Synthetic Liquidity for Long-Tail Claims

Factor Hedging, Residual Breadth, Flow Toxicity, and Quote Feasibility

Grant Stenger

Working paper

June 2026

Abstract

Thin long-tail claim markets need not be supported by one natural liquidity pool per claim. This paper gives a formal mechanical account of when such claims can be quoted. In a one-period dealer model, claim payoffs decompose as

$$X_i = \alpha_i + \beta_i^\top F + \varepsilon_i,$$

so inventory q has hedgeable factor exposure $B^\top q$ and hedging-residual variance $q^\top \Sigma_\varepsilon q$. The first theorem reduces high-dimensional liquidity provision to a lower-dimensional factor hedge plus a residual covariance book. The second theorem shows that claim count enters liquidity cost only through hedge cost, residual covariance, inventory concentration, supported inventory states, and capital constraints. The quote-feasibility theorem then gives a sufficient local break-even condition: spread revenue must cover factor hedge cost, incremental residual capital, adverse-selection loss, model and oracle risk, operating cost, and funding or collateral charges. These results are deliberately not a full equilibrium theory of dealer competition or dynamic strategic quoting; they identify the balance-sheet inequality any richer model must respect.

The paper distinguishes three liquidity regimes. Factor-hedged dealer liquidity uses deep hedge markets to absorb common risk. Bookmaker or insurer liquidity uses broad non-toxic residual flow and effective breadth to diversify independent claims. Subsidized information-market liquidity uses an outside budget, such as a scoring-rule loss bound, to buy price discovery when neither factor hedging nor residual-flow diversification is available. The distinctive empirical implication is residual price discovery: factor hedging commoditizes systematic exposure but leaves ε_i as the locus of specialist alpha, adverse selection, dealer spread, and factor-adjusted learning. The central conclusion is limited but important: liquidity fragmentation is not mechanically proportional to claim count. Liquidity scales with hedgeable factor dimension, residual effective breadth, flow toxicity, inventory concentration, model error, operating cost, collateral, and capital constraints. Synthetic liquidity is normal-time liquidity, not infinite liquidity.

Keywords: liquidity; market making; market microstructure; factor models; long-tail claims; derivatives; adverse selection; incomplete markets; prediction markets; scoring rules.
JEL codes: D47, D53, G12, G13, G14, G18.

Contents

1 Introduction

A large research program on market creation faces an immediate objection. Even if many residual payoff directions are valuable enough to justify claim creation, how can the resulting claims trade? Attention is finite. Dealer balance sheets are finite. Many candidate claims have little natural two-sided flow. If every small-business revenue share, creator-royalty claim, local weather trigger, product-line derivative, or event contract requires its own independent crowd, liquidity appears to fragment before the market frontier can move very far.

This paper gives the mechanical answer to that objection. The answer is not that every long-tail claim will be liquid. It is that long-tail claims do not require one natural liquidity pool per claim. There are at least three liquidity mechanisms. In the first, many claims share common exposure to a lower-dimensional set of hedgeable factors, so a market maker quotes individual claims by hedging their common component in deep factor markets and charging for the residual component that remains. In the second, a bookmaker, insurer, or dealer warehouses a broad book of independent or weakly correlated residuals because the flow is mostly non-toxic; here independence is an asset, not a problem. In the third, a sponsor or protocol subsidizes liquidity for information production, as in scoring-rule markets. The binding objects are factor hedgeability, residual covariance, inventory concentration, flow toxicity, adverse selection, model and oracle error, operating cost, collateral, and capital. Claim count matters only through its effect on these objects.

The core decomposition is simple. Let claim i have payoff

$$X_i = \alpha_i + \beta_i^\top F + \varepsilon_i,$$

where $F \in \mathbb{R}^d$ is a vector of hedgeable factor payoffs or returns, $\beta_i \in \mathbb{R}^d$ is the claim's factor loading, and ε_i is the residual. For a dealer with inventory $q \in \mathbb{R}^N$,

$$q^\top X = q^\top \alpha + (B^\top q)^\top F + q^\top \varepsilon.$$

The factor exposure is the d-dimensional vector $B^\top q$. The residual risk is summarized by the quadratic form

$$\mathrm{Var}(q^\top \varepsilon) = q^\top \Sigma_\varepsilon q.$$

If the factor exposure can be hedged in deep markets, the dealer's liquidity problem is no longer N unrelated market-making problems. It is a factor-hedge problem plus a residual-risk-capacity problem.

