is typically smaller than full-day volume, which raises effective participation and thus δ_{bps}. For instance, from equation (65) we have:

$$f = 1\% \Rightarrow \delta_{\text{bps}} \approx 20, \quad f = 2\% \Rightarrow \delta_{\text{bps}} \approx 28, \quad f = 5\% \Rightarrow \delta_{\text{bps}} \approx 45, \quad f = 10\% \Rightarrow \delta_{\text{bps}} \approx 63.$$

Inverting (65) gives $f \approx (\delta_{\text{bps}}/200)^2$, so even very small bps distortions correspond to extremely small fractions of full-day ADV when the sampling window is thin. This is consistent with Frazzini et al. (2018) and Kyle and Obizhaeva (2016).

C.3 Elasticities

Let X denote the position of the constraint-sensitive sector. Define the elasticity of this sector's position to the margin level M by

$$\varepsilon_{X,M} = -\frac{d \log X}{d \log M}.$$

For a relative increase $\Delta M/M$, the implied fractional position reduction is

$$\frac{L}{X} \approx \varepsilon_{X,M} \frac{\Delta M}{M}. \tag{66}$$

Empirical work on futures margins reports economically meaningful elasticities. To keep the calibration conservative and avoid selecting extreme estimates, we use the range

$$\varepsilon_{X,M} \in \{0.15,\ 0.225,\ 0.30\}.$$

This range is consistent with evidence that margin increases reduce open interest and speculative demand over horizons relevant for margin adjustment, e.g., the evidence in Hedegaard (2014).

C.4 Forced-liquidation multiple

Define the turnover ratio

$$\tau = \frac{XP_t}{ADV}.$$

Since $f = (qP_t)/ADV$, we have $q/X = f/\tau$. Combining with (66) yields

$$G = \frac{L}{q} \approx \varepsilon_{X,M} \frac{\Delta M}{M} \frac{\tau}{f}. \tag{67}$$

This identity shows that G increases when: (i) margin moves are large ($\Delta M/M$ high), (ii) the constrained sector is large relative to market depth (τ high), and (iii) the trigger is a small share of volume (f low).

C.5 Margin increases

Section 3.6 uses relative margin jumps

$$\frac{\Delta M}{M} \in \{20\%,\ 40\%,\ 60\%\},$$

which correspond to moving from a baseline margin ratio $m_0 = 5\%$ to $m_1 \in \{6\%, 7\%, 8\%$ under the local linearization $m = M/P$. Risk-based initial margin models can generate short-horizon margin calls of this order even with anti-procyclicality tools. Stress-episode evidence and model-based simulations document large and discrete changes in initial margin over short horizons; these facts motivate using 20% to 60% as routine stressed and binding scenarios rather than tail hypotheticals.

C.6 Tables

Table 2 Calibration inputs used in Appendix C.

Object	Symbol	Values
Baseline margin ratio	m_0	5%
Post-jump margin ratio	m_1	6%, 7%, 8%
Relative margin jump	$\Delta M/M$	20%, 40%, 60%
Position-margin elasticity	$\varepsilon_{X,M}$	0.15, 0.225, 0.30
Turnover ratio	τ	0.5, 1, 2
Participation rate	f	0.5%, 1%, 2%, 5%
Mark distortion mapping	$\delta_{\text{bps}}(f)$	$200\sqrt{f}$

Table 3
Forced-liquidation multiple $G = L/q$ implied by (67). Each cell lists G for $\varepsilon_{X,M} = 0.15/0.225/0.30$.

