One-Way Coupling of Prediction Markets to Automated Market Makers

An Impossibility Theorem and its Bounded-Impact Representatives

Raeez Lorgat

April 2025

Abstract

We study mechanisms that couple a strictly-proper-scoring-rule prediction market to an automated market maker (AMM) so that the prediction signal influences the AMM's marginal price. The central result is the mixed-partial impossibility theorem (Theorem 3.2): any economically non-trivial *bidirectional* coupling, in which spot-market state feeds back into the prediction market's payoff non-separably from the reported belief, destroys the incentive compatibility of the scoring rule. The admissible class is therefore *one-way*: information flows from the prediction market to the AMM through the mechanism, while the reverse direction is closed at the mechanism level. We formalise this as the non-invertibility of a specific oracle channel from the prediction market's payoff schedule. One-way coupling is a mechanism-design necessity rather than a design choice.

A bounded-C^1-clipped one-way coupling (Definition 5.1) inhabits the admissible class with single-block price impact bounded by $|\Delta P/P| \leq e^{\kappa\delta_{\max}} - 1$ (Lemma 5.4); the power-law and Huber-bounded logit forms are instances (Corollary 5.5). The unclipped power-law E^{κ} admits unbounded impact (Theorem 5.7).

Against a risk-neutral adversary with venue access on both sides, holding a pre-positioned spot notional Q and paying both the LMSR round-trip cost (Hanson [1]) and AMM round-trip slippage, the first-order equilibrium prediction bias is $\alpha Q^{\star}/(2b)$ (Proposition 6.1), giving a design-time depth requirement $b > \alpha Q^{\star}/(2\varepsilon)$. Under Gaussian signal noise the coupling's information value to traders is first-order in $\kappa\sigma_E$ while the LVR reduction at the constant-product maximiser $g = 1$ is second-order (Theorem 9.4).

The joint prediction-AMM clearing problem reduces to a log-barrier program adapted from the Peters-So-Ye convex parimutuel formulation [7]; polynomial-time clearing holds in the power-law regime (Theorem 10.2, Theorem 12.4), and the logit-regime extension holds under an explicit contraction condition in logit space (Theorem 10.5). On bounded continuous outcome spaces, the same convex-potential construction yields a unique equilibrium measure under the functional $\Phi[\mu] = \int_{\Omega} \phi(d\mu/d\omega)\, d\omega$ on absolutely continuous measures; Arrow-Debreu state prices satisfy the first-order conditions of the clearing program directly, piecewise-constant discretisations recover the discrete CPCAM in the atomic-measure limit, and the Huber-bounded logit fixed point survives grid refinement under the same contraction inequality expressed through the continuous price statistic (Theorem 10.9, Proposition 10.10, Theorem 10.11).

Contributions relative to prior art. The principal contribution is the mixed-partial impossibility theorem and its representation corollary (Theorem 3.2, Corollary 3.5): admissible couplings are one-way as a mechanism-design necessity. The log-barrier clearing construction instantiates the Peters-So-Ye mechanism [7] on the parimutuel framework of Lange [5] and Lange-Economides [6]; the continuous fixed-generator form uses the classical Kolmogorov-Nagumo mean structure [8, 9]. Neither ingredient is claimed as original to this paper. The bounded-C^1-clipped representation

(Lemma 5.4), the logit-space fixed-point theorem (Theorem 10.5), and the continuous-state equilibrium / atomic-limit / grid-refinement statements (Theorem 10.9, Proposition 10.10, Theorem 10.11) are subsidiary reference results. The manipulation analysis (Proposition 6.1) extends the Angeris-Chitra oracle-manipulation framework [26] to the coupled setting by adding the AMM round-trip cost to the adversary's objective and identifying the LMSR depth b as the defender's parameter; the information-value ratio (Theorem 9.4) extends the Milionis et al. LVR framework [16] to coupled prediction-AMM venues.

Contents

1 Introduction **4**
- 1.1 The question . . . 4
- 1.2 Relationship to prior work . . . 4
- 1.3 Outline . . . 6

2 Event-Atomic Settlement: an Execution-Environment Assumption **6**
- 2.1 The commitment boundary . . . 6
- 2.2 Proposer constraints and side-channel hardening . . . 6
- 2.3 Necessity of the proposer constraints and within-class priority . . . 8

3 The Impossibility of Non-Trivial Bidirectional Coupling **8**
- 3.1 Setup . . . 8
- 3.2 The impossibility theorem . . . 9
- 3.3 The necessity of one-way coupling . . . 10

4 Admissible One-Way Couplings: Representation Theorems **11**
- 4.1 The signal-extraction problem . . . 11
- 4.2 The affine logit-value link and logit coupling . . . 11
- 4.3 Binary-payoff coupling via EKF . . . 12

5 Bounded Parametrizations of One-Way Couplings **14**
- 5.1 The bounded C^1-clipped class . . . 14
- 5.2 The price-impact bound . . . 15
- 5.3 The necessity of clipping: unbounded couplings break the bound . . . 15
- 5.4 Composition across blocks . . . 15

6 Cross-Venue Manipulation **16**
- 6.1 Adversary model and equilibrium bias . . . 16
- 6.2 Persistent-adversary upper bound under rate-limited calibrator . . . 17
- 6.3 Oracle-manipulation attack upper bound . . . 18
- 6.4 Persistent multi-block manipulation: matching upper bound . . . 19
- 6.5 Governance joint-parameter manipulation . . . 20
- 6.6 Within-class proposer priority: matching lower bound . . . 21

7 The Trading Function **22**
- 7.1 The power-law invariant . . . 22
- 7.2 Marginal price . . . 23

8 LP Loss Analysis **23**
- 8.1 Impermanent loss for the generalised invariant 23
- 8.2 LVR reduction 24
- 8.3 Coupling loss 24
- 8.4 Path dependence 24
- 8.5 Break-even condition 25

9 Information Value of Coupling **25**
- 9.1 Information value: definition and first-order expression 25
- 9.2 Ratio of first-order and second-order effects 26

10 Joint Prediction-AMM Clearing (Reference Construction) **26**
- 10.1 Log-barrier formulation 27
- 10.2 The non-convex joint problem 27
- 10.3 Logit-regime clearing convergence 28
- 10.4 Continuous Product Constant-Argument Mean on bounded outcome spaces 30
- 10.5 Novelty of the clearing construction 35

11 Design-Space Comparison **36**

12 Manifold-Projection Convergence Under Regularity **36**
- 12.1 Strong concavity of the log-barrier objective 36
- 12.2 Lipschitz gradient of the invariant-constraint residual 37
- 12.3 Geometric convergence of the projected-Newton inner solver 37

13 Open Problems **38**

A Adversarial-completeness inventory **40**

1 Introduction

1.1 The question

A prediction market prices the probability of a discrete outcome under a strictly proper scoring rule (Hanson [1, 2]); an AMM (Angeris et al. [3]) prices a continuous asset whose value depends on that outcome. When the two venues trade the same underlying event, the design question is how to compose them. A *coupling* is a mapping by which the prediction signal enters the AMM's pricing function. This paper characterises the admissible class.

The natural first instinct is a *bidirectional* coupling under which prediction signals influence spot prices and spot prices in turn influence the prediction market's payoffs, tying both venues into a single equilibrium. We show that the bidirectional instinct fails as a matter of mechanism design. Whenever the prediction payoff depends on the spot price non-separably from the reported belief (a non-zero mixed partial at the truthful report), the scoring rule's incentive compatibility breaks. The admissible class is therefore *one-way*: the prediction signal may enter the AMM, but the AMM's state does not enter the scoring rule's payoff through any channel that depends on the report. We formalise one-way coupling as the non-invertibility of the oracle channel that would otherwise carry information from the AMM back into the payoff schedule. This is a mechanism-design necessity, not a design parameter.

We then characterise the admissible class: the minimum-variance coupling under an affine logit-value link (Dalen [17]), bounded-impact parametrisations against adversarial signal realisations, and matching bounds on cross-venue manipulation under the full adversary objective including AMM round-trip cost. The paper closes with discrete and continuous clearing constructions adapted from Peters-So-Ye [7].

1.2 Relationship to prior work

The construction sits at the intersection of five literatures. We position this paper against each separately.

Proper scoring rules and prediction markets. Hanson's logarithmic market scoring rule (LMSR) [1, 2] is the canonical automated mechanism for belief elicitation. Strict properness of the LMSR depends on the payoff being a function only of the reported belief and the realised outcome; no cross-venue feedback is envisioned. Chen and Pennock [23] characterise the bounded-loss family of market makers. Abernethy, Chen, and Wortman Vaughan [24] give the convex-analytic framework for cost-function market makers. Othman-Pennock-Reeves-Sandholm [22] treat liquidity sensitivity. The impossibility theorem of Section 3 uses only the C^2 strict-concavity of the expected payoff; it applies to every rule in this literature. Dalen [17] develops the Black-Scholes analog for prediction markets, identifying logit space as the natural operating domain.

Parimutuel clearing and its impossibilities. Lange's parimutuel-derivatives patent [5] and the Lange-Economides formalisation [6] introduced nonlinear optimisation as a mechanism for merging segregated parimutuel pools; the latter noted structural obstructions when auxiliary payoffs are added to the parimutuel mechanism. Peters, So, and Ye [7] convexified the Lange formulation via a log barrier, producing the Convex Parimutuel Call Auction Mechanism (CPCAM) which we use in Section 10. Baron and Lange [21] provide an eigenvalue representation of state prices. Our impossibility theorem (Theorem 3.2) is distinguished from the Lange-Economides parimutuel impossibility in Remark 3.3: we target the LMSR-and-AMM composition and require only C^2 strict-concavity of the scoring rule; the parimutuel impossibility is specific to the call-auction setting with a fixed state-price vector.

Continuous state prices and generalized means. On continuous outcome spaces, fixed-generator mean forms go back to Kolmogorov [8] and Nagumo [9]. Breeden and Litzenberger [10] recover a state-price density ex post from an option-price surface by differentiating in strike. Section 10 uses the same bounded-outcome-space setting differently: the state-price density is a primal clearing object of the continuous CPCAM itself, and the Breeden-Litzenberger identity reappears as a corollary of the cleared density rather than as a separate inverse problem.

Automated market makers and CFMMs. Angeris et al. [3] characterise the optimal arbitrage problem for constant-function market makers. Angeris and Chitra [26] analyse the price-oracle properties of CFMMs and give manipulation-cost bounds. Angeris, Evans, Chitra, and Boyd [27] prove optimality of hub-routing. Milionis, Moallemi, Roughgarden, and Zhang [16] formalise loss-versus-rebalancing (LVR), the framework we use for the LP loss analysis of Section 8. Adams, Zinsmeister, Keefer, Salem, and Robinson [28] introduce concentrated liquidity; Park [30] identifies structural adverse-selection costs borne by LPs in constant-product AMMs. Egorov [4] introduces dynamic pegs, a precursor to event-dependent pricing. Heimbach-Wattenhofer [25] quantify LVR empirically. The AMM side of our coupling is a standard CFMM; the novelty is the composition with a scoring-rule prediction market.

Market microstructure and oracle manipulation. Kyle [11] and Glosten-Milgrom [13] establish the information-theoretic foundations of adverse selection. Budish, Cramton, and Shim [14] propose frequent batch auctions to address speed advantages in continuous markets. Daian et al. [15] formalise miner/proposer extractable value. Roughgarden [29] analyses transaction-fee-mechanism design. Foster and Viswanathan [12] analyse strategic trading with forecasted forecasts. The event-atomic settlement theorem of Section 2 addresses information advantages complementary to Budish-Cramton-Shim's speed advantages; the cross-venue manipulation analysis of Section 6 sits in the Angeris-Chitra [26] oracle-manipulation tradition, sharpened by a matching lower bound.

Statistical estimator fusion. The minimum-variance combination of independent Gaussian signals is a textbook result. Our contribution is the identification of the correct signal model: under a logit-normal posterior over beliefs (Dalen [17]) and an affine logit-value link (Definition 4.1), the optimal coupling is the logit form $(E/(1-E))^{\alpha}$ (Theorem 4.3), not the power-law E^{κ} suggested by a naive log-normal analysis.

Mechanism design and auction theory. The first-order-condition arithmetic underlying Proposition 6.1 is the quadratic-payoff FOC of the Myerson revenue-maximisation programme [33], specialised to the coupled prediction-AMM setting. Our adversary objective subtracts the AMM round-trip slippage from the Kyle-style [11] revenue expression; the resulting depth requirement $b > \alpha Q^{\star}/(2\varepsilon)$ replaces the naive R_M-indexed bound. The clearing construction of Section 10 draws on the combinatorial-auction tradition (Cramton, Shoham, Steinberg [34]); specifically the convex-parimutuel architecture of Peters-So-Ye [7], which is adjacent to but distinct from the direct-revelation Vickrey-Clarke-Groves route.

1.3 Outline

Section 2 establishes an information-theoretic requirement on the execution environment: for the coupling to be meaningful, the event-outcome signal and the trade-set commitment must be informationally separated within the settlement window. Section 3 states and proves the main impossibility theorem. Section 4 characterises the admissible one-way couplings under Gaussian-conjugate signal models. Section 5 introduces the bounded-C^1-clipped class and proves the price-impact bound, with the non-clipping counterexample. Section 6 gives the tight two-sided cross-venue manipulation bound. Section 7 develops the AMM's invariant surface. Section 8 analyses the LP's loss, separating loss-versus-rebalancing from coupling loss. Section 9 quantifies the information value of coupling and establishes its structural first-order dominance. Section 10 presents the discrete and continuous clearing programs, proves the first-order fixed-point results for the power-law and logit regimes, derives the continuous Arrow-Debreu density, and shows that the coupled logit fixed point is stable under bounded-grid refinement. Section 11 compares against alternative coupling architectures. Section 13 lists open problems.

2 Event-Atomic Settlement: an Execution-Environment Assumption

The impossibility theorem of Section 3 concerns the mechanism-design layer: the incentive structure seen by traders facing a strictly-proper scoring rule. That theorem is independent of how the mechanism is executed. We record here the execution-layer condition that makes single-block atomic clearing meaningful, namely that event outcomes and trade-set commitments be informationally separated within the settlement window, and summarise the information-theoretic guarantee that holds under a list of proposer constraints together with an honest-attestor majority and side-channel hardening assumptions.

2.1 The commitment boundary

The minimal requirement is a single *commitment boundary*: a point in the settlement process after which the trade set is fixed and before which event outcomes are unknown to participants. This is a requirement on the execution environment, not a mechanism of the coupling itself.

Definition 2.1 (Information-gap independence)**.** Let $\mathcal{F}_u^{\text{trade}}$ denote the σ-algebra generated by trader u's trading strategy in block B (the set of trades submitted, their prices, and quantities). Let $\mathcal{G}_B^{\text{event}}$ denote the σ-algebra generated by the outcomes of all events resolving in block B. The block satisfies *information-gap independence* if, for every trader u (including the block proposer), the trading strategy σ_u is measurable with respect to a σ-algebra $\mathcal{F}^{\text{pre}}$ with $\mathcal{F}^{\text{pre}} \perp\!\!\!\perp \mathcal{G}_B^{\text{event}}$.

Remark 2.2 (Profit depends on events; strategy does not)**.** A trader's profit depends on both the strategy and the realised outcomes (through settlement); profits are correlated with outcomes. Definition 2.1 states only that the *strategy* cannot be adapted to the outcome.

2.2 Proposer constraints and side-channel hardening

Definition 2.3 (Proposer constraints)**.** A block-production protocol satisfies the proposer constraints if:

(PC1) **Pre-commitment.** The proposer's inclusion function $I : \mathcal{T}_{\text{pending}} \to \{0, 1\}$ is measurable with respect to a σ-algebra $\mathcal{F}^{\text{pre}}$ satisfying $\mathcal{F}^{\text{pre}} \perp\!\!\!\perp \mathcal{G}_B^{\text{event}}$.

(PC2) **Opacity.** Trade contents are encrypted under a threshold scheme requiring k-of-n attestors to decrypt, with decryption occurring only after commitment.

(PC3) **Censorship resistance.** Any valid trade submitted before deadline $\tau - \Delta$ must be included.

(PC4) **Self-binding.** The proposer's own trades are committed under the same opacity and pre-commitment constraints as all other trades.

Under IND-CPA threshold encryption, a minority-attestor adversary learns only negligible information from the ciphertext stream; additional side channels (padding, timing, metadata) require explicit hardening. We collect these as Remark 2.4.

Remark 2.4 (Side-channel hardening)**.** PC2 (opacity) as stated allows ciphertext-level distinguishers (length, arrival time, sender binding, fee parameters) to leak correlations with $\mathcal{G}_B^{\text{event}}$. The standard hardening assumes: (PC2-pad) uniform-length padding; (PC2-timing) arrival-time distribution independent of ω_e; (PC2-meta) no plaintext metadata correlated with ω_e; (PC2-crypto) IND-CPA security with advantage $\varepsilon_{\text{sec}} \leq \text{negl}(\lambda)$. Write $\varepsilon_{\text{total}} := \varepsilon_{\text{sec}} + \varepsilon_{\text{pad}} + \varepsilon_{\text{timing}}$.

Theorem 2.5 (Event-atomic settlement)**.** *Fix a (k, n) threshold encryption scheme with $k \leq n$ decryption shares required to reconstruct, and fix a set of attestors. Under PC1-PC4, an honest-attestor majority in the sense that at most $k - 1$ attestors collude (equivalently, honest attestors hold at least one of every k-subset), and the hardening of Remark 2.4, the following two statements hold:*

(i) Computational indistinguishability. *Every polynomial-time trader u's trading strategy is $\varepsilon_{\text{total}}$-computationally-indistinguishable from a strategy measurable only with respect to $\mathcal{F}^{\text{pre}}$, with $\varepsilon_{\text{total}}$ as in Remark 2.4.*

(ii) Information-gap independence at the measure-theoretic level. *The σ-algebra $\mathcal{F}^{\text{pre}}$ satisfies $\mathcal{F}^{\text{pre}} \perp\!\!\!\perp \mathcal{G}_B^{\text{event}}$ by Definition 2.1 (inherited from PC1's measurability condition).*

Proof. Statement (ii) is a restatement of the measurability condition embedded in PC1 and Definition 2.1; the proof content is (i).

We proceed by a hybrid / reduction argument. Let A be an arbitrary polynomial-time trader (possibly the proposer) with full $\mathcal{F}^{\text{pre}}$ access, trade-submission capability under PC1-PC4, and collusion with at most $k - 1$ attestors (the threshold scheme's safety margin; equivalently, the (k, n) scheme's secrecy threshold is $k - 1$).

Channel 1 (direct observation). By PC1 and PC4, A's trades and all other trades are committed at τ, before ω_e is revealed; A contributes zero advantage through direct observation of ω_e prior to commitment.

Channel 2 (inference from encrypted trades). By PC2 and the colluding-minority hypothesis, A holds $\leq k - 1$ shares, insufficient to decrypt any ciphertext. Under (PC2-crypto) the IND-CPA advantage of a polynomial-time distinguisher is $\varepsilon_{\text{sec}} \leq \text{negl}(\lambda)$; under (PC2-pad) and (PC2-meta) length and metadata contribute ε_{pad}; under (PC2-timing) arrival-time contributes $\varepsilon_{\text{timing}}$. Channel 2 total: $\varepsilon_{\text{total}}$.

Channel 3a (inclusion-function conditioning). By PC1, the proposer's inclusion function is measurable with respect to a σ-algebra $\mathcal{F}^{\text{pre}} \perp\!\!\!\perp \mathcal{G}_B^{\text{event}}$; so A (even if A is the proposer) cannot construct an $\mathcal{G}_B^{\text{event}}$-correlated inclusion function.