This is familiar in listed options. Many strikes and maturities have sparse natural flow, but dealers quote them because much of the risk can be hedged using underlying shares, index futures, rates, volatility books, and related options. The same architecture can extend to long-tail claims. A claim on one restaurant's revenue may load on neighborhood spending, food-input prices, labor costs, reservation traffic, and a basket of comparable restaurants. A creator-royalty claim may load on platform-wide advertising demand, genre cohorts, audience-decay curves, and comparable catalogs. A parametric crop claim may load on traded commodity prices, weather indexes, freight, and regional basis factors. The dealer does not need each claim to support a deep order book. It needs a hedge basis, a residual covariance model, and capital to warehouse what remains.

The same covariance mathematics also explains a different class of liquid markets: sports books, casinos, some insurance pools, and other bookmaker-style venues. In those settings the claims may be largely independent of traded hedge factors. Liquidity comes from broad, mostly uninformed, small-stake flow and from residual diversification across many independent events. This is the opposite comparative static from factor hedging: a factor-hedged dealer wants common components it can offload, while a bookmaker wants independent non-toxic risks it can diversify. Both mechanisms fail when flow becomes informed, one-sided, correlated, or too concentrated.

The paper is deliberately narrower than a theory of market creation. The first paper in this program studies which residual payoff subspaces are valuable enough to become markets. This paper assumes claims exist or can be created and asks how they can be quoted. The bridge is projection language, but the terminology matters. Paper 1's valuation-residual surplus is the dispersion-weighted value of payoff directions that existing markets fail to span. Paper 2's hedging residual risk is the dealer risk that factor hedges fail to eliminate. Both objects are residuals after projection onto an existing basis, but they answer different questions.

The paper is also narrower than a full market-microstructure equilibrium. The quote-feasibility theorem is a one-period break-even capacity condition for a dealer with a given book, hedge basis, cost functions, and risk constraints. It does not solve for equilibrium spreads under dealer competition, strategic order submission, endogenous information acquisition, latency races, or dynamic inventory control. Those objects matter for actual markets. The purpose here is to isolate the mechanical capacity condition that any such richer model must respect.

Reader's map: formal spine

The paper is organized around a short formal spine.

- **R1. Factor-hedge decomposition.** Theorem 4.1 identifies the exact object a dealer hedges, $B^\top q$, and the exact residual variance it cannot hedge, $q^\top \Sigma_\varepsilon q$. Proposition 4.3 shows that beta error reappears as residual risk.
- **R2. Residual breadth and liquidity dimension.** Proposition 5.5 and theorem 5.7 show that diversification expands capacity through effective residual breadth and supported inventory states, not through raw claim labels. Proposition 5.9 gives the residual-diversification mechanism for bookmaker and insurer liquidity.
- **R3. Quote feasibility.** Theorem 6.2 states the local break-even inequality. A quote is mechanically feasible when spread revenue covers factor hedge cost, incremental residual capital, adverse-selection loss, model and oracle risk, operating cost, and funding or collateral charges. Proposition 6.7 turns the inequality into a size-at-spread capacity bound.
- **R4. Residual information.** Theorem 8.1, proposition 8.2, and corollary 8.4 show that factor hedging localizes claim-specific information in ε_i. The same residual wedge is specialist alpha, dealer adverse-selection loss, informed-flow spread, and factor-adjusted price discovery.
- **R5. Liquidity regimes.** Propositions 7.1, 7.4, 11.1 and 11.2 separate factor-hedged dealer liquidity, bookmaker/insurer residual-diversification liquidity, and subsidized information-market liquidity. These are classification and benchmark results, not a separate equilibrium theory.
- **R6. Birth, stress, and concentration.** Proposition 9.1 gives the coordination problem; Proposition 10.8 gives the stress decomposition; and Proposition 10.10 shows why normal-time quote feasibility need not imply systemic robustness.