	$\tau = XP_t/ADV$		
f	0.5	1	2
Panel A: $\Delta M/M = 20\%$			
0.5%	3.0/4.5/6.0	6.0/9.0/12.0	12.0/18.0/24.0
1%	1.5/2.25/3.0	3.0/4.5/6.0	6.0/9.0/12.0
2%	0.75/1.125/1.5	1.5/2.25/3.0	3.0/4.5/6.0
5%	0.30/0.45/0.60	0.60/0.90/1.20	1.20/1.80/2.40
Panel B: $\Delta M/M = 40\%$			
0.5%	6.0/9.0/12.0	12.0/18.0/24.0	24.0/36.0/48.0
1%	3.0/4.5/6.0	6.0/9.0/12.0	12.0/18.0/24.0
2%	1.5/2.25/3.0	3.0/4.5/6.0	6.0/9.0/12.0
5%	0.60/0.90/1.20	1.20/1.80/2.40	2.40/3.60/4.80
Panel C: $\Delta M/M = 60\%$			
0.5%	9.0/13.5/18.0	18.0/27.0/36.0	36.0/54.0/72.0
1%	4.5/6.75/9.0	9.0/13.5/18.0	18.0/27.0/36.0
2%	2.25/3.375/4.5	4.5/6.75/9.0	9.0/13.5/18.0
5%	0.90/1.35/1.80	1.80/2.70/3.60	3.60/5.40/7.20

Table 4

ROI lower bound in percent using (64) with $\delta_{\text{bps}}(f)$ from (65) and the conservative denominator $m_1 q P_t$. Each cell lists ROI% for $\varepsilon_{X,M} = 0.15/0.225/0.30$. Negative entries correspond to $G < 2$.

	$\tau = XP_t/ADV$		
f	0.5	1	2
Panel A: $m_0 = 5\% \to m_1 = 6\%$ ($\Delta M/M = 20\%$)			
0.5%	2.8/7.1/11.4	11.4/19.9/28.5	28.5/45.6/62.6
1%	-2.0/1.0/4.0	4.0/10.1/16.1	16.1/28.2/40.3
2%	-7.1/-5.0/-2.8	-2.8/1.4/5.7	5.7/14.2/22.8
5%	-15.3/-14.0/-12.6	-12.6/-9.9/-7.2	-7.2/-1.8/3.6
Panel B: $m_0 = 5\% \to m_1 = 7\%$ ($\Delta M/M = 40\%$)			
0.5%	9.8/17.1/24.4	24.4/39.0/53.7	53.7/83.0/112.3
1%	3.5/8.6/13.8	13.8/24.2/34.5	34.5/55.2/75.9
2%	-2.4/1.2/4.9	4.9/12.2/19.5	19.5/34.2/48.8
5%	-10.8/-8.5/-6.2	-6.2/-1.5/3.1	3.1/12.3/21.6
Panel C: $m_0 = 5\% \to m_1 = 8\%$ ($\Delta M/M = 60\%$)			
0.5%	14.9/24.6/34.2	34.2/53.4/72.6	72.6/111.0/149.5
1%	7.5/14.3/21.1	21.1/34.7/48.3	48.3/75.5/102.7
2%	1.1/5.9/10.7	10.7/20.3/29.9	29.9/49.1/68.3
5%	-7.4/-4.4/-1.4	-1.4/4.7/10.8	10.8/23.0/35.1

D Optimal Margin Design: Supplementary Derivations

This appendix collects the derivations and proof sketches for Section 6. Throughout, r denotes variance, $\sigma = \sqrt{r}$ the associated volatility, and z the state. The rule is $M = g(r)$ in margin-level units and $m = M/P$ in margin-ratio units.

D.1 First-best target and comparative statics

First-order conditions

Under $\ell(x) = x_+$ and $C(x) = x^2/2$, and for a state such that $M > M_-(z)$,

$$\partial_M \mathbb{E}[(D_{t+1} - M)_+ \mid r, z] = -\Pr(D_{t+1} > M \mid r, z)\,,$$

$$\partial_M \frac{1}{2}\Big(c_X^{\text{eff}}(z)\big(M - M_-(z)\big)\Big)^2 = \big(c_X^{\text{eff}}(z)\big)^2\big(M - M_-(z)\big),$$

and

$$\partial_M H(1 - N(M, M_-(z); z)) = -H'(1 - N(M, M_-(z); z))\,\partial_1 N(M, M_-(z); z).$$

Setting the derivative of $\Psi(r, M; z)$ to zero yields (35). The positivity of $p(M; z)$ in (36) follows from $H' > 0$ and $\partial_1 N < 0$.