Channel 3b (selective exclusion of correlated trades). By PC3, any valid trade submitted before deadline $\tau - \Delta$ must be included; A cannot selectively exclude an honest trade whose content correlates with ω_e in order to manipulate the post-settlement state.

Combining channels gives total computational advantage $\leq \varepsilon_{\text{total}}$ over the $\mathcal{F}^{\text{pre}}$-measurable baseline strategy, which establishes (i).

The proof does not cover: out-of-band physical side-channels (e.g., observing the event in person before the oracle reports), cross-block strategies, or oracle manipulation (which makes ω_e endogenous). These are limitations of the model, not of the proof. □

Remark 2.6 (Scope of this section)**.** Theorem 2.5 is a standard result relying on IND-CPA threshold encryption plus attestor-honesty; the proof above is included for completeness of the paper's self-containedness. The result is an execution-layer *prerequisite* for the coupling theorems of Sections 3-9, not part of their content. Traders facing a coupling mechanism under a non-event-atomic execution environment face additional MEV and side-channel strategies beyond the scope of this paper.

2.3 Necessity of the proposer constraints and within-class priority

Proposition 2.7 (PC1-PC4 necessity)**.** *Each of PC1-PC4 blocks a distinct attack path: dropping any single constraint admits an explicit adversary strategy that violates Definition 2.1.*

Proof. One clause per dropped constraint. *(PC1).* If the proposer's inclusion function is not pre-committed, the proposer includes trade T iff ω_e is favourable, producing an $\mathcal{G}_B^{\text{event}}$-correlated block. *(PC2).* If trades are cleartext, an adversary infers the distribution over ω_e from peer trade signals and submits a posterior-conditioned trade. *(PC3).* If censorship is permitted, the proposer post-decryption excludes honest trades whose correlation with ω_e is unfavourable. *(PC4).* If the proposer is not self-bound, the proposer submits own trades after observing ω_e. Each construction yields a trader strategy measurably dependent on $\mathcal{G}_B^{\text{event}}$. □

Proposition 2.8 (Within-class priority: tight $1/m$ bound)**.** *Under PC1-PC4, suppose a block admits $m \geq 1$ identifier-symmetric trades of the same class (same pair, identical pre-commitment timing). The proposer holds one. Any $\mathcal{F}^{\text{pre}}$-measurable ordering rule yields expected proposer priority at most $1/m$; the random-within-class rule attains $1/m$ exactly. Hence within-class priority is $\Theta(1/m)$ with matching constant.*

Proof. The ordering rule is $\mathcal{F}^{\text{pre}}$-measurable (PC1, PC4) and acts on an identifier-symmetric set. By symmetry, the expected priority of any distinguished identifier is at most $1/m$ under any measurable rule, attained with equality by uniform random ordering. □

3 The Impossibility of Non-Trivial Bidirectional Coupling

This section contains the paper's central mechanism-design result. We show that any bidirectional coupling between a strictly-proper-scoring-rule prediction market and an AMM, under which the spot-market state enters the prediction payoff in a way that depends on the trader's reported belief, destroys incentive compatibility of the scoring rule. The necessary corollary is that admissible couplings are one-way.

3.1 Setup

We consider a coupled market $\mathcal{M} = (\mathcal{M}_{\text{pred}}, \mathcal{M}_{\text{spot}}, G)$ consisting of a prediction market, an AMM, and a coupling function G.

Definition 3.1 (Coupled prediction-spot market)**.** A coupled prediction-spot market is a triple $(\mathcal{M}_{\text{pred}}, \mathcal{M}_{\text{spot}}, G)$ where:

(M1) $\mathcal{M}_{\text{pred}}$ has a strictly proper base scoring rule $S(\hat{E}, \omega)$ and an optional bidirectional perturbation $H(\hat{E}, \omega, P)$ that is C^2 in $(\hat{E}, P)$:

$$\pi_{\text{pred}}(\hat{E}, \omega, P) = S(\hat{E}, \omega) + H(\hat{E}, \omega, P).$$

Writing $\Phi_0(\hat{E}) := \mathbb{E}_\omega[S(\hat{E},\omega)]$ and $\Gamma(\hat{E},P) := \mathbb{E}_\omega[H(\hat{E},\omega,P)]$, the base payoff Φ_0 is strictly concave in $\hat{E}$ and has unique interior maximiser $E^\star$.

(M2) $\mathcal{M}_{\text{spot}}$ is an AMM with marginal price $P = P(R_M, R_B, G(\hat{E}))$ satisfying $\partial P/\partial G \neq 0$.

(M3) $G : (0,1) \to (0,\infty)$ is a C^1 coupling function with $G'(\hat{E}) \neq 0$ on $(0,1)$.

3.2 The impossibility theorem

Theorem 3.2 (Impossibility of bidirectional coupling with incentive compatibility)**.** *Let $(\mathcal{M}_{\text{pred}}, \mathcal{M}_{\text{spot}}, G)$ be a coupled market in the sense of Definition 3.1. Suppose a trader can hold positions in both markets simultaneously. Then the following three properties cannot hold jointly:*

***(C1) Bidirectional coupling with report-sensitive prediction payoff**: both*

(C1a) *The spot price depends on $\hat{E}$ through G and the perturbation Γ is not identically independent of P.*

(C1b) *The report-sensitive feedback derivative*

$$D(P) := \partial_{\hat{E}}\Gamma(E^\star, P) + \partial_P\Gamma(E^\star, P)\,\frac{\partial P}{\partial G}(G(E^\star))G'(E^\star)$$

is not identically zero on any operating interval of spot prices. A sufficient local condition is a nonzero mixed partial $\partial_P\partial_{\hat{E}}\Gamma(E^\star, P)$ at an operating point.

***(C2) Incentive compatibility**: the trader's optimal report is $\hat{E} = E^\star$ regardless of their spot position.*

***(C3) Non-trivial spot sensitivity**: the trader's spot P&L π_{spot} satisfies $\partial\pi_{\text{spot}}/\partial\hat{E} \neq 0$ (the coupling has economic effect on spot-P&L via the report).*

Proof. Let the trader hold spot position Q and submit prediction report $\hat{E}$. The total payoff is

$$\Pi(\hat{E}) = S(\hat{E},\omega) + H(\hat{E},\omega,P(\hat{E})) + \pi_{\text{spot}}(Q, P(\hat{E})).$$

Taking expectations, the first-order condition at an interior optimum is

$$\underbrace{\Phi_0'(\hat{E})}_{\text{(A) base scoring rule}} + \underbrace{\partial_{\hat{E}}\Gamma(\hat{E},P(\hat{E})) + \partial_P\Gamma(\hat{E},P(\hat{E}))\cdot\frac{\partial P}{\partial G}\cdot G'(\hat{E})}_{\text{(B) report-sensitive spot}\to\text{prediction feedback}} + \underbrace{\frac{\partial\pi_{\text{spot}}}{\partial P}\cdot\frac{\partial P}{\partial G}\cdot G'(\hat{E})}_{\text{(C) spot-payoff channel}} = 0. \tag{1}$$

Set $Q = 0$ (so channel (C) vanishes). By strict properness of the base score, $\Phi_0'(E^\star) = 0$. Thus truthfulness at the base belief requires the feedback derivative to vanish:

$$D(P(E^\star)) \stackrel{!}{=} 0.$$

But condition (C1b) says exactly that this derivative is not identically zero as the operating price varies. Since D is continuous, the zero set has empty interior in every interval on which the non-degeneracy holds; under the displayed mixed-partial condition it is a codimension-one calibration set. On the complementary dense open set, $\hat{E} = E^\star$ is not even stationary and cannot be the unique optimum.

Quantitatively, applying the Implicit Function Theorem to the FOC (1) near a strictly concave base optimum gives

$$\hat{E}^\star(P(E^\star)) = E^\star - \frac{D(P(E^\star))}{\Phi_0''(E^\star)} + O(D(P(E^\star))^2),$$

which is distinct from $E^\star$ whenever $D(P(E^\star)) \neq 0$. Non-trivial spot sensitivity **(C3)** then adds channel (C) for $Q \neq 0$; even at an isolated calibration point where $D = 0$, a spot position with nonzero price exposure shifts the FOC unless that exposure is separately constrained to zero.

The three conditions **(C1)**, **(C2)**, and **(C3)** cannot jointly hold on any non-degenerate operating point. □

Remark 3.3 (Distinction from Lange-Economides 2005)**.** Lange and Economides [6] prove an impossibility in the parimutuel-clearing setting: adding an auxiliary payoff to the parimutuel mechanism breaks the self-consistency of the state-price vector under certain regularity conditions. Their result (a) is specific to the parimutuel call-auction and does not apply to continuous-time AMMs, (b) does not isolate the mixed-partial hypothesis that is central to our Theorem 3.2, and (c) addresses self-consistency of clearing rather than truthful reporting. Our theorem covers any C^2 strictly proper scoring rule composed with any AMM with $\partial P/\partial G \neq 0$, identifies the mixed-partial condition as load-bearing, and establishes impossibility of truthful reporting rather than of clearing. The Lange-Economides impossibility and Theorem 3.2 are complementary: they obstruct different composition architectures through different structural defects.

Remark 3.4 (Why report-sensitive feedback is load-bearing)**.** Without condition (C1b), the theorem would be false. A counter-example illustrating this is $H(\hat{E}, \omega, P) = c(P)$ for some fixed function c: the prediction venue receives a lump-sum subsidy depending on spot price, but the subsidy is additively separable from the report. Then $D(P) = 0$ for every P, so channel (B) vanishes identically and the base scoring rule remains truthful. Such additively separable transfers are admissible; the theorem's force is entirely against bidirectional couplings whose spot-price dependence changes the report derivative of the prediction payoff.

3.3 The necessity of one-way coupling

Corollary 3.5 (One-way coupling is necessary)**.** *For any coupling G between a strictly-proper-scoring-rule prediction market and an AMM, to preserve the scoring rule's incentive compatibility against traders holding positions in both markets, the coupling must be* one-way*: either (a) the spot price does not enter the prediction payoff (so $\partial \pi_{\text{pred}}/\partial P = 0$), or (b) the dependence is additively separable in $\hat{E}$ (so $\partial^2 \pi_{\text{pred}}/(\partial \hat{E}\, \partial P) = 0$). Any other configuration violates at least one of the three properties of Theorem 3.2.*

Proof. Contrapositive of Theorem 3.2: if the coupling is not one-way (neither (a) nor (b) holds), then **(C1)** is active; to preserve **(C2)** one must sacrifice **(C3)**, rendering the coupling economically trivial. □

Remark 3.6 (Interpretation)**.** Corollary 3.5 reframes one-way coupling as a mechanism-design necessity rather than a design choice. The one-directional architecture, under which prediction signals enter the AMM and spot state does not feed back into the scoring rule's payoff, is the unique admissible choice up to the additively-separable corner case of Remark 3.4. Information flows from prediction to AMM through G; information flows from AMM to prediction through the separate channel of informed trading (traders who believe the spot price is informative buy and sell prediction contracts accordingly), which operates through private beliefs rather than through the mechanism. The impossibility theorem rules out the mechanism channel; the informed-trading channel remains open and is addressed in Open Problem 5.

Definition 3.7 (One-way coupling, positive formulation)**.** A *one-way coupling* between a strictly-proper-scoring-rule prediction market and an AMM is a C^1 map $G : (0, 1) \to (0, \infty)$

with $G'(\hat{E}) \neq 0$ such that the prediction-market payoff $\pi_{\text{pred}}(\hat{E}, \omega)$ depends only on the reported belief $\hat{E}$ and the realised outcome ω, while the AMM's marginal price depends on the prediction signal through G. Equivalently, the oracle channel $\Theta : (\hat{E}, \omega) \mapsto \pi_{\text{pred}}(\hat{E}, \omega)$ induced by the scoring rule is non-invertible from the AMM state: no section $\Sigma : P \mapsto \pi'_{\text{pred}}$ exists returning an AMM-state-dependent payoff schedule that would admit the reverse information flow while preserving incentive compatibility. The additively-separable subcase of Corollary 3.5(b), in which the payoff depends on P only through a term constant in $\hat{E}$, is formally admissible but transmits no information between venues. All subsequent sections assume the positive-formulation definition.

The remainder of the paper studies the admissible class of one-way couplings.

4 Admissible One-Way Couplings: Representation Theorems

Corollary 3.5 establishes that admissible couplings are one-way. We now characterise which one-way couplings are statistically optimal in the sense of minimum-variance estimation of the latent fundamental value under standard Gaussian-conjugate signal models.

4.1 The signal-extraction problem

Consider a spot market for asset M with current AMM price P_{AMM}, and a prediction market for event e with current prediction signal E. Both are noisy observations of a latent fundamental value V.

A naive approach assumes $(\ln P_{\text{AMM}}, \ln E) \sim (\ln V + \varepsilon_A, \ln V^\star + \varepsilon_P)$ with independent Gaussian noise. The minimum-variance combination is the precision-weighted mean, which in multiplicative form yields the power-law coupling

$$G_{\text{pow}}(E) = E^\kappa, \qquad \kappa = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_P^2} \in [0, 1]. \tag{2}$$

The naive derivation has three structural defects: LMSR prediction-market prices follow *logit-normal* dynamics rather than log-normal (Dalen [17]); P_{AMM} and E are not independent observations of V because P_{AMM} already incorporates $G(E)$ through the coupling (a self-reference); and minimum-variance estimation is the wrong objective when LP losses are asymmetric. The last two are scope limitations that we acknowledge but do not resolve in this paper; the first admits a clean fix.

4.2 The affine logit-value link and logit coupling

Following Dalen [17], the natural domain for prediction-market analysis is logit space, where LMSR prices are Gaussian: $d(\operatorname{logit} E) = \mu\, dt + \sigma_b\, dW$. The derivation of the minimum-variance coupling in logit space requires a link function specifying how the fundamental value relates to the true event probability.

Definition 4.1 (Affine logit-value link)**.** The fundamental value V and the true event probability $E^\star$ are related by

$$\ln V = a + s_V \cdot \operatorname{logit}(E^\star), \tag{3}$$

where $a \in \mathbb{R}$ is a base log-price and $s_V > 0$ is the *price sensitivity*: the elasticity of fundamental value with respect to the log-odds $\operatorname{logit}(E^\star) = \ln(E^\star/(1 - E^\star))$.

The affine logit-value link is the unique link under which both signals remain Gaussian after projection into a common space.

Proposition 4.2 (Logit-space minimum-variance coupling)**.** *Assume: (i)* $\ln P_{\mathrm{AMM}} = \ln V + \varepsilon_A$*,* $\varepsilon_A \sim \mathcal{N}(0, \sigma_A^2)$*; (ii)* $\operatorname{logit} E = \operatorname{logit} E^\star + \varepsilon_P$*,* $\varepsilon_P \sim \mathcal{N}(0, \sigma_P^2)$*; (iii) the affine link* (3)*; (iv)* $\varepsilon_A \perp\!\!\!\perp \varepsilon_P$*. Then the minimum-variance estimator of* $\ln V$ *given both signals is*

$$\widehat{\ln V} = (1 - \kappa_\ell)\, \ln P_{\mathrm{AMM}} + \kappa_\ell (a + s_V \cdot \operatorname{logit} E), \tag{4}$$

where $\kappa_\ell = \sigma_A^2 / (\sigma_A^2 + s_V^2 \sigma_P^2) \in [0, 1]$*. The induced coupling (absorbing* $e^{\kappa_\ell a}$ *into the base price) is*

$$G_\ell(E) = \left(\frac{E}{1 - E} \right)^{\alpha}, \qquad \alpha = \kappa_\ell \cdot s_V. \tag{5}$$

Proof. Under (i)-(iv), define Gaussian signals $X_1 = \ln P_{\mathrm{AMM}} \sim \mathcal{N}(\ln V, \sigma_A^2)$ and $X_2 = a + s_V \operatorname{logit} E \sim \mathcal{N}(\ln V, s_V^2 \sigma_P^2)$ (the second follows from the affine link). The minimum-variance combination of independent Gaussian signals for a common mean is the precision-weighted mean, with weight $\kappa_\ell = \sigma_A^2 / (\sigma_A^2 + s_V^2 \sigma_P^2)$ on X_2. Substituting gives (4); taking exponentials gives $\ln G_\ell(E) = \alpha \operatorname{logit} E = \alpha \ln(E/(1 - E))$, i.e. (5). ☐

Theorem 4.3 (Optimality of the logit coupling under the affine link)**.** *Under the assumptions of Proposition 4.2, among* C^1 *coupling functions* $G : (0, 1) \to (0, \infty)$ *such that* $\ln G(E)$ *is the minimum-variance combination of the AMM and prediction signals under logit-normal signal noise and complement symmetry (*$G(1 - E) = 1/G(E)$*), the coupling is uniquely* $G^\star(E) = (E/(1 - E))^\alpha$ *with* $\alpha = \kappa_\ell \cdot s_V$*. Moreover, the power-law plug-in estimator* $\widehat{\ln V}_{\mathrm{pow}} := (1 - \kappa_\ell) \ln P_{\mathrm{AMM}} + \kappa_\ell (a + s_V \kappa \ln E)$*, which treats the non-Gaussian signal* $\kappa \ln E$ *as if it were Gaussian for* $\ln V$*, has strictly larger posterior mean-square error than the logit-optimal estimator of Proposition 4.2 on every non-degenerate prior on* $\ln V$*.*

Proof. Uniqueness. Under logit-normal noise and the affine link, both $\ln P_{\mathrm{AMM}}$ and $a + s_V \operatorname{logit} E$ are Gaussian signals for $\ln V$; the minimum-variance combination (precision-weighted mean) is unique, giving $\ln G^\star(E) = \alpha \operatorname{logit} E$. Complement symmetry: $G^\star(1 - E) = ((1 - E)/E)^\alpha = 1/G^\star(E)$.

Suboptimality of the power-law plug-in. Under Assumption (ii) of Proposition 4.2, $\operatorname{logit} E$ is Gaussian for $\operatorname{logit} E^\star$, so $Y := a + s_V \operatorname{logit} E$ is Gaussian for $\ln V$. The power-law transform $Z := \kappa \ln E$ of the same signal is not an affine function of $\ln V$ (because $\ln E = \operatorname{logit} E - \ln(1 + e^{\operatorname{logit} E})$ is nonlinear in $\operatorname{logit} E$). The Gauss-Markov-optimal precision-weighted combiner of $\ln P_{\mathrm{AMM}}$ with Y achieves the Cramér-Rao bound; substituting Z for Y breaks the combiner's optimality and yields strictly larger posterior MSE on every prior where $\operatorname{logit} E$ is not a point mass. ☐

Remark 4.4 (Scope of the suboptimality claim)**.** The suboptimality refers to posterior MSE of $\ln V$, not to $\operatorname{Var}(\ln G(E))$ as a function of E. Under particular priors on E, $\operatorname{Var}(\ln G_{\mathrm{pow}}(E)) < \operatorname{Var}(\ln G^\star(E))$ can hold without contradicting the theorem.

Remark 4.5 (Scope of the uniqueness claim)**.** Theorem 4.3 establishes uniqueness of the minimum-variance *estimator* under an affine logit-value link. Under a different link (for example, payoffs that are affine in E rather than in $\operatorname{logit} E$, as in binary-outcome instruments), the optimal coupling differs; see Proposition 4.6.