Vocabulary discipline

The program uses "residual" in several related but distinct senses. Paper 1's *valuation residual* is a missing payoff direction after projection onto the existing market span $\mathcal{H}$, weighted by the valuation-dispersion operator. Paper 2's *hedging residual* is the dealer risk left after projection onto hedgeable factor markets. Paper 3's *admissibility residual* is the externality or constraint burden left after technical feasibility. Keeping these residuals separate is load-bearing: a claim can be valuable in Paper 1, locally quotable in Paper 2, and still inadmissible in Paper 3.

The conclusion is intentionally disciplined. The paper does not claim that all long-tail claims will trade, that factor hedging eliminates adverse selection, or that residual diversification survives crises without assumptions. It claims that liquidity fragmentation is not mechanically proportional to claim count. Liquidity scales with the dimensions and costs of the risks a dealer must actually bear.

2 Related Literature and Positioning

The paper sits between factor pricing, derivative replication, dealer inventory models, adverse-selection microstructure, transaction-cost theory, liquidity coordination, and automated market-making.

Factor structure and replication. The payoff decomposition follows the arbitrage-pricing tradition of Ross (1976): high-dimensional asset returns may be driven by a lower-dimensional set of systematic factors plus idiosyncratic residuals. This paper uses the factor structure for liquidity production rather than primarily for expected-return restrictions. The replication logic is also close to Black and Scholes (1973) and Merton (1973). In option pricing, a derivative's hedgeable component can be manufactured from more liquid instruments. Here, the same idea is applied to long-tail state-contingent claims whose underlying exposures are not necessarily exchange-traded assets.

Dealer inventory and market-making. Classical market-microstructure models emphasize inventory risk, order-processing cost, and dealer balance sheets (Garman, 1976; Stoll, 1978; Ho and Stoll, 1981; Grossman and Miller, 1988). The quote-feasibility condition below is in this tradition, but the inventory is a cross-claim book in a high-dimensional claim space. The dealer's state variable is not just inventory in one security; it is the vector q, its aggregate factor exposure $B^\top q$, and its residual covariance exposure $q^\top \Sigma_\varepsilon q$. Dynamic inventory models such as Avellaneda and Stoikov (2008) provide a continuous-time benchmark for optimal quotes under inventory risk. This paper remains one-period and deliberately reduced-form to isolate the liquidity-capacity mechanism. It should be read as a mechanical break-even and capacity model, not as a substitute for equilibrium models of dealer competition, limit-order-book dynamics, search, or strategic quoting.

Immediacy, natural flow, and liquidity coordination. Demsetz (1968) frames the bid-ask spread as the cost of transacting and supplying immediacy. That perspective matters for long-tail markets because some flow is not primarily hedging or information production; it is demand for immediacy, entertainment, leverage, or convenience. Pagano (1989) emphasizes that liquidity depends on participant entry and can have coordination properties. This paper keeps the quote condition reduced-form, but it treats flow composition and liquidity coordination as primitives that determine whether a local quote becomes a durable market.

Adverse selection and information. Informed-trading models show that spreads compensate liquidity providers for trading against better-informed agents (Glosten and Milgrom, 1985; Kyle, 1985). Grossman and Stiglitz (1980) show that information acquisition cannot be costlessly competed away in equilibrium. In this paper, factor risk can be hedged mechanically, but residual information remains. Synthetic liquidity therefore does not solve adverse selection; it localizes it in ε_i.

Security design and incomplete markets. Paper 1 in this program is closer to the endogenous-security and incomplete-markets literature: Allen and Gale (1994), Duffie and Rahi (1995), Pesendorfer (1995), Demange and Laroque (1995), Ohashi (1995), Hara (1995), and Bisin (1998). The present paper does not rederive market-creation value. It takes valuable claims as given and asks how dealers can quote them. Security-design models with asymmetric information, including DeMarzo and Duffie (1999), are relevant because retention, tranching, collateral, and disclosure can reduce residual adverse selection.

Scoring rules and automated market makers. For claims with no hedgeable underlying, a different liquidity mechanism is needed. Market scoring rules, especially Hanson's logarithmic market scoring rule, provide bounded-loss liquidity for pure information claims (Hanson, 2003, 2007). Convex-cost formulations of prediction-market making connect this mechanism to online learning and automated market-making design (Abernethy et al., 2013). Practical liquidity-sensitive market makers are discussed by Othman et al. (2013). The citation details for the latter two should be verified before journal submission; they are included here as citation candidates for the automated-market-making literature.