Quantile and expected-shortfall targets

If $\lambda_L = \lambda_P = 0$, the first-order condition becomes

$$\Pr\Big(D_{t+1} > M^{\text{fb}}(r; z) \mid r, z\Big) = \kappa_M.$$

Using $D_{t+1} = \sqrt{r}\,Y_{t+1}$ gives the quantile rule in (37). If instead the designer minimizes a tail-loss objective subject to a fixed tail probability p, then positive homogeneity of expected shortfall implies

$$M^{\text{fb}}_{ES}(r; z) = \sqrt{r}\,ES_p(Y_{t+1} \mid z)\,.$$

Hence both VaR-based and ES-based first-best targets inherit the same square-root scaling.

Slope with respect to r

Define

$$F(r, M; z) = \Pr(D_{t+1} > M \mid r, z) - \kappa_M - \lambda_L\big(c_X^{\text{eff}}(z)\big)^2\big(M - M_-(z)\big) - \lambda_P p(M; z).$$

At an interior optimum, $F(r, M^{\text{fb}}(r; z); z) = 0$. The implicit function theorem gives

$$\partial_r M^{\text{fb}}(r; z) = -\frac{\partial_r F(r, M; z)}{\partial_M F(r, M; z)}\bigg|_{M = M^{\text{fb}}(r; z)}.$$

Notice that

$$-\partial_M F(r, M; z) = f_D(M \mid r, z) + \lambda_L \big(c_X^{\text{eff}}(z)\big)^2 + \lambda_P \partial_1 p(M; z) \equiv \Delta(r; z) > 0.$$

Under the square-root scaling,

$$\Pr(D_{t+1} > M \mid r, z) = 1 - F_Y\left(\frac{M}{\sqrt{r}} \middle| z\right),$$

so

$$\partial_r \Pr(D_{t+1} > M \mid r, z) = \frac{M}{2r^{3/2}} f_Y\left(\frac{M}{\sqrt{r}} \middle| z\right) > 0.$$

Therefore the first-best target is increasing in r.

Comparative statics with respect to market functioning and participation

Differentiating with respect to c_X^{eff} gives

$$\partial_{c_X^{\text{eff}}} M^{\text{fb}}(r; z) = -\frac{2\lambda_L c_X^{\text{eff}}(z)\big(M^{\text{fb}}(r; z) - M_-(z)\big)}{\Delta(r; z)},$$

which is negative whenever margins increase. If a parameter η enters the participation block only through $p(M; z, \eta)$, then

$$\partial_\eta M^{\text{fb}}(r; z) = -\frac{\lambda_P \partial_\eta p(M^{\text{fb}}(r; z); z, \eta)}{\Delta(r; z)}.$$

Thus any shift that makes participation more margin-sensitive lowers the first-best target.

D.2 The constrained design problem

Fix z and define

$$\mathcal{G}_{\text{IC}}(z) = \big\{0 \le g'(r) \le \bar{s}(r; z) \text{ for } r \in \mathcal{I}\big\}. \tag{68}$$

KKT conditions

Introduce the multipliers $\mu(r; z)$ and $\nu(r; z)$ for the upper and lower derivative constraints. The Lagrangian is

$$\mathcal{L}(g, \mu, \nu; z) = \int_{\mathcal{I}} \big[f(r \mid z)\, \Psi\big(r, g(r); z\big) + \mu(r; z)\big(g'(r) - \bar{s}(r; z)\big) - \nu(r; z) g'(r)\big]\, dr. \tag{69}$$

If we vary g by a small perturbation h, integration by parts gives

$$\delta\mathcal{L} = \int_{\mathcal{I}} \big[f(r \mid z)\, \partial_M \Psi\big(r, g(r); z\big) - \mu_r(r; z) + \nu_r(r; z)\big]\, h(r)\, dr.$$

Hence any interior optimum satisfies

$$f(r \mid z)\, \partial_M \Psi\big(r, M^*(r; z); z\big) - \mu_r(r; z) + \nu_r(r; z) = 0, \tag{70}$$

together with complementary slackness

$$\mu(r;z)\big(M^{*\prime}(r;z)-\bar{s}(r;z)\big)=0, \tag{71}$$

$$\nu(r;z)M^{*\prime}(r;z)=0. \tag{72}$$

Tracking, saturation, and pooling.