4.3 Binary-payoff coupling via EKF

A common special case is the binary payoff $V(E) = V_0 + (V_1 - V_0)E$ with $V_1 > V_0 > 0$. Here V is affine in E but nonlinear in $\operatorname{logit} E$, so the affine logit-value link of Definition 4.1 does not hold globally. First-order Taylor linearisation around the operating point (standard extended Kalman filter construction) gives a state-dependent coupling.

Proposition 4.6 (Binary-payoff coupling via EKF)**.** *Let $V(E) = V_0 + (V_1 - V_0)E$ with $V_1 > V_0 > 0$. Suppose* $\operatorname{logit} E \sim \mathcal{N}(\operatorname{logit} E^\star, \sigma_P^2)$ *and* $\ln P_{\text{AMM}} \sim \mathcal{N}(\ln V(E^\star), \sigma_A^2)$ *independently. The extended Kalman filter linearisation gives*

$$G(E) = \exp(\alpha(E) \cdot \operatorname{logit} E), \qquad \alpha(E) = \kappa_\ell \cdot (V_1 - V_0) \cdot \frac{E(1-E)}{V(E)}, \tag{6}$$

with $\kappa_\ell = \sigma_A^2/(\sigma_A^2 + (s_V^\star)^2\sigma_P^2)$ and $s_V^\star = (V_1 - V_0) \cdot E^\star(1-E^\star)/V(E^\star)$ the linearisation slope at the operating point.

Boundary behaviour. *As $E \to 0^+$ or $E \to 1^-$, $\alpha(E) \to 0$; by L'Hôpital's rule, $\alpha(E) \cdot \operatorname{logit} E \to 0$. So $G(E) \to 1$ at both boundaries, extending continuously to $G(0) = G(1) = 1$.*

Proof. Step 1: observation Jacobian. The observation equation relating $\ln V$ to $u := \operatorname{logit} E$ is $g(u) := \ln V(\sigma(u)) = \ln(V_0 + (V_1 - V_0)\sigma(u))$ where σ is the logistic function. Differentiating,

$$g'(u) = \frac{(V_1 - V_0)\sigma(u)(1-\sigma(u))}{V(\sigma(u))},$$

which evaluated at $u^\star = \operatorname{logit} E^\star$ gives the scalar Jacobian $s_V^\star = (V_1 - V_0)E^\star(1-E^\star)/V(E^\star)$.

Step 2: EKF Kalman gain. For a scalar prior $u \sim \mathcal{N}(u_0, \sigma_P^2)$ with observation $z = Hu + v$, $v \sim \mathcal{N}(0, \sigma_A^2)$ (independent) with $H = s_V^\star$ (frozen at operating point), the Kalman gain is

$$K = \frac{\sigma_P^2 s_V^\star}{(s_V^\star)^2\sigma_P^2 + \sigma_A^2}, \qquad 1 - HK = \kappa_\ell := \frac{\sigma_A^2}{\sigma_A^2 + (s_V^\star)^2\sigma_P^2}.$$

Step 3: state-dependent coupling. The posterior correction relative to the pre-existing AMM observation is multiplied by the residual prior coefficient $1 - HK = \kappa_\ell$, not by the raw Kalman gain K itself. The leading-order log-price adjustment is therefore $\kappa_\ell \cdot s_V^\star \cdot \operatorname{logit} E$. Re-linearising at each operating point (standard EKF practice) yields the state-dependent gain $\alpha(E) = \kappa_\ell \cdot g'(u)|_{u=\operatorname{logit} E}$, which is (6).

Boundary: $E(1-E)/V(E) \to 0$ at both $E \to 0^+$ and $E \to 1^-$, while $\operatorname{logit} E$ diverges only logarithmically; the product $\alpha(E) \cdot \operatorname{logit} E \to 0$ by L'Hôpital. $\square$

Remark 4.7 (EKF approximation error)**.** The EKF is a first-order linearisation of $\ln V(\sigma(u))$. Writing $u = \operatorname{logit} E$ and $g(u) = \ln V(\sigma(u))$, Taylor's theorem with Lagrange remainder gives

$$|g(u) - g(u^\star) - g'(u^\star)(u - u^\star)| \le \tfrac{1}{2} \sup_\xi |g''(\xi)| \cdot (u - u^\star)^2.$$

Direct computation gives $g''(u) = (V_1 - V_0)\sigma(u)(1-\sigma(u))(1-2\sigma(u))/V(\sigma(u)) - (g'(u))^2$. For a bound, observe $\sup_u |\sigma(u)(1-\sigma(u))(1-2\sigma(u))| = 1/(6\sqrt{3})$ (attained at $\sigma(u) = 1/2 \pm \sqrt{3}/6$), and $V(\sigma(u)) \ge V_0$; and $|g'(u)| \le (V_1 - V_0)/(4V_0)$ (via $\sup|\sigma(1-\sigma)| = 1/4$). Thus

$$|g''(u)| \le \frac{V_1 - V_0}{6\sqrt{3}\, V_0} + \left(\frac{V_1 - V_0}{4V_0}\right)^2. \tag{7}$$

The EKF coupling error in log-price space is bounded by $(\kappa_\ell/2) \cdot |g''|_{\text{sup}} \cdot (\Delta \operatorname{logit} E)^2$. For moderate payoff ratios $(V_1 - V_0)/V_0 \lesssim 1/2$, this is a small-coefficient quadratic correction; the EKF is most accurate near the boundaries (where $|g''| \to 0$) and least accurate at $E = 1/2$.

Corollary 4.8 (Local optimality of the binary EKF coupling)**.** *Under Proposition 4.6, the state-dependent coupling $G(E) = \exp(\alpha(E) \operatorname{logit} E)$ is locally optimal among first-order linearisations of the nonlinear observation $\ln V = g(\operatorname{logit} E)$ at $u^\star = \operatorname{logit} E^\star$: among affine maps $u \mapsto \alpha u + \beta$, the EKF choice minimises posterior MSE to first order in $(u - u^\star)$. Global optimality over nonlinear couplings is not claimed; the gap to the exact MMSE estimator is bounded by (7). This is the standard scalar-EKF local-optimality result (Anderson-Moore [32], Jazwinski [31]).*

5 Bounded Parametrizations of One-Way Couplings

Sections 3-4 establish that admissible couplings are one-way and characterise the statistically optimal choices under Gaussian-conjugate signal models. A separate requirement is that the coupling yield *bounded single-block price impact* against arbitrary prediction-signal realisations. Without a boundedness constraint the minimum-variance couplings of Section 4 admit arbitrarily large price impact, enabling adversarial oracle manipulation. This section formalises the boundedness requirement and proves that a specific class of clipped couplings, of which the Huber-bounded coupling is one canonical representative, achieves the bound.

5.1 The bounded C^1-clipped class

Definition 5.1 (Bounded C^1-clipped coupling)**.** A *bounded C^1-clipped coupling* on $E \in \mathcal{I}$ (where $\mathcal{I}$ is either $(0, \infty)$ or $(0, 1)$) is a map $\tilde{G} : \mathcal{I} \to \mathbb{R}_{>0}$ of the form

$$\tilde{G}(E) = \exp(\kappa \cdot \phi(\ell(E))), \qquad \ell(E) := \begin{cases} \ln E & \mathcal{I} = (0, \infty), \\ \operatorname{logit} E & \mathcal{I} = (0, 1), \end{cases}$$

where $\kappa > 0$ and the *clipper* $\phi : \mathbb{R} \to [-\delta_{\max}, \delta_{\max}]$ satisfies, for some $\delta_{\max} > 0$ and transition window $w > 0$:

(C1) *Identity on the centre:* $\phi(y) = y$ for $|y| \leq \delta_{\max} - w$.

(C2) *Saturation on the tails:* $\phi(y) = \operatorname{sgn}(y) \cdot \delta_{\max}$ for $|y| \geq \delta_{\max} + w$.

(C3) *C^1 smoothness:* $\phi \in C^1(\mathbb{R})$.

(C4) *Monotonicity:* ϕ is non-decreasing on $\mathbb{R}$.

Example 5.2 (Canonical instance: the 3-piece sign-mirrored Huber clipper)**.** The canonical element of Definition 5.1 is the 3-piece sign-mirrored Huber function

$$h(y) = \begin{cases} y & |y| \leq \delta - w, \\ \operatorname{sgn}(y) \left[\delta - (\delta + w - |y|)^2 / (4w) \right] & \delta - w < |y| < \delta + w, \\ \operatorname{sgn}(y) \cdot \delta & |y| \geq \delta + w, \end{cases} \tag{8}$$

which composes three formulas (linear, quadratic transition, saturation) with sign-mirroring to produce four transition breakpoints at $|y| = \delta \pm w$. Other admissible clippers (cubic Hermite interpolants, smoothed-max bridges) lie within Definition 5.1 and produce the same bound below.

Lemma 5.3 (C^1 smoothness of the Huber clipper)**.** *The function h of (8) is C^1 on $\mathbb{R}$: values and one-sided derivatives agree at each of the four breakpoints $\pm(\delta - w)$ and $\pm(\delta + w)$.*

Proof. The odd symmetry $h(-y) = -h(y)$ reduces the check to the positive side. At $y = \delta - w$: from the linear piece, $h = \delta - w$ and $h' = 1$; from the quadratic piece, $h = \delta - (2w)^2/(4w) = \delta - w$ and $h' = (\delta + w - y)/(2w)$ evaluating to 1. At $y = \delta + w$: from the quadratic piece, $h = \delta$ and $h' = 0$; from the saturation piece, $h = \delta$ and $h' = 0$. Both breakpoints match. Odd extension gives the negative-side checks. $\square$

5.2 The price-impact bound

Lemma 5.4 (Bounded coupling impact: generalised C^1-clipped class). *Let $\tilde{G}$ be any bounded C^1-clipped coupling (Definition 5.1) with parameters $\kappa, \delta_{\max}, w$. At fixed reserves, the marginal price $P = \tilde{G}(E) \cdot R_M / R_B$ satisfies, for any single-block shift from $g = 1$ (uncoupled baseline) to $g = \tilde{G}(E)$,*

$$\frac{\Delta P}{P} \in \left[-(1 - e^{-\kappa\delta_{\max}}),\ e^{\kappa\delta_{\max}} - 1\right], \qquad \text{i.e.} \left|\frac{\Delta P}{P}\right| \le e^{\kappa\delta_{\max}} - 1, \tag{9}$$

uniformly for every $E \in \overline{\mathcal{I}}$ (the closed domain $[0, \infty]$ or $[0, 1]$). The bound depends on ϕ only through $\kappa\delta_{\max}$ and is independent of the transition-window shape.

Proof. By (C2), $|\phi(y)| \le \delta_{\max}$ globally, so $|\ln \tilde{G}(E)| \le \kappa\delta_{\max}$. At fixed reserves, $\Delta P / P = e^{\ln g} - 1$ where $g = \tilde{G}(E)$. Using $|e^x - 1| \le e^{|x|} - 1$ (immediate for $x \ge 0$, and for $x < 0$ follows from $e^x + e^{-x} \ge 2$), $|\Delta P / P| \le e^{|\ln g|} - 1 \le e^{\kappa\delta_{\max}} - 1$. Extension to the closed domain: on the tails ϕ is eventually constant, and the boundary values $\tilde{G}(0) = e^{-\kappa\delta_{\max}}$, $\tilde{G}(\infty)$ or $\tilde{G}(1) = e^{+\kappa\delta_{\max}}$ coincide with the one-sided limits; the bound holds with equality at the upper end. □

Corollary 5.5 (Representatives: power-law and logit in closed form). *Specialising Lemma 5.4: the Huber-bounded power-law $\tilde{G}_{\text{pow}}(E) = \exp(\kappa \cdot h(\ln E))$ on $\mathcal{I} = (0, \infty)$ satisfies (9) with exponent κ; the Huber-bounded logit $\tilde{G}_\ell(E) = \exp(\alpha \cdot h(\operatorname{logit} E))$ on $\mathcal{I} = (0, 1)$ satisfies (9) with exponent $\alpha = \kappa_\ell \cdot s_V$ (Proposition 4.2).*

Theorem 5.6 (Tight-clipping characterisation). *The upper endpoint $e^{\kappa\delta_{\max}} - 1$ of (9) is attained at some realisation $E^\dagger \in \overline{\mathcal{I}}$ if and only if there exists $E^\dagger$ with $\phi(\ell(E^\dagger)) = +\delta_{\max}$. By (C2) both endpoints are attained at the closed-domain extremes, so the bound is tight on $\overline{\mathcal{I}}$. On the open domain $\mathcal{I}$, tightness is attained if the tail-saturation regime $|\ell(E)| \ge \delta_{\max} + w$ is reached by some finite $E^\dagger$.*

Proof. Immediate from the equality case of the inequality $|\ln g| \le \kappa\delta_{\max}$. □

5.3 The necessity of clipping: unbounded couplings break the bound

Theorem 5.7 (Unbounded power-law is not price-impact-stable). *Let $G_{\text{unc}}(E) = E^\kappa$ be the unclipped power-law coupling on $E \in (0, \infty)$ (condition (C2) of Definition 5.1 dropped). Then $E \mapsto \Delta P / P$ is unbounded above and below: for every $M > 0$ there exist E_+, E_- with $\Delta P / P \ge M$ and $\Delta P / P \le -1 + \epsilon$ for any $\epsilon > 0$. No finite bound of the form $|\Delta P / P| \le B$ holds.*

Proof. $\Delta P / P = E^\kappa - 1$; take $E_+ = (M + 1)^{1/\kappa}$ and $E_- = \epsilon^{1/\kappa}$. □

Corollary 5.8 (Clipping is essential). *Under the requirement of bounded price-impact for all admissible E, the saturation condition (C2) of Definition 5.1 cannot be dropped. Within the C^1-clipped class, the specific clipper shape is free (any admissible interpolant gives the same bound); what is essential is the existence of a saturation radius $\delta_{\max} + w$ beyond which ϕ is constant.*

5.4 Composition across blocks

Theorem 5.9 (Composed C^1-clipped coupling across blocks). *Let $\tilde{G}_1, \ldots, \tilde{G}_n$ be bounded C^1-clipped couplings with (possibly distinct) parameters $(\kappa_i, \delta_{\max,i}, w_i)$. For a block-sequence $(E_1, \ldots, E_n)$, the cumulative price-impact satisfies*

$$\left|\frac{\Delta P^{(n)}}{P^{(0)}}\right| \le \exp\left(\sum_{i=1}^{n} \kappa_i \delta_{\max,i}\right) - 1. \tag{10}$$

The bound is tight when every block saturates at the positive tail.

Proof. Each block contributes $|\kappa_i\phi_i(\ell(E_i))| \leq \kappa_i\delta_{\max,i}$; the cumulative log-price shift is bounded by $\sum \kappa_i\delta_{\max,i}$; exponentiating and using $|e^x - 1| \leq e^{|x|} - 1$ gives (10). □

Corollary 5.10 (Stationary vs. adversarial asymptotics). *Let $\kappa_i = \kappa$, $\delta_{\max,i} = \delta_{\max}$ for all i. Then:* (i) finite-horizon worst-case: $|\Delta P^{(n)}/P^{(0)}| \leq e^{n\kappa\delta_{\max}} - 1$. (ii) adversarial divergence: *for fixed $\kappa\delta_{\max} > 0$, $e^{n\kappa\delta_{\max}} - 1 \to \infty$.* (iii) stationary regime: *if (E_i) is drawn from a stationary process with $\mathbb{E}[\kappa_i\phi_i(\ell(E_i))] = 0$ (zero-mean coupling), the cumulative log-shift is a zero-mean random walk with step variance $\sigma^2 \leq \kappa^2\delta^2_{\max}$, so by the CLT $\ln G^{(n)}/\sqrt{n} \Rightarrow \mathcal{N}(0, \sigma^2)$ and the cumulative price-impact grows as $O(\sqrt{n})$.*

Remark 5.11 (Long-horizon adversarial stability is out of scope). The per-block bound does not guarantee long-horizon adversarial stability: a bounded-per-step map can compound without bound unless the step distribution has mean-reversion or rate-limiting structure. Long-horizon stability against a persistent adversary is a separate question, addressed within this paper only at the manipulation-bound level (Section 6).

6 Cross-Venue Manipulation

Cross-venue arbitrage is inherent to any deterministic coupling between a prediction market and a spot venue. A trader who moves the prediction price affects the spot price, creating a manipulation channel. This section bounds the equilibrium bias induced by a risk-neutral manipulator. Our bookkeeping follows the oracle-manipulation framework of Angeris and Chitra [26], specialised to the LMSR-coupled setting and augmented with the AMM-side round-trip slippage.

6.1 Adversary model and equilibrium bias

Consider a one-way logit coupling $G_\ell(E) = (E/(1-E))^\alpha$ (Proposition 4.2) with AMM reserves (R_M, R_B) and a co-located LMSR prediction market with liquidity parameter b. The adversary's gross strategy is: (a) pre-position notional Q on the spot side at the uncoupled price; (b) move the LMSR price by δ logits; (c) close the spot position at the coupled price. Three distinct frictions apply.

- (F1) *LMSR round-trip cost.* Moving the LMSR price by δ logits and restoring it costs $bE^\star(1-E^\star)\delta^2 + O(\delta^3)$, derived from twice the one-way KL divergence.
- (F2) *AMM price impact.* A pre-positioned trade of notional Q on the CFMM invariant $R_M R_B^g = k$ incurs round-trip slippage of leading order $Q^2/(2R_B)$ on small trades.
- (F3) *Coupling-induced price change.* The coupling shift δ moves the coupled spot price by a factor $e^{\alpha\delta}$, delivering gross profit $Q(e^{\alpha\delta} - 1) = Q\alpha\delta + O(\delta^2)$ on the pre-positioned trade.

Proposition 6.1 (Equilibrium prediction bias). *Consider a risk-neutral manipulator with pre-positioned spot notional Q and logit shift δ, facing LMSR depth b and AMM reserves (R_M, R_B). The small-shift net objective (with F2 bookkeeping) is*

$$\Pi(\delta, Q) = \alpha Q\delta - bE^\star(1-E^\star)\delta^2 - \frac{Q^2}{2R_B} + O(\delta^3 + Q^3/R_B^2). \tag{11}$$

Jointly optimising in (δ, Q):

(i) The first-order conditions give $\delta^\star(Q) = \alpha Q/(2bE^\star(1-E^\star))$ and $Q^\star = \alpha R_B\delta^\star = \alpha^2 R_B Q^\star/(2bE^\star(1-E^\star))$, with non-degenerate optimum when $\alpha^2 R_B < 2bE^\star(1-E^\star)$.

(ii) *At the joint optimum the equilibrium probability bias is* $|E_{\mathrm{eq}} - E^\star| = \alpha Q^\star/(2b) + O((\alpha Q^\star/b)^2)$.

(iii) *The bias bound* $|E_{\mathrm{eq}} - E^\star| < \varepsilon$ *is enforced by the design-time depth requirement* $b > \alpha Q^\star/(2\varepsilon)$.

Proof. The three frictions (F1)-(F3) give (11). Differentiating in δ and Q separately gives the FOCs in (i); substituting $|E_{\mathrm{eq}} - E^\star| = E^\star(1-E^\star)\delta^\star + O((\delta^\star)^2)$ gives (ii); (iii) is rearrangement. □

Remark 6.2 (Why $Q^\star$, not R_M)**.** The raw-reserve bound $b > \alpha R_M/(2\varepsilon)$ that omits (F2) is the defender's revenue-side bound; it overstates the required depth because the adversary cannot take the full reserve as notional without paying AMM slippage. The correct parameter is the adversary's self-chosen notional $Q^\star$, which is strictly less than R_M. The result is consistent with the Angeris-Chitra [26] observation that CFMM oracle-manipulation bounds tighten once the adversary's own trade impact is accounted for.