Positioning. The novelty is not the existence of factor models, dealer inventory costs, informed trading, transaction costs, or LMSR. The contribution is the combination: a liquidity model in which sparse long-tail claims become quotable because systematic risk maps into a deep factor basis, or because broad non-toxic residual flow diversifies, or because a sponsor subsidizes information-market depth. The most distinctive standalone prediction is residual price discovery: after controlling for factor moves, signed order flow should forecast changes in the residual component of claim value, and that predictability should weaken after disclosure, oracle, or verification improvements. The capacity propositions make precise the statement that claim count can grow faster than natural liquidity-pool count, while also making precise where the statement breaks.

3 Environment

3.1 Claims, factors, and residuals

There is one terminal date. Random payoffs live in $L^2(\mathbb{P})$. The risk-free asset is available, and discounting is suppressed. There are N long-tail claims indexed by $i = 1, \dots, N$. Claim i has payoff

$$X_i = \alpha_i + \beta_i^\top F + \varepsilon_i,$$

where $F \in \mathbb{R}^d$ is a vector of hedgeable factor payoffs or returns, $\alpha_i \in \mathbb{R}$, $\beta_i \in \mathbb{R}^d$, and $\varepsilon_i \in L^2(\mathbb{P})$. In vector form,

$$X = \alpha + BF + \varepsilon,$$

where $X, \alpha, \varepsilon \in \mathbb{R}^N$ and $B \in \mathbb{R}^{N \times d}$ has row $\beta_i^\top$.

Assumption 3.1 (Residual orthogonality)**.** *The factor representation is chosen so that*

$$\mathbb{E}[\varepsilon] = 0, \qquad \mathrm{Cov}(\varepsilon, F) = 0, \qquad \mathrm{Var}(\varepsilon) = \Sigma_\varepsilon,$$

where Σ_ε is symmetric positive semidefinite.

Assumption 3.1 can be obtained by linear projection when factors are square-integrable and have nonsingular covariance. It does not require that factors price all claims perfectly. The whole point is that ε_i remains.

The factor vector F need not be a literal vector of primitive traded assets. It can be an exposure basis implemented through futures, ETFs, swaps, options, indexes, reinsurance, structured notes, offsetting claims, or dynamic baskets. The model treats the cost and imperfection of these hedges explicitly.

3.2 Dealer inventory and hedge costs

A dealer holds inventory $q \in \mathbb{R}^N$, where $q_i > 0$ means the dealer is long claim i. The unhedged payoff of the book is

$$q^\top X = q^\top \alpha + (B^\top q)^\top F + q^\top \varepsilon.$$

The *aggregate factor exposure* is

$$b(q) = B^\top q \in \mathbb{R}^d.$$

The dealer can trade factor hedges. Let $K_F(b) \geq 0$ denote the all-in cost of implementing factor exposure $-b$, including bid-ask spread, market impact, borrow cost, funding, margin, basis cost, and any residual hedge mismatch. A deep factor market is one in which $K_F(b)$ is small and slowly increasing over the relevant exposure region, not one in which hedging is free.

Residual risk is charged through a capital function

$$\mathcal{K}_R(q) = \psi(q^\top \Sigma_\varepsilon q),$$

where ψ is increasing and convex. The canonical quadratic specification is

$$\mathcal{K}_R(q) = \frac{\eta}{2} q^\top \Sigma_\varepsilon q,$$

where $\eta > 0$ is the dealer's effective residual-risk charge. This reduced-form parameter can represent risk aversion, regulatory capital, value-at-risk limits, expected-shortfall limits, internal balance-sheet costs, or shadow cost of scarce risk capital.

The dealer may also face a hard risk budget

$$\mathcal{K}_R(q) \leq K_R^{\max},$$

and a collateral or margin budget

$$M_F(B^\top q) + M_R(q) \leq C^{\max}.$$

Spreads become infinite, or quotes disappear, when the post-trade book violates these constraints.

3.3 Quotes and order signs

Let e_i be the i-th unit vector. A customer order changes dealer inventory by xe_i, where $x \in \mathbb{R}$. Thus $x > 0$ means the dealer buys claim i from the customer, and $x < 0$ means the dealer sells claim i to the customer. Let $s_i(x; q) \geq 0$ be the side-specific half-spread per unit relative to the dealer's marked mid. Spread revenue from the order is $|x| \, s_i(x; q)$.