If $0<M^{*\prime}(r;z)<\bar{s}(r;z)$ on an interval, complementary slackness implies $\mu=\nu=0$ there. Stationarity then reduces to

$$\partial_M\Psi\big(r,M^*(r;z);z\big)=0,$$

which identifies $M^*(r;z)$ with the pointwise first-best target $M^{\mathrm{fb}}(r;z)$. This is the tracking region.

If $M^{*\prime}(r;z)=\bar{s}(r;z)$, the upper derivative constraint binds and the rule rises at the maximal feasible rate. Under the realized-variance envelope,

$$M^{*\prime}(r;z)=\frac{1}{\alpha(z)c_X^{\mathrm{eff}}(z)\sqrt{r}},$$

so integration yields the closed form

$$M^*(r;z)=a(z)+\frac{2}{\alpha(z)c_X^{\mathrm{eff}}(z)}\sqrt{r}$$

on any slope-saturated interval.

If $M^{*\prime}(r;z)=0$ on an interval $[a,b]$, the rule is locally flat. To pin down the pooled level k, perturb the flat segment by a common shift ε. Feasibility is preserved for small $|\varepsilon|$, so optimality requires

$$0=\frac{d}{d\varepsilon}\bigg|_{\varepsilon=0}\int_a^b f(r\mid z)\,\Psi(r,k+\varepsilon;z)\,dr=\int_a^b f(r\mid z)\,\partial_M\Psi(r,k;z)\,dr,$$

which is exactly (43).

Quadratic approximation around the first-best target.

A useful local approximation is

$$\Psi(r,M;z)\approx\Psi\big(r,M^{\mathrm{fb}}(r;z);z\big)+\frac{1}{2}h(r;z)\big(M-M^{\mathrm{fb}}(r;z)\big)^2,$$

where $h(r;z)=\partial_{MM}\Psi\big(r,M^{\mathrm{fb}}(r;z);z\big)>0$. On a pooling interval $[a,b]$, the condition (43) becomes

$$k=\frac{\int_a^b h(r;z)f(r\mid z)M^{\mathrm{fb}}(r;z)\,dr}{\int_a^b h(r;z)f(r\mid z)\,dr}.$$

Thus the pooled level is the weighted average of the first-best target over the interval. The weights depend on both the density of states and the local curvature of the objective.

D.3 Empirical bounds

This subsection converts the exact local cap into empirical objects that can be calibrated from market depth, liquidation elasticities, and sampled-price distortions. Start from the exact local cap in levels under $\kappa = 0$:

$$\frac{dM}{dr_P} \le \frac{1}{\alpha c_X^{\text{eff}} |\Delta p/P| P}, \tag{73}$$

where r_P is the price-level variance input and $\Delta p/P$ is the last sampled return in decimal units. Let

$$r = \frac{r_P}{P^2}, \qquad \sigma = \sqrt{r}, \qquad m(r) = \frac{g(P^2 r)}{P}.$$

Holding the current price fixed over the local perturbation gives

$$\frac{dm}{dr} = P\frac{dM}{dr_P}, \qquad \frac{dm}{d\sigma} = \frac{dM}{d\sigma_P}, \qquad \sigma_P = P\sigma.$$

Using $|\Delta p/P| \le \sigma$, the exact cap implies

$$\frac{dm}{d\sigma} \le \frac{2}{\alpha c_X^{\text{eff}}}, \qquad \frac{dm}{dr} \le \frac{1}{\alpha c_X^{\text{eff}} \sigma}, \qquad \frac{dm}{d\sigma^{ann}} \le \frac{2}{\sqrt{252}\, \alpha c_X^{\text{eff}}}. \tag{74}$$

Next we need to rewrite αc_X^{eff} in terms of standard measurable quantities. Let

$$\epsilon_{X,M} = -\frac{d \log X}{d \log M}, \qquad c_X^{\text{eff}} = \epsilon_{X,M} \frac{X}{M}, \qquad \tau = \frac{XP}{ADV}, \qquad m = \frac{M}{P}.$$