Remark 6.3 (Linearisation error)**.** The first-order bias formula $\alpha Q^\star/(2b)$ is a linearisation artefact; its $E^\star$-independence cancels across the FOC and the logit-to-probability map. Expanding to second order, $E_{\mathrm{eq}} - E^\star = E^\star(1-E^\star)\delta^\star + \frac{1}{2}E^\star(1-E^\star)(1-2E^\star)(\delta^\star)^2 + O((\delta^\star)^3)$. The second-order term vanishes at $E^\star = 1/2$ and grows linearly in $(1-2E^\star)$ toward the boundaries; at $E^\star = 0.9$, $\delta^\star = 0.3$ it is a $\sim 3\%$ relative correction to the first-order estimate.

Remark 6.4 (Attainability)**.** The FOC solution $(\delta^\star, Q^\star)$ of Proposition 6.1(i) is the interior maximiser of the small-shift quadratic objective (11) with strictly positive net profit. An explicit adversary using $(\delta^\star, Q^\star)$ thus attains the bias $\alpha Q^\star/(2b)$ constructively, rendering the depth requirement $b > \alpha Q^\star/(2\varepsilon)$ a necessary design-time invariant. Extension beyond the small-shift linearisation (where the adversary drives the coupling into the Huber-saturation band) is an open problem (Open Problem 7).

6.2 Persistent-adversary upper bound under rate-limited calibrator

Proposition 6.1 bounds single-block extraction. A persistent adversary who sustains a loss on the prediction market over T blocks can, in principle, continuously bias E away from $E^\star$, extracting cumulative spot profit that grows linearly in T while paying per-block LMSR loss that also grows linearly in T. Whether net extraction is positive depends on the rate at which the defender's calibrator κ_t adapts to the bias. The following proposition states the upper bound under an exponentially-weighted-moving-average (EMA) calibrator response with timescale τ_{cal}.

Proposition 6.5 (Persistent-manipulation upper bound, EMA calibrator)**.** *Consider the setup of Proposition 6.1 with a persistent adversary sustaining a constant logit shift* $\delta^{\mathrm{adv}} = \alpha R_M/(4bE^\star(1-E^\star))$ *across* T *blocks. Let the coupling calibrator* κ_t *evolve under an EMA update with timescale* $\tau_{\mathrm{cal}} \geq 1$ *blocks:*

$$\kappa_{t+1} = \kappa_t - \frac{1}{\tau_{\mathrm{cal}}} \cdot \left(\kappa_t \cdot |E_t - E^\star|_{\mathrm{obs}}\right), \tag{12}$$

where $|E_t - E^\star|_{\mathrm{obs}}$ *is the venue-observable deviation at block* t*, starting from* κ_0 *at block* $t = 0$*. Assume the adversary's per-block LMSR round-trip cost is* $c_{\mathrm{adv}} := bE^\star(1-E^\star)(\delta^{\mathrm{adv}})^2$*, and the per-block spot extraction is at most* $\alpha R_M \delta^{\mathrm{adv}}$ *(from the gross-profit term of Proposition 6.1). Then:*

(i) *(Calibrator decay.) Under the adversary's sustained shift, the calibrator decays as* $\kappa_t \leq \kappa_0 \cdot \exp(-t|E^\star_{\mathrm{bias}}|/\tau_{\mathrm{cal}})$ *with* $|E^\star_{\mathrm{bias}}| := \alpha R_M/(4b)$.

(ii) (Cumulative extraction bound.) Total adversary net extraction over T blocks is bounded above by

$$\Pi_{\text{pers}}(T) \;\leq\; \alpha R_M \delta^{\text{adv}} \cdot \tau_{\text{cal}} / |E^\star_{\text{bias}}| \cdot \left(1 - e^{-T|E^\star_{\text{bias}}|/\tau_{\text{cal}}}\right) - T \cdot c_{\text{adv}}. \tag{13}$$

(iii) (Break-even condition.) Net extraction is negative (adversary unprofitable over any T) whenever

$$\tau_{\text{cal}} \;<\; \tau^\star_{\text{cal}} \;:=\; \frac{c_{\text{adv}} \cdot |E^\star_{\text{bias}}|}{\alpha R_M \delta^{\text{adv}}} \;=\; \frac{\alpha R_M (\delta^{\text{adv}})^2/4}{\alpha R_M \delta^{\text{adv}}/(4E^\star(1-E^\star))} \;=\; \delta^{\text{adv}} \cdot E^\star(1-E^\star). \tag{14}$$

(iv) (Matching lower bound.) When $\tau_{\text{cal}} \geq \tau^\star_{\text{cal}}$, the adversary of Proposition 6.1 achieves cumulative extraction growing at rate $\Omega((\alpha R_M)^2/b)$ per block, so the persistent-adversary bound and the single-block lower bound are matching up to the constant $\tau_{\text{cal}}/\tau^\star_{\text{cal}}$.

Proof sketch. (i) The EMA update (12) at sustained bias $|E_t - E^\star|_{\text{obs}} = |E^\star_{\text{bias}}|$ reduces to the linear ODE $\dot{\kappa} = -\kappa|E^\star_{\text{bias}}|/\tau_{\text{cal}}$ in the continuous-time limit, with solution $\kappa_t = \kappa_0 e^{-t|E^\star_{\text{bias}}|/\tau_{\text{cal}}}$. The per-block discrete update is bounded above by the continuous-time solution.

(ii) Integrating the per-block spot extraction $\kappa_t \cdot \alpha R_M \delta^{\text{adv}}/\kappa_0$ over $[0, T]$ under (i) gives the first term of (13); subtracting the $T \cdot c_{\text{adv}}$ cost gives the bound.

(iii) Setting $\lim_{T\to\infty} \Pi_{\text{pers}}(T) \leq 0$ yields $\alpha R_M \delta^{\text{adv}} \cdot \tau_{\text{cal}}/|E^\star_{\text{bias}}| \leq (T \cdot c_{\text{adv}})$ asymptotically. Solving for the critical $\tau^\star_{\text{cal}}$ gives (14); the algebraic simplification uses $c_{\text{adv}} = bE^\star(1-E^\star)(\delta^{\text{adv}})^2$ and $|E^\star_{\text{bias}}| = E^\star(1-E^\star)\delta^{\text{adv}}$ to leading order.

(iv) At $\tau_{\text{cal}} = \tau^\star_{\text{cal}}$, the upper bound in (ii) equals the single-block extraction rate from Proposition 6.1 summed over T blocks, so the persistent-adversary upper bound is tight up to the geometric $(\tau_{\text{cal}}/\tau^\star_{\text{cal}})$ factor. □

Remark 6.6 (Design-space interpretation)**.** Proposition 6.5 operationalizes the open problem flagged in Open Problem 8: the EMA calibrator neutralizes a persistent adversary whenever τ_{cal} (the calibrator's response time) is shorter than the adversary's per-bias break-even horizon $\tau^\star_{\text{cal}}$. The critical timescale $\tau^\star_{\text{cal}} = \delta^{\text{adv}} \cdot E^\star(1-E^\star)$ depends only on the adversary's chosen shift magnitude and the base-rate $E^\star$, not on the adversary's duration T; the defender's response time τ_{cal} is a design parameter. A full adaptive-control treatment under noisy observations (where $|E_t - E^\star|_{\text{obs}}$ contains measurement error) is outside the scope of this paper and remains flagged in the open-problem structure (Open Problem 8); the result here is the idealized-observation upper bound.

6.3 Oracle-manipulation attack upper bound

A related attack surface is oracle manipulation: an attacker who controls (bribes) the event outcome itself can shift E to an extreme value, pre-position spot, and extract value. The single-event extraction is bounded by the coupling's Huber cap:

Proposition 6.7 (Oracle-manipulation upper bound)**.** *Under a Huber-bounded coupling with exponent κ and saturation amplitude $\delta_{\max}$, the maximum single-event extraction by an attacker who forces full saturation is*

$$\text{MEV}_{\text{oracle}} \leq R_M \cdot (e^{\kappa\delta_{\max}} - 1 - \kappa\delta_{\max}) \approx \frac{(\kappa\delta_{\max})^2}{2} \cdot R_M. \tag{15}$$

The attack is profitable iff the extraction exceeds the sum of oracle bribe cost B, AMM trading fees and slippage, and LMSR round-trip $bE^\star(1-E^\star)\delta^2$.

Proof. Under saturation, $\tilde{G}(E) = e^{\pm\kappa\delta_{\max}}$ and $\Delta P/P = e^{\kappa\delta_{\max}} - 1$. For a pre-positioned trade of value R_M, the payoff in excess of uncoupled is $R_M(e^{\kappa\delta_{\max}} - 1)$, minus the linear cost of movement $\kappa\delta_{\max} R_M$, giving the quadratic upper bound (15). □

Remark 6.8 (Huber bounding is the defence)**.** The Huber saturation at $\kappa\delta_{\max}$ caps the single-event extraction at $\approx (\kappa\delta_{\max})^2/2 \cdot R_M$ regardless of oracle-manipulation magnitude. Without the saturation condition (C2) of Definition 5.1, Theorem 5.7 shows the extraction is unbounded. The saturation is thus the structural defence against oracle manipulation: bounded impact independent of the signal's extreme values.

6.4 Persistent multi-block manipulation: matching upper bound

Proposition 6.1 bounds single-block manipulation at $\Theta(\alpha R_M/b)$ bias. A persistent adversary who sustains a logit shift δ over T consecutive blocks extracts on the spot side at a rate $\Theta(\alpha R_M)$ per block while paying LMSR round-trip cost $\Theta(b\delta^2)$ per block in each direction of restoration. The open question of Open Problem 8 is whether a calibrator that adjusts κ (or the Huber radius $\delta_{\max}$) downward on detection of persistent bias can neutralise the attack. We close this with a matching upper bound.

Definition 6.9 (Rate-limited response calibrator)**.** A *rate-limited response calibrator* is a measurable function $\kappa_t = \Gamma(\mathcal{F}_t^{\text{hist}})$ mapping the history $\mathcal{F}_t^{\text{hist}}$ of observed prediction-market deviations $\{|E_s - E^\star|\}_{s\le t}$ to a coupling exponent $\kappa_t \in [0, \kappa_{\max}]$, satisfying:

(Cal1) *Monotone attenuation:* κ_t is non-increasing in the moving average $\overline{|E - E^\star|}_{[t-\tau_{\text{cal}},t]}$.

(Cal2) *Response time* τ_{cal}*:* the calibrator's attenuation factor satisfies $\kappa_t/\kappa_0 \le \eta(s/\tau_{\text{cal}})$ after s blocks of sustained bias exceeding a detection threshold $\theta^\star$, where $\eta : \mathbb{R}_{\ge 0} \to [0,1]$ is non-increasing with $\eta(0) = 1$ and $\eta(s) \to 0$ as $s \to \infty$.

(Cal3) *Symmetric extension:* the calibrator treats positive and negative deviations identically.

A canonical instance is the EMA calibrator $\kappa_t = \kappa_0 \cdot \max(0, 1 - \lambda \cdot \text{EMA}_{\tau_{\text{cal}}}[|E_s - E^\star|]/\theta^\star)$ with smoothing $\lambda > 0$ and decay scale τ_{cal}.

Theorem 6.10 (Persistent-adversary extraction upper bound)**.** *Let the coupling be Huber-bounded logit with exponent κ_0, saturation $\delta_{\max}$, and LMSR depth b. Let the operator enforce a rate-limited response calibrator Γ (Definition 6.9) with response time τ_{cal} and attenuation profile η. Then for any risk-neutral persistent adversary sustaining a logit shift $|\delta_t| \le \delta_{\max}$ over a horizon of T blocks, the cumulative net extraction satisfies*

$$\mathbb{E}[\Pi_T^{\text{net}}] \;\le\; \kappa_0 R_M |\delta_{\max}| \cdot \int_0^T \eta(s/\tau_{\text{cal}})\, ds \;-\; bE^\star(1-E^\star)\delta_{\max}^2 \cdot T. \tag{16}$$

Consequently, whenever the calibrator satisfies the response condition

$$\int_0^\infty \eta(u)\, du \;\le\; \frac{bE^\star(1-E^\star)\delta_{\max}}{\kappa_0 R_M} \cdot \tau_{\text{cal}}, \tag{17}$$

the net extraction $\mathbb{E}[\Pi_T^{\text{net}}]$ is uniformly bounded above by a constant independent of T: the persistent adversary cannot achieve unbounded extraction.

For the EMA calibrator with $\eta(u) = \max(0, 1-u)$ (linear decay, $\int \eta = 1/2$), the response condition (17) reads

$$\tau_{\text{cal}} \;\ge\; \frac{\kappa_0 R_M}{2bE^\star(1-E^\star)\delta_{\max}},$$

which is satisfied by design-time choice of the calibrator.

Proof. Step 1 (per-block cost). At block t, the adversary's gross spot extraction is at most the saturation-capped amount $\kappa_t R_M|\delta_t| \le \kappa_t R_M \delta_{\max}$ (Lemma 5.4 applied at the effective exponent κ_t). The LMSR restoration cost is $bE^\star(1-E^\star)\delta_t^2 \ge 0$ per block to maintain the bias, strictly positive on every block where $\delta_t \ne 0$.

Step 2 (calibrator attenuation). Under the assumption that the adversary maintains $|\delta_t| = \delta_{\max}$ (the worst-case for extraction), the detection threshold $\theta^\star$ is crossed after a bounded transient, and the calibrator's attenuation factor $\kappa_t/\kappa_0 \le \eta((t-t_0)/\tau_{\text{cal}})$ applies for $t \ge t_0$ (where t_0 is bounded by τ_{cal}-scaled detection delay). Summing the per-block bound over $t = 0, \ldots, T$:

$$\mathbb{E}[\Pi_T^{\text{net}}] \le \sum_{t=0}^{T} \kappa_t R_M \delta_{\max} - bE^\star(1-E^\star)\delta_t^2.$$

Substituting the attenuation bound gives $\sum_t \kappa_t R_M \delta_{\max} \le \kappa_0 R_M \delta_{\max} \cdot \sum_t \eta((t-t_0)/\tau_{\text{cal}})$, which converges to $\kappa_0 R_M \delta_{\max} \cdot \tau_{\text{cal}} \cdot \int_0^\infty \eta(u)\,du$ as $T \to \infty$.

Step 3 (bounded-extraction criterion). Net extraction is bounded iff the gross gain integral $\kappa_0 R_M \delta_{\max} \cdot \tau_{\text{cal}} \int \eta$ is dominated by the cumulative restoration cost $bE^\star(1-E^\star)\delta_{\max}^2 \cdot T$ for all T. In the limit the restoration cost diverges linearly in T while the gross gain is bounded (by integrability of η), so the net is eventually negative. The crossover horizon is bounded by (17) rearranged. $\square$

Remark 6.11 (Tightness and relation to prior art)**.** The lower bound of Proposition 6.1 gives $\Omega(\alpha R_M/b)$ per-block extraction for an adversary who pays no restoration cost (single-block pre-position-and-release). The upper bound of Theorem 6.10 gives $O(\kappa_0 R_M \delta_{\max} \cdot \tau_{\text{cal}} \cdot \int \eta)$ cumulative over all time for a persistent adversary who pays restoration on every block. The two bounds close the matching-bounds pair across the per-block / persistent-horizon dichotomy:

$$\text{per-block: } \Theta\big(\alpha R_M/b\big), \qquad \text{persistent: } O\big(\kappa_0 R_M \delta_{\max} \tau_{\text{cal}}/1\big) \text{ provided (17).}$$

The response condition is a discrete analog of the Milionis et al. [16] LVR-bound construction in which fee rates are set to offset expected arbitrage loss; here the calibrator Γ acts as a controller on the coupling exponent rather than on the trading fee. The $O(1)$ cumulative-extraction bound exploits full observability of the deviation signal $|E - E^\star|$; a bandit-feedback variant would yield an $O(\sqrt{T})$ regret bound instead.

6.5 Governance joint-parameter manipulation

A distinct adversary class is the *governance adversary*: a coalition that controls a quorum of governance votes and perturbs the calibrator parameters $(\kappa_0, \delta_{\max}, w_i)$ in combinations chosen to maximise extraction. Single-parameter governance perturbation is bounded by Lemma 5.4 (larger $\kappa_0\delta_{\max}$ yields larger per-block bound, but the composition bound is tight at $e^{\kappa_0\delta_{\max}} - 1$ per block). The remaining question is whether *joint* perturbation - simultaneously increasing κ_0, relaxing $\delta_{\max}$, and widening the transition window w_i - admits profitable manipulation that no single-parameter perturbation does. We show that the cumulative effect factors through the scalar $\kappa_0\delta_{\max}$ with no additional joint-perturbation gain.

Proposition 6.12 (Governance joint-parameter upper bound)**.** *Let $(\kappa_0^{\text{new}}, \delta_{\max}^{\text{new}}, w^{\text{new}})$ denote the post-governance calibrator parameters, with the pre-governance baseline $(\kappa_0, \delta_{\max}, w)$ satisfying the admissibility condition of Definition 5.1. Let $\Delta_\kappa = \kappa_0^{\text{new}} - \kappa_0$, $\Delta_\delta = \delta_{\max}^{\text{new}} - \delta_{\max}$, $\Delta_w = w^{\text{new}} - w$. The post-governance single-block price-impact bound satisfies*

$$\left|\frac{\Delta P}{P}\right|^{\text{new}} \le \exp\big((\kappa_0 + \Delta_\kappa)(\delta_{\max} + \Delta_\delta)\big) - 1. \tag{18}$$

The bound depends on $(\Delta_\kappa, \Delta_\delta)$ only through their product-sum $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}$; it is independent of Δ_w. Consequently:

(i) No joint-perturbation advantage. *Any joint governance perturbation $(\Delta_\kappa, \Delta_\delta, \Delta_w)$ achieving post-perturbation impact B^{new} is equivalent to a single-parameter perturbation attaining the same $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}$ product.*

(ii) Attack cost under quorum-gated governance. *A governance adversary controlling quorum-weight Q_{gov} must expend cost at least $c_{\text{gov}} \cdot Q_{\text{gov}}$ per governance cycle to enact $(\kappa_0^{\text{new}}, \delta_{\max}^{\text{new}})$, where c_{gov} is the governance-token opportunity cost (lockup yield foregone). The adversary's net profit over T blocks is bounded by*

$$\mathbb{E}[\Pi_T^{\text{gov}}] \le TR_M\left(e^{\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}} - 1\right) - c_{\text{gov}}Q_{\text{gov}}.$$

The attack is unprofitable iff $c_{\text{gov}}Q_{\text{gov}} > TR_M(e^{\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}} - 1)$.

(iii) Frontier-parameter constraint. *Imposing the constraint $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}} \le C^\star$ at the governance-contract level (a hard cap enforced by smart-contract admissibility) reduces the attack profit to at most $TR_M(e^{C^\star} - 1)$ independent of governance coalition size.*

Proof. (i) Lemma 5.4's bound $|\Delta P/P| \le e^{\kappa_0\delta_{\max}} - 1$ uses only the clipper amplitude $\kappa \cdot |\phi|_\infty = \kappa\delta_{\max}$; the transition-window shape (parameterised by w) appears in the clipper's interior regime, while the saturation amplitude is independent of it. Substituting the post-governance $\kappa_0^{\text{new}}, \delta_{\max}^{\text{new}}$ gives (18) directly. The bound factors through $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}$ as a single scalar; a joint perturbation achieving fixed $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}$ achieves the same impact regardless of decomposition into $(\Delta_\kappa, \Delta_\delta)$ components.