The dealer's marked mid under information $\mathcal{D}$ is

$$m_i(\mathcal{D}) = \alpha_i + \beta_i^\top p_F + \mu_i(\mathcal{D}),$$

where p_F is the vector of factor prices or forward values and $\mu_i(\mathcal{D}) = \mathbb{E}[\varepsilon_i \mid \mathcal{D}]$ is the dealer's residual-value estimate. The theorem below is independent of the exact mid construction; it requires only that costs and residual inventory changes be evaluated around the dealer's mark.

Definition 3.2 (Natural and synthetic liquidity). *A claim has* natural liquidity *at size u if two-sided order flow in that claim alone supports execution at finite spread. A claim has* synthetic liquidity *at size u if a dealer can quote it at finite spread by hedging factor exposure in other markets and by financing the residual inventory risk. Natural liquidity is an own-market property. Synthetic liquidity is a property of a claim inside a cross-claim book.*

Definition 3.3 (Flow toxicity). *Let $\delta_i = \mathbb{E}[\varepsilon_i \mid \mathcal{S}] - \mathbb{E}[\varepsilon_i \mid \mathcal{D}]$ be the residual information advantage of the trader relative to the dealer, as formalized below. For a trade event T_i, define residual flow toxicity at claim i by*

$$\mathcal{T}_i = \mathbb{E}[|\delta_i| \mid T_i].$$

Flow is non-toxic at half-spread *s if $\mathcal{T}_i \le s$ after accounting for fees, rebates, and any non-informational trading costs. This is a local reduced-form definition, not a model of why traders choose to trade.*

Remark 3.4 (Flow ecology). *Order flow can come from hedging demand, immediacy demand, entertainment or consumption demand, noise or liquidity shocks, informed trading, speculation under disagreement, or manipulation. The same payoff can be liquid for different reasons depending on this flow mix. A market with abundant non-toxic flow can be easy to make even when the claim has little factor hedgeability. A market with toxic flow can be hard to make even when its factor exposure is easy to hedge.*

3.4 Cost terms

For a signed order xe_i, define:

- $k_i^F(x; q)$: incremental factor-hedge cost;
- $a_i(x; q, \mathcal{D})$: expected adverse-selection loss;
- $m_i(x; q)$: model, basis, and oracle-risk charge;
- $o_i(x)$: operating, settlement, compliance, and dispute-resolution cost;
- $f_i(x; q)$: funding, margin, or collateral cost not already included in k_i^F or $\mathcal{K}_R$.

The quote-feasibility theorem below does not require these costs to be linear, symmetric, or constant across claims. It treats them as primitives of the local quoting problem. A richer equilibrium model would endogenize some or all of them through dealer competition, informed-trader participation, order-arrival intensities, funding constraints, and dynamic inventory choice.

4 Factor-Hedge Decomposition Theorem

Theorem 4.1 (Factor-hedge decomposition). *Suppose $X = \alpha + BF + \varepsilon$ and Assumption 3.1 holds. For any inventory $q \in \mathbb{R}^N$,*

$$q^\top X = q^\top \alpha + (B^\top q)^\top F + q^\top \varepsilon.$$

If the dealer implements factor hedge $-B^\top q$, then the post-hedge risky payoff is $q^\top \varepsilon$, up to hedge execution and basis costs. Its variance is

$$\mathrm{Var}(q^\top \varepsilon) = q^\top \Sigma_\varepsilon q.$$

Proof. Substitute $X = \alpha + BF + \varepsilon$ into $q^\top X$:

$$q^\top X = q^\top \alpha + q^\top BF + q^\top \varepsilon = q^\top \alpha + (B^\top q)^\top F + q^\top \varepsilon.$$

A hedge of $-B^\top q$ offsets the factor payoff. The remaining random component is $q^\top \varepsilon$. Since $\mathrm{Var}(\varepsilon) = \Sigma_\varepsilon$,

$$\mathrm{Var}(q^\top \varepsilon) = q^\top \mathrm{Var}(\varepsilon) q = q^\top \Sigma_\varepsilon q.$$

□

Remark 4.2 (The liquidity source)**.** *The hedge markets are the liquidity source for the common component. The residual covariance matrix and the dealer's capital determine how much of the unhedged component can be warehoused. A sparse claim can be quoted when these two objects are favorable even if the claim's own order book is thin.*