Let $f = qP/ADV$ be the trigger trade as a fraction of average daily traded value, and let

$$\delta(f) = \frac{\delta_{bps}(f)}{10^4}$$

denote the sampled-price distortion in decimal units. Since $\delta(f) = \alpha q/P$,

$$\alpha = \frac{P^2}{ADV} \frac{\delta(f)}{f}. \tag{75}$$

Using $X = \tau ADV/P$ and $M = mP$,

$$c_X^{\text{eff}} = \epsilon_{X,M} \frac{\tau ADV}{mP^2}. \tag{76}$$

Multiplying (75) and (76) yields

$$\alpha c_X^{\text{eff}} = \frac{\epsilon_{X,M}\tau}{m} \frac{\delta(f)}{f} = \frac{\epsilon_{X,M}\tau}{m} \frac{\delta_{bps}(f)}{10^4 f}. \tag{77}$$

Substituting (77) into (74) gives

$$\frac{dm}{d\sigma} \le \frac{2mf}{\epsilon_{X,M}\tau\delta(f)} = \frac{2 \times 10^4 mf}{\epsilon_{X,M}\tau\delta_{bps}(f)}, \tag{78}$$

$$\frac{dm}{d\sigma^{ann}} \le \frac{2mf}{\sqrt{252}\,\epsilon_{X,M}\tau\delta(f)} = \frac{2\times 10^4 mf}{\sqrt{252}\,\epsilon_{X,M}\tau\delta_{bps}(f)}. \tag{79}$$

Dividing through by m yields the relative-cap formulas

$$\frac{1}{m}\frac{dm}{d\sigma} \le \frac{2f}{\epsilon_{X,M}\tau\delta(f)}, \qquad \frac{1}{m}\frac{dm}{d\sigma^{ann}} \le \frac{2f}{\sqrt{252}\,\epsilon_{X,M}\tau\delta(f)}. \tag{80}$$

Therefore, approximately

$$\frac{\Delta m}{m} \le \frac{2f}{\sqrt{252}\,\epsilon_{X,M}\tau\delta(f)}\,\Delta\sigma^{ann}. \tag{81}$$

which is the expression used in the main text.

D.4 Benchmark thresholds

Under the benchmark calibration

$$\epsilon_{X,M} = 0.225, \qquad f = 1\%, \qquad \delta_{bps}(1\%) \approx 20,$$

equation (81) becomes

$$\frac{\Delta m}{m} \le \frac{2.80}{\tau}\,\Delta\sigma^{ann}, \tag{82}$$

where $\Delta\sigma^{ann}$ is measured in decimal annualized-volatility units. A ten-point increase in annualized volatility implies

$$\frac{\Delta m}{m} \le \frac{0.28}{\tau}.$$

The annualized-volatility increase required to support a targeted proportional margin increase η_M is

$$\Delta\sigma^{ann}_{min}(\eta_M) = \frac{\tau}{2.80}\,\eta_M. \tag{83}$$

Measured in annualized-volatility points,

$$\Delta\sigma^{ann}_{min}(\eta_M) \approx 7.14\tau \times \eta_M, \tag{84}$$

with η_M in decimal units.

Table 5 reports the resulting benchmark thresholds. The table makes the comparison with observed initial margin rules straightforward. Moderate short-horizon margin increases can fit inside the cap when τ is modest. But once the constrained sector is large relative to market depth, even a highly risk-sensitive rule can only move margin modestly unless the volatility shock is extremely large.

D.5 Portfolio margining

This subsection records the two extra ingredients that we need to keep track of when discussing portfolio margining. The first modification concerns the cost of a tighter reset. If a one-unit increase in margin induces liquidation vector B_t, then the associated liquidation-driven price move is A_tB_t.

Table 5 Implementability thresholds

	$\tau = 0.5$	$\tau = 1$	$\tau = 2$
Max $\Delta m/m$ for +10 annualized-vol points	56%	28%	14%
Max Δm for +10 annualized-vol points when $m = 5\%$	2.8 pp	1.4 pp	0.7 pp
Annualized-vol points needed for a 20% margin jump	3.6	7.1	14.3
Annualized-vol points needed for a 40% margin jump	7.1	14.3	28.6
Annualized-vol points needed for a 60% margin jump	10.7	21.4	42.9
Annualized-vol points needed for an 80% margin jump	14.3	28.6	57.1

Under a quadratic benchmark, the per-state objective becomes

$$\Psi^P(r, M; z) = \mathbb{E}[\ell(D_{t+1} - M) \mid r, z] + \kappa_M M \tag{85}$$

$$+ \frac{\lambda_L}{2}\|A_t B_t\|_2^2 \big(M - M_-(z)\big)_+^2 + \lambda_P H\big(1 - N^P(M, M_-(z); z)\big). \tag{86}$$

Relative to the single-asset case, the term $(c_X^{\text{eff}})^2$ is replaced by $\|A_t B_t\|_2^2$. Hence larger liquidation price pressure makes the first-best target flatter.