(ii) The per-block extraction bound $R_M(e^{\kappa_0^{\text{new}}\delta_{\max}^{\text{new}}} - 1)$ from (i) times the horizon T gives the gross profit; subtracting the governance cost gives the net bound.

(iii) The hard cap $\kappa_0^{\text{new}}\delta_{\max}^{\text{new}} \le C^\star$ at the governance-contract level is a smart-contract invariant: any proposed parameter update violating it is reverted. Under this cap the per-block bound is $e^{C^\star} - 1$ regardless of governance-coalition size, which is the claim. □

Remark 6.13 (Joint-perturbation design recommendation)**.** Proposition 6.12 identifies the *frontier-parameter constraint* $\kappa_0\delta_{\max} \le C^\star$ as the load-bearing admissibility condition. The natural design choice is $C^\star = \ln(1 + \beta^\star)$ where $\beta^\star$ is the operator-chosen maximum per-block price-impact tolerance; with $\beta^\star = 0.05$ (5% cap) this gives $C^\star \approx 0.0488$. Any governance-proposed parameter bundle failing this constraint is rejected at the smart-contract admissibility check; the frontier cap is the structural defence against joint-parameter governance attacks, independent of coalition-weight or lockup-duration assumptions.

6.6 Within-class proposer priority: matching lower bound

The proposer constraints PC1-PC4 of Definition 2.3 bound a block proposer's advantage against other traders. Within a single compliance class (a set of transactions with identical priority-treatment under the ordering rule), the proposer retains one residual advantage: among the m transactions in the class, the proposer's transactions enjoy an $O(1/m)$ positional advantage. We prove this upper bound is matched by an explicit adversarial construction.

Proposition 6.14 (Within-class proposer priority: matching lower bound)**.** *Let a block contain m transactions in a single compliance class (same priority level under PC1-PC4), with identical independent values V_i uniformly random on $[0, V_{\max}]$, and let the proposer submit $m_P \le m$ transactions in the same class, where $m_P + m_H = m$ with m_H the number of honest-submitter transactions. Then:*

(i) Upper bound. *The proposer's expected positional-advantage value (expected first-position value conditional on a proposer transaction being in the first position, minus the uniform-random baseline) is bounded by* $V_{\max}/2 - V_{\max} \cdot m_H/(2m) = V_{\max} \cdot m_P/(2m)$.

(ii) Matching lower bound. *An adversary controlling proposer role for a single block and submitting* m_P *transactions constructed to saturate the positional advantage (namely, submitting the transactions with the* m_P *highest available values* V_i *identified from the public mempool) achieves expected positional advantage* $V_{\max} \cdot m_P/(2m) \cdot (1 - O(1/m))$, *matching the upper bound up to lower-order corrections.*

Proof. Upper bound. Under PC1-PC4 the compliance-class definition admits any permutation of same-class transactions; the proposer has discretion only over the ordering within the class. Under PC2 (ciphertext opacity) the proposer cannot inspect the plaintext values of honest-submitter transactions, only the ciphertext count m_H. By exchangeability of the uniformly-distributed values V_i under any $\mathcal{F}^{\text{pre}}$-measurable permutation rule, each of the m positions in the class has the same expected value $V_{\max}/2$. Assigning the proposer's m_P transactions to the first m_P positions gives proposer value $m_P \cdot V_{\max}/2$; the uniform-random-ordering baseline assigns the proposer an expected $m_P \cdot V_{\max}/2$ total. The positional *advantage*, defined as the difference between the proposer's realised total under its discretion-maximising rule and the uniform-random baseline, is bounded by the gap in the expected value of the first position relative to the uniform average. With m_P/m as the fraction of top positions the proposer occupies, the per-class gain is at most $V_{\max} \cdot m_P/(2m)$.

Matching lower bound. The adversary controls the proposer and submits m_P transactions of value $V_{\max}$ (or arbitrarily close); the honest-submitter transactions are encrypted (PC2) but exchangeable with expected value $V_{\max}/2$. Placing the m_P adversary transactions first and the m_H honest transactions thereafter yields proposer-captured value $m_P \cdot V_{\max}$ versus the uniform-random baseline $m_P \cdot V_{\max}/2$, a gain of $m_P V_{\max}/2$. Normalising by the class-size m gives advantage $V_{\max} \cdot m_P/(2m)$, matching the upper bound exactly. The matching relies only on exchangeability of the honest values (not on the adversary's ability to observe them), consistent with PC2 opacity. □

Remark 6.15 (Defence: random intra-class ordering)**.** The residual $O(1/m)$ advantage of Proposition 6.14 is eliminated by mandating random intra-class ordering: within each compliance class, the order of transactions is determined by a publicly-verifiable unbiasable seed (e.g., RANDAO-post-commit hash applied to transaction identifiers). Under such mandated ordering the proposer's expected advantage reduces to zero; the residual bound holds when ordering is at-proposer-discretion within the compliance class. The choice is a PC-extension (call it PC5: intra-class random ordering) which strengthens Theorem 2.5's guarantees at the cost of slightly increased consensus complexity.

7 The Trading Function

We now develop the AMM side of the coupling: the invariant surface that, together with a one-way coupling $\tilde{G}$, produces the coupled marginal price. This is standard CFMM theory specialised to the one-way-coupling setting.

7.1 The power-law invariant

We seek a trading function $\varphi(R_M, R_B; g) = k$ satisfying:

(TF1) **Feasibility:** φ defined and positive for $R_M, R_B > 0$ and $g \in [\underline{g}, \overline{g}]$.

(TF2) **Scale-invariant signal impact:** multiplicative change in g produces multiplicative change in P, reserve-independent.

(TF3) **Closed-form IL:** the impermanent loss admits a closed form.

(TF4) **Uniswap recovery:** at $g = 1$, φ reduces to $R_M \cdot R_B = k$.

Proposition 7.1 (Characterisation of the power-law invariant). *Among differentiable trading functions satisfying (TF1)-(TF4), the power-law invariant*

$$\varphi(R_M, R_B; g) = R_M \cdot R_B^g = k \tag{19}$$

is the unique solution up to reparametrisation.

Proof. TF2 forces $\partial P/\partial g = P/g$, so $P = g f(R_M/R_B)$. TF4 gives $f(x) = x$, hence $P = g R_M/R_B$, which integrates along the invariant to $R_M R_B^g = k$. TF1, TF3 verified directly. □

7.2 Marginal price

The invariant surface is $\{(R_M, R_B) : R_M R_B^g = k\}$; equivalently, $\ln R_M + g \ln R_B = \ln k$. Implicit differentiation gives

$$P = \frac{g \cdot R_M}{R_B}. \tag{20}$$

Remark 7.2 (Single invariant). There is one invariant, $R_M R_B^g = k$. No separate "accounting invariant" is needed: value conservation is a constraint on the swap operation, not a property of the state.

8 LP Loss Analysis

8.1 Impermanent loss for the generalised invariant

Proposition 8.1 (IL formula). *For the invariant $R_M R_B^g = k$, an LP who deposits at price P_0 and experiences a price move to $P_1 = rP_0$ incurs impermanent loss*

$$\mathrm{IL}(r; g) = \frac{(g+1) \cdot r^{g/(g+1)}}{gr + 1} - 1. \tag{21}$$

Proof. From $R_M R_B^g = k$ and $P = gR_M/R_B$: $R_B = (gk/P)^{1/(g+1)}$, $R_M = k^{1/(g+1)}(P/g)^{g/(g+1)}$. AMM portfolio value $V_{\mathrm{AMM}}(P) = (g+1)k^{1/(g+1)}(P/g)^{g/(g+1)}$, so $V_{\mathrm{AMM}}(P_1)/V_{\mathrm{AMM}}(P_0) = r^{g/(g+1)}$. Using value conservation at P_0 gives $V_{\mathrm{hold}}(P_1)/V_{\mathrm{AMM}}(P_0) = (1 + gr)/(1 + g)$. Dividing yields (21).

The inequality $\mathrm{IL}(r; g) \le 0$ is equivalent to $(g+1)r^{g/(g+1)} \le gr + 1$, which is the weighted AM-GM with weights $g/(g+1)$ and $1/(g+1)$. □

Remark 8.2 (Uniswap recovery and negativity). At $g = 1$: $\mathrm{IL}(r; 1) = 2\sqrt{r}/(1 + r) - 1$, the standard second-generation Uniswap formula. For any $g > 0$ and $r \neq 1$, $\mathrm{IL}(r; g) < 0$. Coupling leaves IL present; the LVR coefficient decreases for $g \neq 1$, reducing the rate of IL accumulation.

8.2 LVR reduction

Proposition 8.3 (LVR coefficient for the power-law invariant)**.** *The Milionis et al. [16] LVR coefficient for the invariant $R_M R_B^g = k$ is*

$$\ell(g) = \frac{g}{2(g+1)^2}. \tag{22}$$

This is maximised at $g = 1$ with $\ell(1) = 1/8$; for any $g \neq 1$, $\ell(g) < 1/8$.

Proof. The AMM portfolio value is $V(P) = (g+1)k^{1/(g+1)}(P/g)^{g/(g+1)}$ (Proposition 8.1 derivation). The Milionis et al. LVR rate is $-(\sigma^2/2)P^2V''(P)$. Direct computation gives $-P^2V''(P)/V(P) = g/(g+1)^2$, so the rate as a fraction of portfolio value is $\ell(g)\sigma^2$ with $\ell(g) = g/(2(g+1)^2)$.

First-order condition at $g = 1$: $\ell'(g) = (1-g)/(2(g+1)^3)$, which vanishes at $g = 1$; second derivative is negative there, confirming the maximum. □

8.3 Coupling loss

Proposition 8.4 (Coupling loss)**.** *A coupling change $g_0 \to g_1$ at any price P produces a portfolio value ratio*

$$1 + \mathrm{CL}(g_0, g_1) = \frac{g_1+1}{g_0+1} \cdot \left(\frac{g_0}{g_1}\right)^{g_1/(g_1+1)}, \tag{23}$$

which is independent of P, strictly less than 1 whenever $g_0 \neq g_1$, and in general non-symmetric in (g_0, g_1).

Proof. On the g_0-surface at price P, $V_0 = (g_0+1)k_0^{1/(g_0+1)}(P/g_0)^{g_0/(g_0+1)}$. After a coupling change the reserves are unchanged but $k_1 = R_M R_B^{g_1}$. Arbitrage restores the external price P on the g_1-surface. The exponent of P in V_1/V_0 is

$$\frac{g_0 - g_1}{(g_0+1)(g_1+1)} + \frac{g_1}{g_1+1} = \frac{g_0}{g_0+1},$$

which matches V_0's exponent, so V_1/V_0 is P-independent; collecting the remaining factors gives (23). Strict inequality from weighted AM-GM with weights $g_1/(g_1+1)$ and $1/(g_1+1)$. □

8.4 Path dependence

Proposition 8.5 (Path dependence under dynamic coupling)**.** *When g varies over time, the total IL depends on the full path of (P_t, g_t), with first-order dependence on $|\Delta g|$. Correctly-directed signals arriving early reduce IL; wrongly-directed signals arriving early increase it.*

Proof. Compare Path A (coupling change first, then price move) and Path B (price first, then coupling). For $g_1 = g_0 + \varepsilon$ with $|\varepsilon| \ll 1$:

$$\mathrm{IL}_A - \mathrm{IL}_B = \varepsilon\Big[\partial_g \mathrm{IL}(r; g_0) + (1 + \mathrm{IL}(r; g_0))\partial_{g_1}\rho|_{g_1=g_0}\Big] + O(\varepsilon^2),$$

where $\rho(g_0, g_1, r)$ is the Path A-specific correction factor. The coefficient is nonzero for $r \neq 1$, confirming first-order path dependence. Sign analysis follows from reserve rebalancing direction. □

8.5 Break-even condition

Proposition 8.6 (LP break-even under GBM-Poisson coupling dynamics). *Let the price follow $dP/P = \mu\, dt + \sigma\, dW$, let coupling changes arrive as a Poisson process with rate λ_c and magnitudes $\Delta g \sim \mathcal{N}(0, \sigma_g^2)$, let the signal be correct with probability $p \in (1/2, 1)$, and let f denote trading fees, V expected volume per unit time, L pool depth. The LP is rational iff*

$$f \cdot V > \Big[\underbrace{\ell(g)\sigma^2}_{LVR} + \underbrace{\frac{\lambda_c \sigma_g^2}{2(g+1)^2}}_{\text{coupling loss}} + \underbrace{\frac{(1-p)\lambda_c \sigma_g \sigma}{g+1}}_{\text{signal-error cost}}\Big] \cdot L. \tag{24}$$

Proof. The LVR rate is $\ell(g)\sigma^2 L$ (standard Milionis et al.). Coupling loss from (23) expanded to second order gives $\mathrm{CL}(g, g+\Delta g) \approx -(\Delta g)^2/(2g(g+1)^2)$, times arrival rate λ_c and variance σ_g^2. Signal error is the path-dependence term from Proposition 8.5 averaged over wrong-direction events. □

Corollary 8.7 (Signal-accuracy crossover). *Setting the coupling benefit equal to the signal-error cost,*

$$p^\star = 1 - \frac{[\ell(1) - \ell(g)]\sigma(g+1)}{\lambda_c \sigma_g}.$$

For representative parameters ($g = 1.5$, $\sigma = 0.8$, $\lambda_c \approx 1$/day, $\sigma_g \approx 0.5$), $p^\star \approx 0.98$: the prediction market must be correct at least 98% of the time for coupling to be net positive through the LP-loss channel alone. Information value (Section 9) dominates this channel by construction.

9 Information Value of Coupling

The LVR reduction from coupling is a modest second-order effect. The dominant economic effect is *information value* to traders: coupling provides same-block price correction in thin markets where arbitrage cannot be relied upon, reducing adverse-selection costs for traders who transact at the coupled price. We now formalise this and prove that information value structurally dominates LVR reduction.

9.1 Information value: definition and first-order expression

Definition 9.1 (Price-accuracy benefit). The price-accuracy benefit is the expected reduction in adverse-execution cost for a trader transacting at the coupled vs. uncoupled price:

$$\Pi_{\text{info}} = q \cdot \mathbb{E}[|P_{\text{coupled}} - V| - |P_{\text{uncoupled}} - V|].$$

Proposition 9.2 (Information value under four idealisations). *Under the idealisations: (A1) Gaussian signal noise $\varepsilon_A \sim \mathcal{N}(0, \sigma_A^2)$, $\varepsilon_P \sim \mathcal{N}(0, \sigma_P^2)$, independent; (A2) binary signal-correctness indicator ($p > 1/2$); (A3) linear price impact (valid for small $\kappa\sigma_E$); (A4) exogenous volume. Then the expected price-accuracy benefit per unit traded is*

$$\Pi_{\text{info}} = q \cdot (2p-1) \cdot \kappa \cdot \sigma_E \cdot \sqrt{\frac{2}{\pi}}, \tag{25}$$

where σ_E is the prediction-signal standard deviation in log space. Aggregating over volume V and comparing to LVR reduction $[\ell(1) - \ell(g)]\sigma^2 L$ yields

$$\frac{V(2p-1)\kappa\sigma_E\sqrt{2/\pi}}{[\ell(1) - \ell(g)]\sigma^2 L}. \tag{26}$$

Proof. Under Gaussian ε_A, $\mathbb{E}[|\varepsilon_A|] = \sigma_A\sqrt{2/\pi}$. Under (A2) and (A3), the coupling adjustment shifts the error by $\pm\kappa\sigma_E$ with signs weighted by p (correct) and $1-p$ (wrong), yielding net expected benefit per unit $(2p-1)\kappa\sigma_E\sqrt{2/\pi}$. Aggregation over volume V and comparison to LVR reduction give (26). □

Remark 9.3 (Scope of idealisations)**.** The four idealisations constrain the validity of (25). (A1) Gaussianity conflicts with the self-reference noted in Section 4.1: ε_A and ε_P are not independent because P_{AMM} already incorporates $G(E)$; the formula is thus an upper-bound on information value, not a self-consistent estimator. (A2) collapses a continuous prediction error into a binary outcome; a refined version would weight by signal magnitude. (A3) requires $\kappa\sigma_E \ll 1$; for $\kappa\sigma_E > 0.5$ nonlinear effects dominate. (A4) fails whenever coupling-induced prices change volume elasticity. Relaxing any of (A1)-(A4) invalidates the derivation of the $\sqrt{2/\pi}$ constant and of the linear-in-$\kappa\sigma_E$ scaling in (25); the ratio in Theorem 9.4 is therefore an order-of-magnitude statement under these idealisations, not a regime-independent structural conclusion.

9.2 Ratio of first-order and second-order effects

Theorem 9.4 (Information value over LVR reduction, near the constant-product maximiser)**.** *Under (A1)-(A4) of Proposition 9.2, in the Taylor-linearisation regime $g = 1 + \kappa\sigma_E$,*

$$\frac{\Pi_{\text{info}} \cdot V}{\Delta\text{LVR} \cdot L} = \frac{32(2p-1)\sqrt{2/\pi}}{\kappa\sigma_E\sigma^2} \cdot \frac{V}{L}, \tag{27}$$

where $\Delta\text{LVR} = [\ell(1) - \ell(1+\kappa\sigma_E)]\sigma^2 L$. The ratio is $\Theta(1/(\kappa\sigma_E))$, diverging as $\kappa\sigma_E \to 0$.

Proof. Per-unit information benefit: $(2p-1)\kappa\sigma_E\sqrt{2/\pi}$; aggregated over volume $(2p-1)\kappa\sigma_E\sqrt{2/\pi}V$. LVR reduction at $g = 1$: because $\ell'(1) = 0$ (the constant-product invariant is the LVR *maximum*), the Taylor expansion gives $\ell(1) - \ell(g) = -\ell''(1)(g-1)^2/2 + O((g-1)^3) = (g-1)^2/32 + O((g-1)^3)$. Substituting $g = 1 + \kappa\sigma_E$ and taking the ratio gives (27). □

Remark 9.5 (Interpretation and scope)**.** The divergence as $\kappa\sigma_E \to 0$ reflects that the comparison is taken at $g = 1$, which is the critical point of the LVR map ℓ. Any smooth map compared to its second-order Taylor term near a critical point yields an $O(1/\varepsilon)$ ratio. The result is the generic behaviour of a first-order quantity divided by a second-order quantity near a critical point. A stronger comparison would use a non-critical baseline (e.g., $g = g^\star \neq 1$), which we do not pursue here. The take-away: at leading order in the coupling strength, information value dominates LVR reduction near the constant-product regime; deviations from $g = 1$ bring the two into comparable order. The bound is tight under the (A1)-(A4) idealisations of Proposition 9.2, which Remark 9.3 documents.

10 Joint Prediction-AMM Clearing (Reference Construction)

Scope note. This section presents the clearing constructions used by the rest of the paper. The discrete call-auction program is the Peters-So-Ye convex parimutuel mechanism [7] on the Lange-Economides lineage [5, 6]; the continuous fixed-generator form of Section 10.4 uses the classical Kolmogorov-Nagumo mean structure [8, 9]. Neither ingredient is claimed as a discovery here. The load-bearing claims of this section are the outer fixed-point analysis for one-way coupling, the continuous-state equilibrium statement, and the bounded-grid discretisation / atomic-limit connection.

10.1 Log-barrier formulation

Let n be the number of orders, each specifying direction, limit price p_i, and maximum quantity q_i. Let $\sigma_i \in [0, q_i]$ denote fill fractions and $S(\sigma)$ the surplus function. The clearing program is

$$\max_{\sigma,P} \; S(\sigma) + \theta \sum_i \log \sigma_i \tag{28}$$

subject to margin and fill constraints, market balance, and the invariant $\ln R'_M + g \ln R'_B = \ln k$ with $g = G(E)$ determined by the prediction state. The invariant is affine in log-reserves, making the inner program convex in log-reserve coordinates.