Proposition 4.3 (Estimated betas turn factor risk into residual risk)**.** *Suppose true payoffs satisfy*

$$X = \alpha + BF + \varepsilon,$$

but the dealer hedges using estimated loading matrix $\widehat{B}$. Let $\Delta B = B - \widehat{B}$ and $\Omega_F = \mathrm{Var}(F)$. After hedging $\widehat{B}^\top q$, the unhedged payoff is

$$q^\top \varepsilon + (\Delta B^\top q)^\top F.$$

If $\mathrm{Cov}(\varepsilon, F) = 0$, its variance is

$$q^\top \Sigma_\varepsilon q + (\Delta B^\top q)^\top \Omega_F (\Delta B^\top q).$$

If $\|\Delta B\|_{op} \le \delta$, then

$$\mathrm{Var}(\textit{post-hedge payoff}) \le q^\top \Sigma_\varepsilon q + \lambda_{\max}(\Omega_F)\delta^2 \|q\|_2^2 .$$

Proof. The dealer hedges $\widehat{B}^\top q$, leaving true factor exposure $(B^\top q - \widehat{B}^\top q) = \Delta B^\top q$. Orthogonality gives zero covariance between $q^\top \varepsilon$ and $(\Delta B^\top q)^\top F$, yielding the variance formula. The bound follows from

$$(\Delta B^\top q)^\top \Omega_F (\Delta B^\top q) \le \lambda_{\max}(\Omega_F) \left\| \Delta B^\top q \right\|_2^2 \le \lambda_{\max}(\Omega_F)\delta^2 \|q\|_2^2 .$$

□

Remark 4.4 (Where data lowers spreads)**.** *Better data lowers spreads through at least three channels: it improves beta estimates, lowers oracle/model risk, and reduces residual adverse selection. In this paper all three enter the quote-feasibility condition separately.*

5 Residual Diversification and Capacity

Factor hedging removes the common component only if residual risk is financeable. The next results characterize when residual risk diversifies and when it does not.

Proposition 5.1 (Spectral residual bound)**.** *If $\lambda_{\max}(\Sigma_\varepsilon) \le \bar{\sigma}^2$, then for every inventory vector q,*

$$q^\top \Sigma_\varepsilon q \le \bar{\sigma}^2 \|q\|_2^2 .$$

Proof. For a symmetric positive semidefinite matrix, the Rayleigh quotient satisfies

$$q^\top \Sigma_\varepsilon q \le \lambda_{\max}(\Sigma_\varepsilon) \|q\|_2^2 .$$

Apply the assumed eigenvalue bound. □

Corollary 5.2 (Independent residuals)**.** *If residuals are independent and* $\mathrm{Var}(\varepsilon_i) \le \sigma^2$ *for all* i*, then*

$$q^\top \Sigma_\varepsilon q \le \sigma^2 \sum_{i=1}^{N} q_i^2.$$

If $q_i = Q/N$ *for all* i*, then*

$$q^\top \Sigma_\varepsilon q \le \sigma^2 \frac{Q^2}{N}.$$

Proof. Independence makes Σ_ε diagonal with diagonal entries bounded by σ^2. The equal-inventory formula follows by substituting $q_i = Q/N$. □

Corollary 5.3 (Equicorrelated residuals)**.** *Suppose residuals have common variance* σ^2 *and common pairwise correlation* $\rho \in [0, 1]$*. If* $q_i = Q/N$ *for all* i*, then*

$$q^\top \Sigma_\varepsilon q = \sigma^2 Q^2 \left(\rho + \frac{1-\rho}{N}\right).$$

As $N \to \infty$*, residual variance per unit of aggregate inventory vanishes if* $\rho = 0$ *and converges to* $\sigma^2 Q^2 \rho$ *if* $\rho > 0$*.*

Proof. With equal inventory,

$$q^\top \Sigma_\varepsilon q = \sum_i q_i^2 \sigma^2 + \sum_{i \ne j} q_i q_j \rho \sigma^2.$$

Substituting $q_i = Q/N$ gives

$$\sigma^2 \frac{Q^2}{N} + \rho \sigma^2 Q^2 \frac{N-1}{N} = \sigma^2 Q^2 \left(\rho + \frac{1-\rho}{N}\right).$$