The second modification concerns implementability. Let d_t be the direction that raises the portfolio-risk input, and let $c_{r,t}$ be the directional sensitivity of that input. From Section 5, a sufficient local no-manipulation condition is

$$\big(g'(r_t) c_{r,t}\big)\psi_t \leq 2, \qquad \psi_t = \max\{0, -d_t^\top A_t B_t\}. \tag{87}$$

Equivalently,

$$g'(r_t) \leq \frac{2}{c_{r,t}\psi_t}. \tag{88}$$

Thus larger trigger-liquidation alignment tightens the admissible slope pointwise.

The portfolio case therefore differs from the single-asset benchmark in exactly two ways: the first-best target becomes flatter when liquidation generates more cross-asset price pressure, as measured by $\|A_t B_t\|_2$; and, the implementability cap becomes tighter when the trigger portfolio is more strongly aligned with the liquidation portfolio, as measured by ψ_t.

D.6 A two-asset illustration

A simple two-asset example shows that portfolio margining need not mechanically make the rule flatter. Depending on cross-impact and on the alignment between the trigger portfolio and the liquidation portfolio, the admissible slope can be either looser or tighter than in the single-asset benchmark. Take two assets and suppose the cross-impact matrix is

$$A = \alpha \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

Let the liquidation response to a one-unit increase in posted margin be split evenly across the two assets:

$$B = \frac{c}{2}\begin{pmatrix}1\\1\end{pmatrix},$$

where $c > 0$ is chosen so that the total liquidation sensitivity matches the single-asset benchmark. That means the constrained sector keeps portfolio weights and proportionally sell equally the two assets. Finally, suppose the local risk-increasing direction is the equally weighted direction

$$d = -\frac{1}{\sqrt{2}}\begin{pmatrix}1\\1\end{pmatrix},$$

so that selling both assets raises the scalar portfolio risk input. Let $c_r > 0$ denote the sensitivity of risk along d, and assume that c_r is the same as in the single-asset comparison benchmark. Under these assumptions,

$$AB = \alpha\frac{c}{2}\begin{pmatrix}1+\rho\\1+\rho\end{pmatrix},$$

so the liquidation price-pressure scale is

$$\|AB\|_2 = \alpha c\frac{1+\rho}{\sqrt{2}}. \tag{89}$$

The trigger-liquidation alignment term is

$$\psi = -d^{\top}AB = \alpha c\frac{1+\rho}{\sqrt{2}}. \tag{90}$$

For comparison, define the corresponding single-asset benchmark price-pressure scale as

$$\xi^{SA} = \alpha c.$$

Then the two-asset example implies

$$\frac{\|AB\|_2}{\xi^{SA}} = \frac{\psi}{\xi^{SA}} = \frac{1+\rho}{\sqrt{2}}. \tag{91}$$

On the cost side, the first-best target is flatter when the liquidation price-pressure scale is larger. Relative to the single-asset benchmark, the two-asset first-best is therefore flatter if

$$\frac{1+\rho}{\sqrt{2}} > 1 \qquad \Longleftrightarrow \qquad \rho > \sqrt{2}-1 \approx 0.414,$$

and steeper if $\rho < \sqrt{2}-1$. Second, on the implementability side, the portfolio slope cap is

$$g'(r) \leq \frac{2}{c_r\psi}.$$

Hence, relative to the single-asset cap

$$g'_{SA}(r) \leq \frac{2}{c_r \xi^{SA}},$$

the ratio of admissible slopes is

$$\frac{\bar{g}'_P(r)}{\bar{g}'_{SA}(r)} = \frac{\xi^{SA}}{\psi} = \frac{\sqrt{2}}{1+\rho}. \tag{92}$$

Thus portfolio margining makes the rule more responsive when $\rho < \sqrt{2} - 1$, and flatter when $\rho > \sqrt{2} - 1$.