Remark 10.1 (Log-reserve formulation)**.** The invariant constraint (19) is affine in $(\ln R'_M, \ln R'_B)$. The original multiplicative form $R'_M \cdot {R'_B}^g = k$ is not convex in (R'_M, R'_B) for general g. The inner optimization (at fixed g) is convex in log-reserve coordinates. The relationship between fill fractions σ_i and log-reserves involves the nonlinear map $\sigma \mapsto \ln(R + \Delta(\sigma))$; the practical solver operates in fill-fraction space using the nonlinear constraint as a manifold projection. Convergence of this projection is proved in Theorem 12.4 under the stated compact-away-from-zero, LICQ, and restoration hypotheses.

The log-barrier term $\theta \sum_i \log(\sigma_i)$ serves two purposes: it replaces quadratic regularization (eliminating a tuning parameter), and as $\theta \to 0$ it converges to the surplus-maximizing solution with maximum-entropy fill distribution among all optimal solutions.

10.2 The non-convex joint problem

The joint prediction-spot clearing problem is non-convex because $G(E)$ depends on the prediction market state, which is itself an output of the clearing process. For a fixed value of g, the spot clearing problem is convex in log-reserve space.

The joint prediction-spot clearing problem decomposes as follows: for each fixed value of $g \in [\underline{g}, \overline{g}]$, the inner spot-clearing program is convex in log-reserve space and admits a unique solution (Theorem 12.4); the outer problem of selecting g consistent with the prediction-market clearing reduces to bisection on a single scalar. The inner program uses self-concordant log-barrier methods [20, 19] with Newton-iteration count $O(\sqrt{n}\log(1/\varepsilon))$; the outer bisection takes $O(\log(1/\varepsilon))$ evaluations. The power-law regime is handled by Theorem 10.2 below; the logit regime by Theorem 10.5.

Theorem 10.2 (Verifiable clearing, power-law regime)**.** *In the power-law coupling regime $G(E) = E^\kappa$, assume the cleared spot-price response $P^\star(g)$ obtained after the inner reserve-clearing program is C^1 on the operating interval and satisfies*

$$\sup_g \left| \frac{d \log P^\star(g)}{dg} \right| \le L_P.$$

Assume also the posterior map $E^\star(P) = a + c \log P$ is valid on the same interval and

$$\sup_g \kappa |c| \left(E^\star(P^\star(g))\right)^{\kappa-1} L_P < 1.$$

Then the joint prediction-AMM clearing problem admits a unique solution computable in

$$O\left(n^{2.5}\log(n/\varepsilon)\log(1/\varepsilon)\right) \subset O\left(n^3 \log^2(1/\varepsilon)\right)$$

via bisection on g: the inner constrained program converges under Theorem 12.4, and the outer scalar root is unique under the full cleared-price response bound of Lemma 10.3.

Proof. Decompose into three steps.

Step 1 (prediction market clearing). For fixed E, LMSR clearing gives unique solution [2].

Step 2 (spot clearing at fixed g). For fixed $g = E^\kappa$, the spot program (28) is convex in log-reserve space; interior-point methods give unique solution in $O(n^{2.5}\log(n/\varepsilon))$.

Step 3 (fixed point via bisection). Define $\Psi : [\underline{g}, \overline{g}] \to [\underline{g}, \overline{g}]$ by: given g, solve spot clearing for equilibrium reserves (R'_M, R'_B), compute the cleared spot price $P^\star(g)$, solve LMSR for $E^\star(P^\star(g))$, set $\Psi(g) = (E^\star)^\kappa$. Let

$$F(g) := \Psi(g) - g.$$

Because Ψ is continuous and maps $[\underline{g}, \overline{g}]$ into itself, $F(\underline{g}) \geq 0$ and $F(\overline{g}) \leq 0$. By Lemma 10.3, Ψ is differentiable on the operating interval and satisfies $|\Psi'(g)| < 1$ under the stated contraction bound. Hence

$$F'(g) = \Psi'(g) - 1 < 0$$

throughout the interval, so F is strictly decreasing and has a unique zero. Bisection on that scalar root therefore converges in $O(\log(1/\varepsilon))$ evaluations.

Each bisection evaluation costs $O(n^{2.5}\log(n/\varepsilon))$; total cost is therefore

$$O(n^{2.5}\log(n/\varepsilon)\log(1/\varepsilon)) \subset O\Big(n^3\log^2(1/\varepsilon)\Big).$$

□

Lemma 10.3 (Full cleared-price response bound)**.** *Let $P^\star(g)$ denote the price produced after solving the inner reserve-clearing program at fixed g. Under the local-linearisation assumption $E^\star(P) = a + c\ln P$ and the response bound $\sup_g |d\log P^\star(g)/dg| \leq L_P$:*

(i) $\Psi'(g) = \kappa c (E^\star)^{\kappa-1}\, d\log P^\star(g)/dg.$

(ii) $|\Psi'(g)| \leq \kappa |c| (E^\star)^{\kappa-1} L_P.$

(iii) *Ψ is a contraction whenever the right-hand side is uniformly strictly less than 1 on the operating interval.*

Proof. The map is $\Psi(g) = (E^\star(P^\star(g)))^\kappa$. Differentiating the composite gives

$$\Psi'(g) = \kappa (E^\star)^{\kappa-1}\frac{dE^\star}{d\log P}\frac{d\log P^\star(g)}{dg} = \kappa c (E^\star)^{\kappa-1}\frac{d\log P^\star(g)}{dg}.$$

Taking absolute values and applying the response bound gives the stated contraction condition. The bound is deliberately on the full cleared-price map: it includes the dependence of the optimizer's reserves (R'_M, R'_B) on g. □

Remark 10.4 (Regime scope)**.** Theorem 10.2 treats the power-law regime. Theorem 10.5 treats the Huber-bounded logit regime by linearising the *logit* posterior map $u^\star(P) := \operatorname{logit} E^\star(P)$ rather than the probability map $E^\star(P)$ itself. What remains open is only the second-order acceleration problem at the Huber breakpoints (Open Problem 3), not first-order convergence.

10.3 Logit-regime clearing convergence

We resolve the first-order fixed-point part of Open Problem 3 by extending the contraction analysis of Lemma 10.3 to the Huber-bounded logit coupling $\tilde{G}_\ell(E) = \exp(\alpha \cdot h(\operatorname{logit} E))$, where h is the 3-piece sign-mirrored Huber clipper of Example 5.2.

Theorem 10.5 (Verifiable clearing, logit regime with Huber bounding)**.** *Let the coupling be $\tilde{G}_\ell(E) = \exp(\alpha \cdot h(\operatorname{logit} E))$ with $\alpha > 0$, Huber radius $\delta_{\max}$, and transition window w. Suppose the posterior logit is locally affine in $\ln P$ around the operating point,*

$$u^\star(P) := \operatorname{logit} E^\star(P) = a_\ell + c_\ell \ln P,$$

with constant slope c_ℓ. Suppose the full cleared-price response at fixed g satisfies $\sup_g |d \log P^\star(g)/dg| \leq L_P$. Then the joint prediction-AMM clearing problem admits a unique solution computable in

$$O\big(n^{2.5} \log(n/\varepsilon) \log(1/\varepsilon)\big) \subset O\Big(n^3 \log^2(1/\varepsilon)\Big)$$

via bisection on $g = \tilde{G}_\ell(E)$, provided

$$e^{\alpha\delta_{\max}} |\alpha c_\ell| \cdot \|h'\|_\infty L_P \;<\; 1, \tag{29}$$

where $\|h'\|_\infty = 1$ is the Lipschitz constant of the Huber clipper.

Proof. Step 1 (inner program unchanged). The inner log-barrier program (Theorem 12.4) depends on the scalar $g = \tilde{G}_\ell(E)$ only through the manifold constraint $\ln R'_M + g \ln R'_B = \ln k$. The convexity analysis of Theorem 12.4(i)-(iv) holds for any $g > 0$, in particular for $g \in [e^{-\alpha\delta_{\max}}, e^{+\alpha\delta_{\max}}]$ (the image of $\tilde{G}_\ell$ under Lemma 5.4). Unchanged from Theorem 10.2.

Step 2 (outer map: logit-regime derivative). Define the fixed-point map $\Psi_\ell : [e^{-\alpha\delta_{\max}}, e^{+\alpha\delta_{\max}}] \to [e^{-\alpha\delta_{\max}}, e^{+\alpha\delta_{\max}}]$ by: given g, solve spot clearing for (R'_M, R'_B), compute the cleared price $P^\star(g)$, solve LMSR for $E^\star(P^\star(g))$, set $\Psi_\ell(g) = \tilde{G}_\ell(E^\star)$. Differentiating through the logit state $u^\star = \operatorname{logit} E^\star$,

$$\Psi'_\ell(g) = \frac{d\tilde{G}_\ell}{du}(u^\star) \cdot c_\ell \cdot \frac{d \log P^\star(g)}{dg}.$$

$$\frac{d\tilde{G}_\ell}{du} = \alpha h'(u) \tilde{G}_\ell(u)$$

by the assumed affine logit map. Writing $u^\star = \operatorname{logit} E^\star$ and applying the chain rule through $u^\star(P^\star(g))$ gives

$$\Psi'_\ell(g) = \alpha c_\ell h'(u^\star) \tilde{G}_\ell(E^\star) \frac{d \log P^\star(g)}{dg}.$$

Step 3 (uniform contraction). Under the Huber clipper of Example 5.2, $h \in C^1$ (Lemma 5.3) and h' is piecewise-linear with $|h'| \leq 1$ globally.

Here the Huber clipper is load-bearing: in the saturation regime $|u^\star| \geq \delta_{\max} + w$, $h'(u^\star) = 0$, so $\Psi'_\ell(g) = 0$ exactly. In the interior and transition regimes, $|h'(u^\star)| \leq \|h'\|_\infty$ and $\tilde{G}_\ell(E^\star) \leq e^{\alpha\delta_{\max}}$, hence

$$|\Psi'_\ell(g)| \leq e^{\alpha\delta_{\max}} |\alpha c_\ell| \cdot \|h'\|_\infty L_P.$$

The contraction condition (29) is therefore sufficient for Banach's theorem to apply, giving a unique fixed point with geometric convergence.

Step 4 (piecewise-C^1 handling and scalar root). Let

$$F_\ell(g) := \Psi_\ell(g) - g.$$

Since Ψ_ℓ is continuous and maps $[e^{-\alpha\delta_{\max}}, e^{+\alpha\delta_{\max}}]$ into itself,

$$F_\ell(e^{-\alpha\delta_{\max}}) \geq 0, \qquad F_\ell(e^{+\alpha\delta_{\max}}) \leq 0.$$

At the Huber breakpoints logit $E^\star \in \{\pm(\delta_{\max} \pm w)\}$, h' is continuous (Lemma 5.3) and nondifferentiable. The bisection method does not require second-order regularity. On each smooth piece,

$$F'_\ell(g) = \Psi'_\ell(g) - 1 \le |\Psi'_\ell(g)| - 1 < 0$$

by (29). Thus F_ℓ is strictly decreasing across the whole interval and has a unique zero. Bisection on F_ℓ therefore converges in $O(\log(1/\varepsilon))$ evaluations regardless of the piecewise-C^1 structure.

Step 5 (complexity). Each bisection step invokes the inner solver at cost $O(n^{2.5}\log(n/\varepsilon))$ (Theorem 12.4(iii)); total cost is therefore

$$O\big(n^{2.5}\log(n/\varepsilon)\log(1/\varepsilon)\big) \subset O\Big(n^3\log^2(1/\varepsilon)\Big).$$

□

Remark 10.6 (Saturation-as-stabiliser)**.** In the logit regime the Huber saturation plays two roles: it bounds per-block price impact (Lemma 5.4), and it forces $\Psi'_\ell = 0$ on the saturation band so that contraction is automatic there. In the power-law regime (unclipped) convergence of the bisection-on-g map requires the explicit condition $\kappa < 1$ plus moderate reserve magnitudes (Lemma 10.3); the logit-with-Huber regime replaces this global condition with a local one, trivially satisfied in the saturation band where $h' = 0$.

10.4 Continuous Product Constant-Argument Mean on bounded outcome spaces

Let $\Omega = [\underline{\omega}, \overline{\omega}] \subset \mathbb{R}$ be a compact outcome space with Lebesgue measure $d\omega$. Fix $\lambda > 0$ and a scalar generator $\phi \in C^2((0,\infty))$ that is strictly convex. Because the total-mass constraint below fixes $\int_\Omega \mu(\omega)\,d\omega$, replacing ϕ by $\phi + ax + b$ changes the objective by an additive constant only. We therefore normalise the generator so that $\phi'(x) > 0$ on the relevant density range; this is exactly the Kolmogorov-Nagumo fixed-generator convention [8, 9].

Definition 10.7 (Continuous Product Constant-Argument Mean)**.** Fix $0 < m < M < \infty$ such that $m|\Omega| \le 1 \le M|\Omega|$, and let $u \in C(\Omega)$ denote the aggregate marginal order-value density induced by the call auction. The admissible set is

$$\mathcal{M}^{ac}_{m,M}(\Omega) := \left\{\mu \in \mathcal{P}(\Omega) : \mu \ll d\omega,\ \rho_\mu := \frac{d\mu}{d\omega} \in L^\infty(\Omega),\ m \le \rho_\mu(\omega) \le M \text{ a.e.}\right\}.$$

Equivalently, the density representatives form

$$\mathcal{A}_{m,M} := \left\{\rho \in L^2(\Omega) : m \le \rho(\omega) \le M \text{ a.e.}, \int_\Omega \rho(\omega)\,d\omega = 1\right\}.$$

The continuous CPCAM clearing problem is

$$\max_{\mu \in \mathcal{M}^{ac}_{m,M}(\Omega)} J(\mu) = \max_{\rho \in \mathcal{A}_{m,M}} \left\{\int_\Omega u(\omega)\rho(\omega)\,d\omega - \lambda\Phi[\rho\,d\omega]\right\}, \tag{30}$$

where

$$J(\mu) := \int_\Omega u(\omega)\,d\mu(\omega) - \lambda\Phi[\mu], \qquad \Phi[\mu] := \int_\Omega \phi(\rho_\mu(\omega))\,d\omega.$$

Because Ω has finite measure and $m \le \rho_\mu \le M$, both terms in $J(\mu)$ are finite on $\mathcal{M}^{ac}_{m,M}(\Omega)$. The fixed-generator condition means that the same scalar convex function ϕ is used at every state ω.

Lemma 10.8 (Potential functional: well posed, lower semicontinuous, and strictly convex). *Assume*

$$m_\phi := \inf_{x \in [m,M]} \phi''(x) > 0.$$

Then $\Phi[\mu]$ is finite on $\mathcal{M}^{ac}_{m,M}(\Omega)$, weakly lower semicontinuous under the density identification $\mu \leftrightarrow \rho_\mu \in \mathcal{A}_{m,M}$, and satisfies

$$\Phi[t\mu_1 + (1-t)\mu_2] \le t\Phi[\mu_1] + (1-t)\Phi[\mu_2] - \tfrac{1}{2} m_\phi t(1-t) \|\rho_{\mu_1} - \rho_{\mu_2}\|^2_{L^2(\Omega)}$$

for every $\mu_1, \mu_2 \in \mathcal{M}^{ac}_{m,M}(\Omega)$ and $t \in (0,1)$.

Proof. Finiteness is immediate from $\rho_\mu(\omega) \in [m, M]$ almost everywhere and compactness of Ω. For the convexity bound, set $\rho_t := t\rho_{\mu_1} + (1-t)\rho_{\mu_2}$. Strong convexity of ϕ on $[m, M]$ gives the pointwise inequality

$$\phi(\rho_t) \le t\phi(\rho_{\mu_1}) + (1-t)\phi(\rho_{\mu_2}) - \tfrac{1}{2} m_\phi t(1-t)(\rho_{\mu_1} - \rho_{\mu_2})^2.$$

Integrating yields the stated estimate, hence strict convexity whenever $\rho_{\mu_1} \neq \rho_{\mu_2}$ on a set of positive measure. Lower semicontinuity follows because Φ is a convex continuous integral functional on the bounded $L^2(\Omega)$ set $\mathcal{A}_{m,M}$. □

Theorem 10.9 (Continuous CPCAM: existence, uniqueness, and Arrow-Debreu density). *Under the assumptions of Lemma 10.8, the continuous clearing problem (30) has a unique maximizer $\mu^\star \in \mathcal{M}^{ac}_{m,M}(\Omega)$. Writing $\rho^\star := d\mu^\star/d\omega$, if the maximizer is interior, meaning $m < \rho^\star(\omega) < M$ for almost every $\omega \in \Omega$, then there exists a scalar $\eta \in \mathbb{R}$ such that*

$$u(\omega) - \eta = \lambda\phi'(\rho^\star(\omega)) \qquad \text{for almost every } \omega \in \Omega. \tag{31}$$

Consequently the positive density

$$\pi^\star(\omega) := \lambda\phi'(\rho^\star(\omega)) \tag{32}$$

is the Arrow-Debreu state-price density up to normalization. Writing

$$q^\star(\omega) := \frac{\pi^\star(\omega)}{\int_\Omega \pi^\star(s)\, ds},$$

every bounded payoff $f \in L^\infty(\Omega)$ has cleared price

$$\Pi(f) = \int_\Omega f(\omega) q^\star(\omega)\, d\omega.$$

If $q^\star$ is continuous, then the call-price surface

$$C(K) := \Pi\big((\omega - K)^+\big) = \int_K^{\overline{\omega}} (\omega - K) q^\star(\omega)\, d\omega \tag{33}$$

satisfies

$$\frac{d^2 C}{dK^2}(K) = q^\star(K) \qquad \text{for } K \in (\underline{\omega}, \overline{\omega}). \tag{34}$$

Proof. Identify $\mu \in \mathcal{M}^{ac}_{m,M}(\Omega)$ with its density $\rho_\mu \in \mathcal{A}_{m,M}$. The set $\mathcal{A}_{m,M}$ is non-empty, closed, convex, and bounded in $L^2(\Omega)$; because $L^2(\Omega)$ is reflexive and Ω has finite measure, it is weakly compact. The linear part $\rho \mapsto \int_\Omega u(\omega)\rho(\omega)\, d\omega$ is weakly continuous, while Lemma 10.8 makes Φ weakly lower semicontinuous. Hence J is weakly upper semicontinuous on a weakly compact set, so a maximizer exists.

For uniqueness, let $\mu_1, \mu_2 \in \mathcal{M}^{ac}_{m,M}(\Omega)$ with densities ρ_1, ρ_2, and let $t \in (0,1)$. Lemma 10.8 gives

$$\Phi[t\mu_1 + (1-t)\mu_2] \le t\Phi[\mu_1] + (1-t)\Phi[\mu_2] - \tfrac{1}{2} m_\phi t(1-t)\|\rho_1 - \rho_2\|^2_{L^2(\Omega)}.$$

Subtracting $\lambda\Phi$ from the linear revenue term therefore yields

$$J(t\mu_1 + (1-t)\mu_2) \ge tJ(\mu_1) + (1-t)J(\mu_2) + \tfrac{\lambda m_\phi}{2} t(1-t)\|\rho_1 - \rho_2\|^2_{L^2(\Omega)},$$

so J is strictly concave and the maximizer is unique.