□

Definition 5.4 (Effective residual breadth)**.** *For a nonzero inventory vector* q*, define gross inventory* $Q_1(q) = \|q\|_1$*. If residuals have common variance* σ^2*, define the effective residual breadth of* q *by*

$$N_{\mathrm{eff}}(q; \Sigma_\varepsilon) = \frac{\sigma^2 Q_1(q)^2}{q^\top \Sigma_\varepsilon q},$$

with $N_{\mathrm{eff}} = +\infty$ *when* $q^\top \Sigma_\varepsilon q = 0$*. Equivalently,*

$$q^\top \Sigma_\varepsilon q = \sigma^2 \frac{Q_1(q)^2}{N_{\mathrm{eff}}(q; \Sigma_\varepsilon)}.$$

For independent equal-weight inventory, $N_{\mathrm{eff}} = N$. For equicorrelated equal-weight inventory,

$$N_{\mathrm{eff}} = \frac{N}{1 + (N-1)\rho},$$

which converges to $1/\rho$ when $\rho > 0$. This is the effective number of independent residual bets in the book.

Proposition 5.5 (Residual capital capacity)**.** *Suppose residual capital charge is $\eta q^\top \Sigma_\varepsilon q/2$ and the dealer has residual-risk budget $K > 0$. If residuals are independent with common variance σ^2 and $q_i = Q/N$, then feasible aggregate inventory satisfies*

$$Q \leq \sqrt{\frac{2KN}{\eta\sigma^2}}.$$

If residuals are equicorrelated with correlation ρ, then

$$Q \leq \sqrt{\frac{2K}{\eta\sigma^2\,(\rho + (1-\rho)/N)}}.$$

More generally, for any book with common residual variance σ^2, feasible gross inventory satisfies

$$Q_1(q) \leq \sqrt{\frac{2K\, N_{\text{eff}}(q;\Sigma_\varepsilon)}{\eta\sigma^2}}.$$

Proof. Feasibility requires $\eta q^\top \Sigma_\varepsilon q/2 \leq K$. Substitute the variance formulas from Corollaries 5.2 and 5.3. The general expression follows from the definition of N_{eff}. □

Remark 5.6 (Capacity is about breadth, not labels)**.** *Adding more claim labels does not increase capacity unless it increases effective residual breadth, lowers residual covariance, reduces inventory concentration, or improves hedging and capital terms. Equal-weight independent residuals increase capacity at rate $\sqrt{N}$. Positive common residual correlation creates a capacity plateau.*

Theorem 5.7 (Liquidity dimension bound)**.** *Consider a family of claim sets indexed by N, each with payoff representation $X = \alpha + BF + \varepsilon$, $F \in \mathbb{R}^d$, and residual covariance $\Sigma_{\varepsilon,N}$. Let $\mathcal{Q}_N \subseteq \mathbb{R}^N$ be the inventory states in which a dealer is required to quote. Suppose there exist constants $H, R < \infty$, independent of N, such that for every $q \in \mathcal{Q}_N$,*

$$K_F(B^\top q) \leq H, \qquad q^\top \Sigma_{\varepsilon,N} q \leq R.$$

Then the dealer's balance-sheet liquidity cost

$$L_N(q) = K_F(B^\top q) + \frac{\eta}{2} q^\top \Sigma_{\varepsilon,N} q$$

is uniformly bounded over N and $q \in \mathcal{Q}_N$:

$$\sup_N \sup_{q \in \mathcal{Q}_N} L_N(q) \leq H + \frac{\eta R}{2}.$$

Thus claim count does not enter liquidity cost directly; it enters only by changing factor exposure, hedge cost, residual covariance, inventory concentration, and the inventory states the dealer must support.

Conversely, without uniform bounds on factor hedge cost and residual risk, there is no claim-count-invariance result. In particular, concentrated one-sided inventory and positive residual correlation can make required spreads or capital grow with aggregate exposure even when factor dimension d is fixed.