A calibration

It is useful to translate (92) into the benchmark numbers we have used in the main calibration. Under the single-asset calibration with $\tau = 1$, a ten-point increase in annualized volatility permits at most a 28% increase in the margin ratio. In the present two-asset example, the corresponding bound becomes

$$\max \frac{\Delta m}{m} = 28\% \times \frac{\sqrt{2}}{1+\rho} \qquad \text{for a +10 annualized-volatility-point move.} \tag{93}$$

Three cases illustrate the comparison:

$$\rho = 0 \quad \Longrightarrow \quad \max \frac{\Delta m}{m} \approx 39.6\%,$$

$$\rho = 0.5 \quad \Longrightarrow \quad \max \frac{\Delta m}{m} \approx 26.4\%,$$

$$\rho = 1 \quad \Longrightarrow \quad \max \frac{\Delta m}{m} \approx 19.8\%.$$

Starting from a 5% margin ratio, these correspond to maximal manipulation-proof increases of about 1.98, 1.32, and 0.99 percentage points of notional, respectively. If liquidation is spread across weakly linked assets, portfolio margining can make the public rule more responsive than in the single-asset benchmark. But once liquidation loads on positively correlated assets and the trigger direction is aligned with the resulting liquidation pressure, both the first-best target and the implementability cap become flatter. In that case, portfolio netting tightens the design problem.

E How CCP margin engines map into the model

Real-world initial margin systems differ in engineering detail, but their architecture is stable and is described in rulebooks, methodology notes, and parameter files. In most implementations, a core engine produces a scalar risk number from current marks and a prescribed set of risk scenarios or recent price history. We denote this core input by Γ_t. An overlay layer then applies anti-procyclicality tools, floors, add-ons, rounding, and governance adjustments, producing the posted requirement

$$M_{t+1} = g(\Gamma_t).$$

Our model is not a claim that any one CCP uses realized variance. It is a claim about the composite mapping from sampled, transaction-based marks into posted margin, and about what happens when that mapping is locally sensitive in states where initial margin constraints bind.

Two clarifications matter once we allow portfolio margining, as in Section 5. First, the marks are naturally a vector, even for a single clearing service: different contracts, maturities, and options share a portfolio margin system. The core statistic Γ_t is then a portfolio object computed from a vector of marks and from portfolio positions, while M_{t+1} remains a single scalar requirement. Second, a binding call is a funding event. Meeting it can force liquidation in a subset of contracts that need not coincide with the subset of contracts that most efficiently moves the statistic. This is why the multi-asset extension introduces a liquidation direction that summarizes, how a scalar margin change turns into a vector of forced trades.

E.1 SPAN-style margining: CME SPAN and other venues

SPAN is a scenario-based portfolio margin system. At time t, let Θ_t collect the public SPAN inputs, including scenario definitions, scan ranges, spread parameters, delivery parameters, and short-option-minimum parameters (CME Group, 2019b,a, 2026; London Metal Exchange, 2015, 2026). Given current marks P_t, SPAN computes portfolio losses under a finite set of scenarios S and takes the worst one:

$$\Gamma_t^{\text{scan}} = \max_{s \in S} l_s(P_t; \Theta_t).$$

Posted margin is then a further transformation of this scan-risk number:

$$M_{t+1} = g(\Gamma_t^{\text{scan}}; \Theta_t),$$

where g collects the overlay layer, including spread charges and credits, delivery add-ons, short-option-minimum floors, option-value adjustments, rounding, and related conventions (CME Group, 2019b,a; London Metal Exchange, 2015).

This fits our framework directly. First, SPAN is price-based: current marks enter the scenario losses $l_s(P_t; \Theta_t)$. Second, it is portfolio-based: Γ_t^{scan} is computed from a vector of marks and positions, while M_{t+1} is a scalar requirement. Third, it is non-smooth for two reasons. The scan-risk term is a maximum over finitely many scenario losses, so it has kinks when the worst scenario changes. The map g adds further non-smoothness through floors, thresholds, credits, and rounding.