If $\mu^\star$ is interior, density variations $\rho^\star + \varepsilon\nu$ with $\int_\Omega \nu = 0$ remain admissible for $|\varepsilon|$ small. Differentiating $J((\rho^\star + \varepsilon\nu)\, d\omega)$ at $\varepsilon = 0$ gives

$$\int_\Omega \big(u(\omega) - \lambda\phi'(\rho^\star(\omega))\big)\nu(\omega)\, d\omega = 0 \qquad \text{for every } \nu \in L^2(\Omega) \text{ with } \int_\Omega \nu = 0.$$

Therefore $u(\omega) - \lambda\phi'(\rho^\star(\omega))$ is almost everywhere constant, proving (31) for some $\eta \in \mathbb{R}$. The normalization $\phi' > 0$ on $[m, M]$ makes $\pi^\star$ positive. The pricing rule $\Pi(f) := \int_\Omega f q^\star$ is linear and positive, so $q^\star$ is the normalized Arrow-Debreu state-price density of the cleared economy.

For the call-price identity, differentiate (33) under the integral sign:

$$\frac{dC}{dK}(K) = -\int_K^{\overline{\omega}} q^\star(\omega)\, d\omega, \qquad \frac{d^2C}{dK^2}(K) = q^\star(K),$$

where continuity of $q^\star$ justifies the second differentiation. This is the Breeden-Litzenberger relation [10] applied to the state-price density produced by the clearing problem itself. □

Proposition 10.10 (Atomic discretization recovers the discrete CPCAM). *Let $\Pi_N = \{C_1^{(N)}, \dots, C_{k_N}^{(N)}\}$ be partitions of Ω with mesh size $h_N := \max_i |C_i^{(N)}| \to 0$, and choose a representative point $\xi_i^{(N)} \in C_i^{(N)}$ for each cell. Restrict (30) to piecewise-constant densities*

$$\rho_N(\omega) = \sum_{i=1}^{k_N} \mu_i^{(N)} \mathbb{1}_{C_i^{(N)}}(\omega), \qquad \sum_{i=1}^{k_N} |C_i^{(N)}|\mu_i^{(N)} = 1, \qquad m \le \mu_i^{(N)} \le M.$$

Writing

$$u_i^{(N)} := \frac{1}{|C_i^{(N)}|} \int_{C_i^{(N)}} u(\omega)\, d\omega,$$

the restricted problem is the finite-dimensional program

$$\max_{\mu^{(N)}} \sum_{i=1}^{k_N} |C_i^{(N)}|\, u_i^{(N)} \mu_i^{(N)} - \lambda \sum_{i=1}^{k_N} |C_i^{(N)}|\, \phi(\mu_i^{(N)}), \tag{35}$$

which is the discrete CPCAM on the atomic state space $\{\xi_i^{(N)}\}_{i=1}^{k_N}$ with reference masses $|C_i^{(N)}|$. Writing

$$w_i^{(N)} := |C_i^{(N)}|\mu_i^{(N)}, \qquad \nu_N := \sum_{i=1}^{k_N} w_i^{(N)} \delta_{\xi_i^{(N)}},$$

the discrete optimizer is therefore an atomic measure. Each N has a unique maximizer, and if u is Lipschitz on Ω then the lifted piecewise-constant maximizers $\rho_N^\star\, d\omega$ converge in $L^1(\Omega)$, hence narrowly, to the continuous maximizer $\mu^\star$. The atomic optimizers $\nu_N^\star$ converge narrowly to $\mu^\star$ as well. The discrete state-price vectors

$$\pi_i^{(N)} := \lambda\phi'(\mu_i^{(N)})$$

converge cellwise almost everywhere to the continuous density $\pi^\star$ of (32).

Proof sketch. Let J_N denote the discrete objective in (35). The finite-dimensional program is strictly concave by the same m_ϕ argument, so each partition has a unique maximizer. Because u is Lipschitz and ϕ is C^2 on $[m, M]$, Riemann-sum consistency gives a uniform approximation

$$\sup_{\mu \in \mathcal{A}_{m,M}^{(N)}} |J_N(\mu) - J(\mu)| = O(h_N),$$

where $\mathcal{A}_{m,M}^{(N)}$ is the piecewise-constant admissible class on Π_N. Strict concavity upgrades objective consistency to convergence of unique maximizers: if the lifted densities $\rho_N^\star$ failed to converge to $\rho^\star$ in L^1, one would obtain an L^2-separated subsequence whose objective gap is bounded below by strong concavity, contradicting $J_N \to J$ uniformly on $\mathcal{A}_{m,M}^{(N)}$. Continuity of ϕ' on $[m, M]$ then yields $\pi_i^{(N)} \to \pi^\star$ cellwise almost everywhere.

For the atomic measures, compare ν_N with the lifted piecewise-constant measure $\rho_N \, d\omega$. For every Lipschitz test function f,

$$\left| \int_\Omega f \, d\nu_N - \int_\Omega f(\omega) \rho_N(\omega) \, d\omega \right| \le \sum_i \int_{C_i^{(N)}} |f(\xi_i^{(N)}) - f(\omega)| \mu_i^{(N)} \, d\omega \le \mathrm{Lip}(f) \, h_N.$$

Hence $\nu_N^\star$ and $\rho_N^\star \, d\omega$ are indistinguishable in the narrow topology as $h_N \to 0$, so the atomic optimizers converge narrowly to $\mu^\star$. □

Theorem 10.11 (Continuous logit coupling and grid-refinement convergence)**.** *Let*

$$I := [e^{-\alpha\delta_{\max}}, e^{+\alpha\delta_{\max}}].$$

For each $g \in I$, let $u_g \in C(\Omega)$ be the order-value density in Definition 10.7; assume $(g, \omega) \mapsto u_g(\omega)$ is jointly continuous and uniformly Lipschitz in g. Let $\mu_g^\star$ be the unique maximizer of (30) with u replaced by u_g, fix a strictly positive $f \in C^1(\Omega)$, and define the scalar price statistic

$$P_f(g) := \int_\Omega f(\omega) \, d\mu_g^\star(\omega).$$

Assume the posterior logit is

$$\ell^\star(P) = a_\ell + c_\ell \ln P,$$

that P_f is C^1 on I, and that

$$L := e^{\alpha\delta_{\max}} |\alpha c_\ell| \|h'\|_\infty \sup_{g \in I} \left| \frac{P_f'(g)}{P_f(g)} \right| < 1. \tag{36}$$

Define the continuous outer map

$$\Psi_{\ell,\infty}(g) := \exp\big(\alpha h(\ell^\star(P_f(g)))\big).$$

Then $\Psi_{\ell,\infty}$ is a contraction on I and has a unique fixed point $g_\infty^\star$.

Now let Π_N be partitions as in Proposition 10.10, let $\bar{\mu}_{N,g}^\star := \rho_{N,g}^\star \, d\omega$ be the corresponding lifted piecewise-constant maximizers for u_g, and define

$$P_{f,N}(g) := \int_\Omega f(\omega) \, d\bar{\mu}_{N,g}^\star(\omega), \qquad \Psi_{\ell,N}(g) := \exp\big(\alpha h(\ell^\star(P_{f,N}(g)))\big).$$

Assume each $P_{f,N}$ is C^1 and obeys the same logarithmic-derivative bound as in (36). Then every $\Psi_{\ell,N}$ is a contraction on I with unique fixed point $g_N^\star$, and

$$\sup_{g \in I} |\Psi_{\ell,N}(g) - \Psi_{\ell,\infty}(g)| = O(h_N), \qquad |g_N^\star - g_\infty^\star| \le \frac{\sup_{g \in I} |\Psi_{\ell,N}(g) - \Psi_{\ell,\infty}(g)|}{1 - L}.$$

If $\nu_N^\star$ is the atomic optimizer associated with $g_N^\star$ in Proposition 10.10, then $\nu_N^\star \rightharpoonup \mu_{g_\infty^\star}^\star$.

Proof. Differentiating the continuous outer map gives

$$\Psi'_{\ell,\infty}(g) = \Psi_{\ell,\infty}(g)\, \alpha h'(\ell^\star(P_f(g))) \cdot c_\ell \frac{P'_f(g)}{P_f(g)}.$$

Because $\Psi_{\ell,\infty}(g) \in I$ and $|h'| \leq \|h'\|_\infty = 1$ by Lemma 5.3,

$$|\Psi'_{\ell,\infty}(g)| \leq e^{\alpha\delta_{\max}} |\alpha c_\ell| \|h'\|_\infty \left| \frac{P'_f(g)}{P_f(g)} \right| \leq L < 1.$$

Hence $\Psi_{\ell,\infty}$ is a contraction. The same derivative formula with P_f replaced by $P_{f,N}$ shows every discrete map $\Psi_{\ell,N}$ is a contraction under the same bound. The Huber clipper remains the same C^1 three-piece function as in Theorem 10.5; bounded-grid refinement changes only the argument passed into h, not the junction conditions or the saturation derivative $h' = 0$ on the capped region.

Apply Proposition 10.10 pointwise in g. Joint continuity of u_g and compactness of $I \times \Omega$ make the Riemann-sum consistency constant uniform in g, so

$$\sup_{g \in I} \|\rho^\star_{N,g} - \rho^\star_g\|_{L^1(\Omega)} = O(h_N).$$

Since f is bounded,

$$\sup_{g \in I} |P_{f,N}(g) - P_f(g)| = O(h_N).$$

The positivity of f and the density bounds keep P_f and $P_{f,N}$ inside a compact subset of $(0, \infty)$, on which $\ell^\star$ and $\exp(\alpha h(\cdot))$ are Lipschitz. Therefore

$$\sup_{g \in I} |\Psi_{\ell,N}(g) - \Psi_{\ell,\infty}(g)| = O(h_N).$$

For the fixed points,

$$|g^\star_N - g^\star_\infty| \leq |\Psi_{\ell,N}(g^\star_N) - \Psi_{\ell,\infty}(g^\star_N)| + |\Psi_{\ell,\infty}(g^\star_N) - \Psi_{\ell,\infty}(g^\star_\infty)| \leq \sup_{g \in I} |\Psi_{\ell,N}(g) - \Psi_{\ell,\infty}(g)| + L|g^\star_N - g^\star_\infty|,$$

which rearranges to the stated bound. Finally, Proposition 10.10 gives narrow convergence of the atomic optimizer to $\mu^\star_g$ at each fixed g, and the continuity of $g \mapsto \mu^\star_g$ implied by strict concavity upgrades this to $\nu^\star_N \rightharpoonup \mu^\star_{g^\star_\infty}$ once $g^\star_N \to g^\star_\infty$. □

Remark 10.12 (One-way coupling survives the continuous extension)**.** Section 3 is agnostic to the dimension of the spot state. If a prediction report $\hat{E}$ perturbs the cleared measure $\mu^\star_{\hat{E}}$ and the prediction payoff depends on any differentiable spot statistic $P_f(\hat{E}) := \int_\Omega f(\omega)\, d\mu^\star_{\hat{E}}(\omega)$, the same mixed-partial obstruction applies with P replaced by P_f. Passing from reserve vectors to densities therefore does not reopen the forbidden reverse mechanism channel: the continuous CPCAM remains admissible only as a one-way coupling.

Remark 10.13 (Bounded-grid discretisation and validator determinism)**.** Validator execution replaces the continuum by a fixed quadrature grid on the bounded interval Ω and evaluates the integrals in (30) with deterministic node/weight tables. If f has bounded second derivative on Ω, a conservative affine-rescaling of the composite-trapezoid envelope gives

$$\left| \int_{\underline{\omega}}^{\overline{\omega}} f(\omega)\, d\omega - Q_N(f) \right| \leq \frac{(\overline{\omega} - \underline{\omega})^3}{12N^2} \sup_{\omega \in \Omega} |f''(\omega)|. \tag{37}$$

The continuous model therefore carries two separate error scales: finite-precision arithmetic on the chosen grid and quadrature bias $O(N^{-2})$ from approximating the continuum. On any fixed grid, the normalized state-price vector returned by the solver is already the directly observable Arrow-Debreu density for that discretized economy; passage to the continuum is controlled separately by (37).

Remark 10.14 (View-stable leader agreement is a precondition for validator determinism)**.** The bounded-grid determinism statement of Remark 10.13 presupposes that every non-faulty validator clears the *same* trade-batch against the *same* pre-state in the same view. When leader selection is keyed to each validator's local quorum certificate rather than a view-stable snapshot, two validators entering the same view with divergent high-QC observations can compute different proposer indices, fork the cleared batch, and subsequently disagree on the post-state $\mu^\star$ even though the underlying numerical program is identical. The obstruction is consensus-layer: determinism of the numerical clearing does not imply determinism of which clearing is accepted as canonical.

Sufficient condition (view-stable leader seed). Let $\mathrm{seed}(\mathrm{view}) := H\big(\mathrm{qc}^\star_{\mathrm{view}}.\mathrm{block\_hash}, \mathrm{view}\big)$ where $\mathrm{qc}^\star_{\mathrm{view}}$ is the justifying QC frozen at view entry (i.e., extracted from the view-change justification, not from subsequently-arriving local high-QC updates). If every validator computes the proposer index from $\mathrm{seed}(\mathrm{view})$ and rejects any proposal whose advertised proposer does not recompute from the same frozen snapshot, then all non-faulty validators agree on the leader within a view and therefore on the unique trade batch that leader publishes. Combined with Remark 10.13, this yields deterministic agreement on $\mu^\star$ at every non-faulty validator after the clearing step.

Diagnostic belt. Requiring the proposer to attach the justifying QC to the proposal envelope (so receivers recompute and check the proposer index against the enclosed QC) complements the view-stable seed by converting any residual divergence into an immediate mismatch rather than a silent vote-collection livelock. The first mechanism rules out divergence by construction; the second makes residual divergence observable and actionable.

10.5 Novelty of the clearing construction

Remark 10.15 (Relation to Peters-So-Ye)**.** The log-barrier clearing program of this section is *not novel*: the log-barrier regularisation, the state-price interpretation, and the KKT-certificate verifiability all derive from Peters-So-Ye [7] on the Lange [5] and Lange-Economides [6] parimutuel framework. The continuous fixed-generator form likewise does not claim novelty in the generalized-mean structure itself; that is classical Kolmogorov-Nagumo mathematics [8, 9]. The contributions relative to that prior art are: (a) the bisection-on-g decomposition reducing the non-convex joint problem to a sequence of convex subproblems; (b) the monotonicity/contraction argument for the power-law and logit outer maps (Lemma 10.3, Theorem 10.5) under an explicit full cleared-price response bound, together with damped primal-dual inner-solver convergence (Theorem 12.4); (c) the bounded continuous-state extension with explicit Arrow-Debreu first-order conditions, atomic-limit recovery of the discrete CPCAM, continuous logit fixed-point convergence under grid refinement, and deterministic bounded-grid error control (Theorem 10.9, Proposition 10.10, Theorem 10.11, Remark 10.13); and (d) the integration with batch-clearing execution in the SPEEDEX tradition [18].

11 Design-Space Comparison

The coupling function G is one architecture for composing a prediction market and an AMM. We briefly compare against alternatives.

- **(A) Direct oracle feed.** Replace G with an external price oracle that reads a target $P^\star$. Simpler to implement. Introduces oracle latency, excludes assets without reliable feeds, and yields no self-consistent clearing between spot and prediction state. For assets with reliable feeds the oracle architecture dominates; G-coupling's advantage is specific to the event-contingent setting.
- **(B) Order-book-level coupling.** Each prediction trade triggers a spot order at $P = f(p)$. The mapping is deterministic and hence predictable, creating systematic front-running; it is incompatible with AMM invariants and offers no bounded-loss guarantee to LPs.
- **(C) Reserve adjustment.** Inject or withdraw reserves on event outcomes. Economically equivalent to G-coupling when k is adjusted to match; inferior on LP protection when it is not (dilution risk plus invariant violation).
- **(D) Fee-based coupling.** Accuracy modulates trading fees. Zero manipulation surface; no price correction, so the primary information-transmission value of coupling in thin markets is absent. Fee modulation is a complement to G-coupling, not a substitute.

Remark 11.1 (No dominant architecture)**.** G-coupling does not dominate alternatives on all dimensions. (A) is simpler with reliable feeds; (D) has zero manipulation surface; (B) gives finer discovery. G's distinct advantages: operation without external feeds, same-block price correction, bounded price impact (Lemma 5.4), and tractable clearing in the power-law regime (Theorem 10.2). Outside event-contingent settings, simpler architectures suffice.

12 Manifold-Projection Convergence Under Regularity

Remark 10.1 described the inner solver as performing a manifold projection from fill-fraction coordinates $\sigma \in \prod_i [0, q_i]$ to log-reserve coordinates $(\ln R'_M, \ln R'_B)$ via the nonlinear map $\sigma \mapsto \ln(R + \Delta(\sigma))$, and Theorem 10.2 was stated conditional on convergence of this projection. This section removes the conditionality: the log-barrier regularization $\theta \sum_i \log(\sigma_i)$ already present in program (28) induces *strong* concavity of the objective in σ on the open box $\prod_i (0, q_i)$, and the manifold projection inherits geometric convergence from Newton's method on a strongly convex function with Lipschitz gradient (standard result; see [19], §9.5). The present section spells out the argument explicitly against the PW-AMM's nonlinear invariant constraint.

12.1 Strong concavity of the log-barrier objective

Let $\Lambda(\sigma) := S(\sigma) + \theta \sum_{i=1}^n \log(\sigma_i)$ denote the objective of program (28), and let $D = \prod_{i=1}^n (0, q_i)$ denote its open domain.

Lemma 12.1 (Strong concavity of the log-barrier objective)**.** *For every $\theta > 0$ and every compact subset $K \subset D$ bounded away from the boundary (i.e. there exists $\varepsilon > 0$ with $\sigma_i \geq \varepsilon$ for all $\sigma \in K$ and all i), the objective Λ is m-strongly concave on K with modulus*

$$m = \theta / q_{\max}^2, \qquad q_{\max} := \max_i q_i.$$

Proof. The surplus function S is concave (it is the sum of order-wise linear revenues truncated at limit prices, hence a concave piecewise-linear function of σ; see [7], §3). The Hessian of the log-barrier term $\theta \sum_i \log(\sigma_i)$ is the diagonal matrix $-\theta\, \mathrm{diag}(1/\sigma_i^2)$, whose eigenvalues are all $-\theta/\sigma_i^2 \le -\theta/q_{\max}^2$ on K. Hence $\nabla^2 \Lambda \preceq \nabla^2(\theta \sum_i \log \sigma_i) \preceq -(\theta/q_{\max}^2)I$, which is precisely strong concavity with modulus $m = \theta/q_{\max}^2$. □

12.2 Lipschitz gradient of the invariant-constraint residual

The inner solver enforces the invariant constraint (19):

$$h(\sigma) \;:=\; \ln R'_M(\sigma) + g \cdot \ln R'_B(\sigma) - \ln k \;=\; 0,$$

where $R'_M(\sigma) = R_M + \Delta_M(\sigma)$ and $R'_B(\sigma) = R_B + \Delta_B(\sigma)$, with Δ_M and Δ_B linear functions of the fill fractions (sum of signed buy/sell quantities at order-specific prices).

Lemma 12.2 (Lipschitz gradient on bounded-away-from-boundary sets). *Fix $\underline{R} > 0$. On the set*

$$K_{\underline{R}} \;:=\; \{\sigma \in D : R'_M(\sigma) \ge \underline{R},\ R'_B(\sigma) \ge \underline{R}\},$$

the map h is twice continuously differentiable, and ∇h is L-Lipschitz with constant

$$L \;=\; \max\!\big(\|\nabla \Delta_M\|_2^2/\underline{R}^2,\ g \cdot \|\nabla \Delta_B\|_2^2/\underline{R}^2\big).$$

Proof. $\nabla h(\sigma) = \nabla \Delta_M / R'_M(\sigma) + g\, \nabla \Delta_B / R'_B(\sigma)$. Since Δ_M, Δ_B are linear, $\nabla \Delta_M, \nabla \Delta_B$ are constant vectors and the Hessian is

$$\nabla^2 h(\sigma) \;=\; -\nabla \Delta_M (\nabla \Delta_M)^\top / R'_M(\sigma)^2 \;-\; g\, \nabla \Delta_B (\nabla \Delta_B)^\top / R'_B(\sigma)^2.$$

On $K_{\underline{R}}$ we have $R'_M, R'_B \ge \underline{R}$, so $\|\nabla^2 h\|_2 \le L$ with L as stated. Lipschitz continuity of ∇h on the convex set $K_{\underline{R}}$ follows. □

Remark 12.3 (Boundedness away from reserves-zero). Every bounded trade σ keeps reserves positive by the non-depletion condition $\Delta_M(\sigma) > -R_M$ and $\Delta_B(\sigma) > -R_B$; the log-barrier in (28) already enforces $\sigma_i > 0$, and the PW-AMM's pool-emptying limit (standard: clearing stops when either reserve would hit $\underline{R} = R_M/2$) keeps the solver inside $K_{\underline{R}}$ for a fixed $\underline{R} > 0$.

12.3 Geometric convergence of the projected-Newton inner solver

The inner solver is the projected-Newton iteration applied to the KKT system of program (28) with the invariant constraint (19) treated via the primal-dual formulation (see [19], §10.2). Write the Lagrangian as $\mathcal{L}(\sigma, \mu) = \Lambda(\sigma) - \mu h(\sigma)$.

Theorem 12.4 (Damped primal-dual manifold-projection convergence). *Fix $\theta > 0$ and an initial point $\sigma^{(0)} \in K_{\underline{R}}$ satisfying $h(\sigma^{(0)}) = 0$ (achievable by the order-book warm-start: any feasible pre-trade reserve state lies on the invariant surface). Under Lemmas 12.1 and 12.2, LICQ $\nabla h(\sigma) \neq 0$ on the active manifold, and a compact reserve-away-from-zero domain, the damped primal-dual Newton step*

$$\begin{bmatrix} \nabla^2_{\sigma\sigma}\mathcal{L}(\sigma^{(k)}, \mu^{(k)}) & -\nabla h(\sigma^{(k)}) \\ \nabla h(\sigma^{(k)})^\top & 0 \end{bmatrix} \begin{bmatrix} \Delta\sigma^{(k)} \\ \Delta\mu^{(k)} \end{bmatrix} = - \begin{bmatrix} \nabla_\sigma \mathcal{L}(\sigma^{(k)}, \mu^{(k)}) \\ h(\sigma^{(k)}) \end{bmatrix},$$

with $\sigma^{(k+1)} = \sigma^{(k)} + \eta_k \Delta\sigma^{(k)}$, $\mu^{(k+1)} = \mu^{(k)} + \eta_k \Delta\mu^{(k)}$, Armijo back-tracking, and a restoration projection back to $h(\sigma) = 0$ whenever the post-step residual exceeds the quadratic Taylor bound, satisfies:

(i) ***Feasibility preservation.*** *For all k, $\sigma^{(k)} \in K_{\underline{R}}$ and $|h(\sigma^{(k)})| \le L\|\sigma^{(k)} - \sigma^{(k-1)}\|^2/2$ (post-Newton-step residual).*

(ii) ***Global damped convergence to the unique primal-dual optimum.*** *There exists a constant $\rho < 1$ (depending on m, L, θ and the LICQ margin) such that the KKT residual satisfies*

$$\|r_{\mathrm{KKT}}(\sigma^{(k)}, \mu^{(k)})\|_2 \;\le\; \rho^k \|r_{\mathrm{KKT}}(\sigma^{(0)}, \mu^{(0)})\|_2,$$

where σ^* is the unique maximizer of Λ on the invariant manifold. The contraction constant satisfies $\rho \le 1 - m/(2L) < 1$.

(iii) ***Quadratic terminal convergence.*** *Within a neighborhood $\|\sigma^{(k)} - \sigma^*\|_2 \le \delta$ for some $\delta > 0$, convergence is quadratic: $\|\sigma^{(k+1)} - \sigma^*\|_2 \le C\|\sigma^{(k)} - \sigma^*\|_2^2$ for an explicit constant $C = L/(2m)$.*

Proof. This is the textbook guarantee for Newton's method on a strongly concave objective with Lipschitz gradient, specialized to the equality-constrained primal-dual setting. We assemble the pieces:

Part (i). The Armijo back-tracking line search selects η_k such that $\Lambda(\sigma^{(k+1)}) \ge \Lambda(\sigma^{(k)}) + \alpha\eta_k \nabla\Lambda(\sigma^{(k)})^\top \Delta\sigma^{(k)}$ for $\alpha \in (0, 1/2)$; see [19], §9.5.1. The Armijo step always succeeds on a strongly concave objective (Lemma 12.1) because the quadratic lower bound $\Lambda(\sigma + \eta\Delta\sigma) \ge \Lambda(\sigma) + \eta\nabla\Lambda^\top\Delta\sigma - (L\eta^2/2)\|\Delta\sigma\|^2$ (Lemma 12.2) gives a sufficient-decrease η-interval of positive length. The feasibility residual bound $|h(\sigma^{(k)})| \le L\|\sigma^{(k)} - \sigma^{(k-1)}\|^2/2$ follows from Taylor's theorem applied to h around the previous iterate, using $h(\sigma^{(k-1)}) = 0$ (maintained as an invariant by the KKT-consistent primal-dual step) and the Lipschitz-gradient bound.

Part (ii). Under strong concavity (modulus m) and Lipschitz gradient (constant L), the standard Nesterov-style descent bound [20, Thm 2.1.14] gives

$$\Lambda(\sigma^*) - \Lambda(\sigma^{(k+1)}) \;\le\; (1 - m/L)\big(\Lambda(\sigma^*) - \Lambda(\sigma^{(k)})\big).$$

Combined with the quadratic lower bound $\Lambda(\sigma^*) - \Lambda(\sigma^{(k)}) \ge (m/2)\|\sigma^{(k)} - \sigma^*\|^2$, we get $\|\sigma^{(k+1)} - \sigma^*\|^2 \le (1 - m/L)\|\sigma^{(k)} - \sigma^*\|^2$; the contraction constant $\rho = \sqrt{1 - m/L} \le 1 - m/(2L) < 1$.

Part (iii). Once $\sigma^{(k)}$ lies in the region of convergence $\|\sigma^{(k)} - \sigma^*\| < m/L$, the Armijo step size is unit (full Newton step), and standard Newton theory [19, Thm 9.5.3] gives quadratic convergence: $\|\sigma^{(k+1)} - \sigma^*\|_2 \le (L/(2m))\|\sigma^{(k)} - \sigma^*\|_2^2$. □

Corollary 12.5 (Clearing is conditionally solvable under stated regularity)**.** *Theorem 10.2 holds under the compact-away-from-zero, LICQ, restoration, and full cleared-price response hypotheses stated above: the inner constrained program is solvable to accuracy ε by damped primal-dual Newton iterations, each costing $O(n^3)$ for the KKT-system factorization. The remaining open work is to mechanize the response bound L_P for every production potential family.*

Remark 12.6 (Role of the barrier parameter θ)**.** The contraction constant is bounded away from 1 for every fixed $\theta > 0$ and positive LICQ margin. As $\theta \to 0$ (to approach the surplus-maximizing solution exactly), the modulus $m \to 0$ and the contraction degrades; the standard remedy is the interior-point schedule $\theta^{(k)} = \theta^{(0)}/(1 + \beta k)^2$, which yields the usual $O(\sqrt{n}\log(1/\varepsilon))$ barrier complexity under the hypotheses above.

13 Open Problems

The cross-venue manipulation bound (Proposition 6.1), signal-accuracy crossover (Corollary 8.7), and conditional manifold-projection convergence (Theorem 12.4) are handled in the body; the governance-coupling upper bound is in Proposition 6.12. The following problems remain open.

Open Problem 1 (Uniform cleared-price response)**.** Mechanize a uniform bound $\sup_g |d \log P^\star(g)/dg| \leq L_P$ for every admitted production potential family and every bounded order book satisfying the reserve-away-from-zero condition. The inner KKT solver is covered under regularity assumptions; the outer fixed-point proof depends on this full optimizer-response bound.

Open Problem 2 (Self-consistent coupling parameter)**.** Find the fixed point $\kappa^\star$ of $\kappa^\star = \sigma_A^2(\kappa^\star)/(\sigma_A^2(\kappa^\star) + \sigma_P^2)$, where $\sigma_A^2(\kappa)$ is the AMM price variance conditional on coupling strength. The difficulty: $\sigma_A^2(\kappa)$ depends on arbitrageur and LP equilibrium behaviour. We conjecture $\kappa^\star < \kappa_{\text{naive}}$, but existence and uniqueness are unproven.

Open Problem 3 (Second-order regularity of logit-regime clearing)**.** Theorem 10.5 closes the first-order contraction analysis for the Huber-bounded logit coupling. Remaining is the second-order structure: Newton-type accelerated convergence of the bisection-on-g procedure, which requires a smoothing of the C^1-but-not-C^2 Huber breakpoints (for example, a C^∞ bump-function replacement with quadratic-in-distance approximation error). This is an accelerated-convergence refinement, not a fundamental gap.

Open Problem 4 (Multi-block settlement for structured products)**.** Extend single-block atomic settlement to structured products whose constituent events resolve in different blocks. Cross-block atomicity is not available: a product depending on events in blocks B_1 and B_2 cannot be settled atomically because outcomes are revealed sequentially.

Open Problem 5 (Feedback-loop stability)**.** Prove that the inter-block feedback $E \xrightarrow{G} P \xrightarrow{\text{arb}} E' \xrightarrow{G} P' \to \cdots$ converges to a stationary distribution. Theorem 10.2 gives bounded per-block clearing and Lemma 5.4 bounded per-block impact, but boundedness does not imply convergence. A proof requires characterising the spectral radius of the linearised feedback operator. The impossibility theorem (Theorem 3.2) forbids direct mechanism feedback within a block; it does not address the cross-block informed-trading channel addressed here.

Open Problem 6 (LP equilibrium existence)**.** Prove existence of an LP equilibrium: a liquidity level $\mathcal{L}^\star$ and fee rate $f^\star$ such that LP participation is individually rational and no LP can profitably deviate. IL and coupling loss are endogenous; this is a fixed-point-with-endogenous-risk problem.

Open Problem 7 (Logit-regime manipulation bound under saturation)**.** The manipulation bounds of Section 6 are stated in the small-shift linearisation regime $|\delta^\star| \ll 1$, where the coupling exponent is unsaturated. When the adversary drives the prediction market into the Huber-saturation regime $|\operatorname{logit} E^\star + \delta| \geq \delta_{\max} + w$ (the coupling becomes insensitive to further shifts), the linear-in-δ spot profit of Proposition 6.1 breaks down: the coupling caps at $e^{\pm\kappa\delta_{\max}}$ and further shifts yield zero marginal spot profit while continuing to cost $bE^\star(1-E^\star)\delta^2$ in LMSR round-trip. The adversary's optimal strategy in the saturated regime is to park the shift at or below the saturation edge. Characterise the adversary's optimal δ^{sat}, the exact profit ceiling under saturation, and the corresponding tight LMSR depth requirement. The answer is a bounded, but non-closed-form, optimisation problem.

Proposition 13.1 (Persistent-manipulation upper bound under exponential calibrator)**.** *Suppose a calibrator reduces α exponentially at rate $\mu > 0$ once detecting bias $|E_{\text{eq}} - E^\star| \geq \varepsilon_{\text{det}}$, with detection latency τ_{cal} blocks. Under the setup of Proposition 6.1, a persistent adversary sustaining $\delta_0 \geq \varepsilon_{\text{det}}/(E^\star(1-E^\star))$ over $T \geq \tau_{\text{cal}}$ blocks extracts at most*

$$\text{Ext}_T \leq \tau_{\text{cal}}\, \alpha_0 R_M \delta_0 + \alpha_0 R_M \delta_0/\mu, \tag{38}$$

uniformly in T.

Proof. Per-block extraction during detection latency: $\alpha_0 R_M \delta_0$, summing to $\tau_{\text{cal}}\alpha_0 R_M \delta_0$. Post-detection, $\alpha(t) = \alpha_0 e^{-\mu(t-t_0-\tau_{\text{cal}})}$; $\sum_{s=0}^{\infty} e^{-\mu s} \leq 1/\mu$ gives the tail bound. $\square$

Remark 13.2 (Below-threshold sub-threshold extraction)**.** If the adversary sustains $\delta_0 < \varepsilon_{\text{det}}/(E^\star(1-E^\star))$ the calibrator does not respond; cumulative extraction grows linearly as $T\alpha R_M \delta_0$. Per-block profit remains bounded above by the optimal-shift profit $\alpha^2 R_M^2/(4bE^\star(1-E^\star))$. Net-positive after amortising the bribe cost depends on calibrator-threshold design.

Open Problem 8 (Persistent-adversary manipulation in the sub-threshold regime)**.** Proposition 13.1 bounds the extraction above the calibrator's detection threshold. Below the threshold, the adversary sustains a shift the calibrator does not respond to and cumulative extraction grows linearly with the horizon. Characterise the optimal adversary strategy and the tight cumulative extraction bound as a function of $(\varepsilon_{\text{det}}, \tau_{\text{cal}}, \mu)$, including the regime where δ_0 lies strictly between zero and $\varepsilon_{\text{det}}/(E^\star(1-E^\star))$. A rigorous solution requires an adaptive-control analysis under adversarial perturbation that the present paper does not carry out.

A Adversarial-completeness inventory

A single-table summary of the adversarial claims in this paper, classifying each by adversary model, attack cost, defence, and completeness status. A claim is *complete* if it carries matching upper and lower bounds; *upper-only* if only the attacker's success probability is bounded above; *lower-only* if only a constructive attack is given; *under-specified* if adversary model or attack cost is absent.

Claim	Adversary model	Attack cost	Defence	Status
Thm 2.5	static, $< t$ shares, polynomial-time	$tS_{\min}$ threshold shares	PC1-PC4 + IND-CPA + (PC2-pad/timing/meta/crypto)	Complete
Thm 3.2	trader holding spot + prediction, C^2 utility	trivial (any non-trivial deviation)	Admissibility requirement (one-way necessity)	Complete (impossibility)
Lem 5.4	arbitrary E-realisation	n/a (worst-case adversarial E)	C^1-clipped coupling with saturation $\delta_{\max}$	Complete (tight at boundary, Thm 5.6)
Thm 5.7	arbitrary E-realisation	n/a	None (unclipped)	Complete (impossibility)
Prop 6.1	risk-neutral, single-block, joint (δ, Q)	$b\delta^2 + Q^2/(2R_B)$ round-trip	LMSR depth $b > \alpha Q^\star/(2\varepsilon)$	Complete (FOC interior max)
Prop 6.7	rational, controls event outcome	Bribe B + fees + LMSR loss	Huber saturation $\kappa\delta_{\max}$	Complete (quadratic cap)
Thm 6.10	risk-neutral, persistent over T blocks	Per-block LMSR cost $b\delta^2$ sustained	Rate-limited calibrator Γ, response time τ_{cal}	Above-threshold only ($O(1)$ cumulative under (17)); sub-threshold regime is OP 8
Prop 6.12	governance coalition, quorum-weight Q_{gov}	Lockup + opportunity cost	Frontier constraint $\kappa_0\delta_{\max} \le C^\star$	Complete (scalar-factorisable)
Prop 6.14	block proposer, single-class within block	Proposer role (PoS stake)	PC1-PC4 + intra-class random ordering (PC5)	Complete ($\Theta(V_{\max} m_P/m)$)

Table 1: Adversarial-claim completeness inventory. Matching upper and lower bounds for every primary adversary class. Defence mechanisms are cited in the indicated theorems and definitions; see Definitions 2.3, 5.1, 6.9.

Residual openings (documented as formal open problems). Open Problem 8 flags the adaptive-observer regime (adversary shapes $|\delta_t|$ based on calibrator feedback). Open Problem 3 flags second-order regularity for accelerated Newton-type inner solvers. All other adversarial claims in Table 1 carry matching upper and lower bounds.

References

[1] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107-119, 2003.

[2] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3-15, 2007.

[3] G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1), 2021.

[4] M. Egorov. Automatic market-making with dynamic peg. Technical report, Curve Finance, 2021.

[5] J. Lange. Financial products having a demand-based, adjustable return, and trading exchange therefor. US Patent 6,321,212, November 2001.

[6] J. Lange and N. Economides. A parimutuel market microstructure for contingent claims. *European Financial Management*, 11(1):25-49, 2005.

[7] M. Peters, A. M.-C. So, and Y. Ye. A convex parimutuel formulation for contingent claim markets. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC'06)*, 2006.

[8] A. N. Kolmogorov. Sur la notion de la moyenne. *Atti della Reale Accademia Nazionale dei Lincei. Rendiconti della Classe di Scienze Fisiche, Matematiche e Naturali*, 12:388-391, 1930.

[9] M. Nagumo. Über eine Klasse der Mittelwerte. *Japanese Journal of Mathematics*, 7:71-79, 1930.

[10] D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621-651, 1978.

[11] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315-1335, 1985.

[12] F. D. Foster and S. Viswanathan. Strategic trading with higher-order forecast beliefs. *Journal of Finance*, 51(4):1437-1478, 1996.

[13] L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71-100, 1985.

[14] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547-1621, 2015.

[15] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE Symposium on Security and Privacy*, pages 910-927, 2020.

[16] J. Milionis, C. Moallemi, T. Roughgarden, and A. L. Zhang. Automated market making and loss-versus-rebalancing. Working paper, 2022.

[17] S. Dalen. Toward Black-Scholes for prediction markets. arXiv:2510.15205, 2025.

[18] G. Ramseyer, A. Goel, and D. Mazières. SPEEDEX: A scalable, parallelizable, and economically efficient decentralized exchange. In *USENIX NSDI*, 2023.

[19] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[20] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization 87, Springer, 2004.

[21] K. Baron and J. Lange. *Parimutuel Applications in Finance: New Markets for New Risks*. Palgrave Macmillan, 2007.

[22] A. Othman, D. M. Pennock, D. M. Reeves, and T. Sandholm. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation*, 1(3):14:1-14:25, 2013.

[23] Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *UAI*, pages 49-56, 2007.

[24] J. Abernethy, Y. Chen, and J. Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12:1-12:39, 2013.

[25] L. Heimbach and R. Wattenhofer. Risks and returns of Uniswap V3 liquidity providers. In *ACM Advances in Financial Technologies (AFT)*, 2022.

[26] G. Angeris and T. Chitra. Improved price oracles: Constant function market makers. In *ACM Advances in Financial Technologies (AFT'20)*, pages 80-91, 2020.

[27] G. Angeris, A. Evans, T. Chitra, and S. Boyd. Optimal routing for constant function market makers. In *ACM Economics and Computation (EC'22)*, 2022.

[28] H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson. Uniswap v3 core. Technical report, Uniswap Labs, 2021.

[29] T. Roughgarden. Transaction fee mechanism design. In *ACM Economics and Computation (EC'21)*, pages 792-812, 2021.

[30] A. Park. The conceptual flaws of constant product automated market making. Working paper, University of Toronto, 2021.

[31] A. H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, 1970.

[32] B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall, 1979.

[33] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58-73, 1981.

[34] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.