This is exactly the environment of Section 5. A perturbation of marks at time t can move the portfolio risk statistic Γ_t^{scan} and therefore next period's requirement M_{t+1}. If the requirement binds, the induced deleveraging is a vector of forced trades, and those trades need not fall on the contracts that were most effective at moving Γ_t^{scan}. That is the cross-contract wedge emphasized in our multi-asset analysis.

E.2 OCC STANS

OCC STANS is a portfolio-based simulation margin system. Let q_t denote the cleared portfolio at time t, let Z_t collect the marks and risk-state inputs used to seed the simulation, and let $L^{(n)}(q_t; Z_t)$ denote simulated portfolio losses. A natural reduced-form representation is

$$\Gamma_t^{\text{STANS}} = \text{ES}_{0.99}\big(\{L^{(n)}(q_t; Z_t)\}_{n=1}^{N}\big), \qquad M_{t+1} = g(\Gamma_t^{\text{STANS}}),$$

where g collects stress components and other overlays (The Options Clearing Corporation, 2026; U.S. Securities and Exchange Commission, 2024).

This fits Section 5 directly. The risk statistic is computed from a vector of marks and positions, while the output is a scalar requirement. If the requirement binds, the induced deleveraging is a vector of forced trades and need not fall on the contracts that are most effective at moving Γ_t^{STANS}.

For our mechanism, the key point is that Γ_t^{STANS} depends on the marks used to value the portfolio and seed the simulation. When those inputs are built from sampled transaction-based marks at predictable times, a feasible perturbation of a sampled mark can change the simulated loss distribution and hence the posted requirement. The composite mapping can also be non-smooth because OCC adds stress components, max-type conservative choices, and other overlays (The Options Clearing Corporation, 2026). This is the same portfolio-level, potentially non-smooth price-based environment studied in Section 5.

E.3 LCH SwapClear PAIRS

LCH describes PAIRS as a filtered-historical-simulation margin model with volatility scaling and an expected-shortfall-type tail risk measure (LCH, 2026; LCH.Clearnet Limited, 2015). In our notation, the useful reduced-form representation is

$$\Gamma_t^{\text{PAIRS}} = (\text{tail functional of portfolio losses under scaled historical moves}), \qquad M_{t+1} = g(\Gamma_t^{\text{PAIRS}}),$$

where g collects the overlay layer, including add-ons, floors, and intraday margining features documented in the procedures (LCH Limited, 2025; Reserve Bank of Australia, 2019; LCH Ltd, 2022).

This fits Section 5 directly. PAIRS is portfolio-based: it maps a vector of marks and risk-factor inputs into a scalar margin requirement. If the requirement binds, the induced deleveraging is a portfolio response and need not occur in the instruments that are most effective at moving the risk statistic.

For our mechanism, the key point is narrower than any claim of practical manipulability. PAIRS depends on current marks both through portfolio valuation and through the volatility-scaling step.

If an economically important input mark is locally sensitive to tradable prices at a predictable sampling time, and if the scaling step loads positively on recent mark changes, then $\Gamma_t^{\mathrm{PAIRS}}$ inherits local sensitivity to the last sampled mark. That is the only property we need for the general logic.

The remaining ingredients are the same as elsewhere in the paper. A locally mark-sensitive risk input feeds into a scalar posted requirement through g. If the resulting requirement binds, meeting a higher call induces forced portfolio adjustment, with liquidation potentially falling in a different set of instruments from those that moved the risk measure. PAIRS overlays such as floors, add-ons, and intraday triggers can reduce smooth local pass-through in some regions, but they also shift adjustment toward non-smooth events. That is the same continuity tradeoff studied in the main theorem.

E.4 European anti-procyclicality measures as transformations inside g

Under EMIR, CCPs must address margin procyclicality, for example through a buffer, stressed weights, or a long-lookback floor (European Union, 2013; European Securities and Markets Authority, 2018, 2022, 2023). In our notation, these measures are transformations inside g: they take a core risk number Γ_t and map it into posted margin M_{t+1}. This is the only feature that matters for our mechanism. Floors and buffers flatten pass-through when they bind and create kinks when they cease to bind. Threshold-based release rules can add jump-like behavior. Stressed weights damp sensitivity to calm-period data, but can introduce regime changes when the stressed component or its weight changes.