Proof. The bound is immediate from the definition of $L_N(q)$. The converse follows from counterexamples. If $q = Qe_i$, residual variance is $Q^2(\Sigma_\varepsilon)_{ii}$, independent of how many other claim labels exist. If $q_i = Q/N$ and residuals are equicorrelated with $\rho > 0$, Corollary 5.3 shows that residual variance converges to $\sigma^2 Q^2 \rho$, not zero. If factor exposure $B^\top q$ grows beyond the depth of factor markets, $K_F(B^\top q)$ can also grow without bound. □

Remark 5.8 (What the liquidity dimension bound does not say). *The theorem does not say that millions of claims can be quoted at tight spreads automatically. It says that the dimensionality of the hedge problem is d, while the residual problem is governed by covariance and inventory geometry. If adding claims worsens covariance, concentrates inventory, increases adverse selection, or crowds factor hedges, liquidity deteriorates.*

Proposition 5.9 (Bookmaker liquidity from diversified non-toxic flow). *Consider a family of claims with negligible hedgeable factor exposure over the relevant book, so $B^\top q \approx 0$. Suppose residual capital is quadratic, residuals have common variance σ^2, inventory is spread across effective residual breadth N_{eff}, and flow toxicity satisfies $\mathcal{T}_i \le \bar{a}$ for the marginal claims being quoted. For equal trade size u and gross inventory Q_1, a sufficient per-unit spread for local break-even quoting is bounded by*

$$s_i \ge \bar{a} + \bar{c}_{\text{op}} + \bar{c}_{\text{fund}} + \frac{\eta u}{2}\sigma^2 + \eta\sigma^2 \frac{Q_1}{N_{\text{eff}}},$$

where $\bar{c}_{\text{op}}$ and $\bar{c}_{\text{fund}}$ bound operating and funding/collateral charges per unit. Thus independent or weakly correlated non-toxic flow can support liquidity even without hedgeable factor exposure, while toxic flow or low effective breadth destroys the mechanism.

Proof. Set the factor hedge cost to zero in Theorem 6.2. Bound adverse-selection loss per unit by $\bar{a}$, operating and funding/collateral charges by $\bar{c}_{\text{op}}$ and $\bar{c}_{\text{fund}}$, and use $(\Sigma_\varepsilon q)_i \approx \sigma^2 Q_1/N_{\text{eff}}$ for a diversified common-variance book. The finite-trade own-variance term contributes $\eta u \sigma^2/2$. □

Remark 5.10 (Why independence can be valuable). *In factor-hedged liquidity, common components are useful because they can be offloaded into deep hedge markets. In bookmaker-style liquidity, common components are dangerous because they defeat diversification. The sportsbook case is therefore not an exception to the residual-capacity logic. It is the pure residual-diversification case: β_i is small, ε_i is the book, and the business works only while flow remains broad, balanced, and sufficiently uninformed.*

6 Reduced-Form Quote Feasibility and Spread Decomposition

A dealer quotes when spread revenue covers incremental costs and constraints remain satisfied. This section derives a sufficient one-period break-even condition. It is not a necessary condition for every possible trading protocol and not a full equilibrium theory of quoted spreads. It is a capacity test: given a book, a hedge basis, a risk charge, and cost terms, can the dealer fill this order without expected loss net of the modeled charges?

For a signed order xe_i, define the incremental residual variance

$$\Delta V_i(x; q) = (q + xe_i)^\top \Sigma_\varepsilon (q + xe_i) - q^\top \Sigma_\varepsilon q.$$

Expanding gives

$$\Delta V_i(x; q) = 2x(\Sigma_\varepsilon q)_i + x^2(\Sigma_\varepsilon)_{ii}. \tag{1}$$

Definition 6.1 (Local break-even quote feasibility). *Fix inventory q, order xe_i, information $\mathcal{D}$, cost terms, and hard risk and collateral constraints. A quote is* locally break-even feasible *if filling the order leaves the post-trade book inside the constraints and generates nonnegative expected profit after factor hedge cost, residual capital charge, adverse-selection loss, model/oracle risk, operating cost, and funding or collateral cost. Local feasibility is a sufficient capacity property of a particular dealer book. It is not an equilibrium selection rule for spreads.*

Theorem 6.2 (Quote-feasibility condition)**.** *Suppose residual capital is quadratic,*

$$\mathcal{K}_R(q) = \frac{\eta}{2} q^\top \Sigma_\varepsilon q.$$

A sufficient local break-even condition for filling a signed order xe_i, starting from inventory q, is

$$|x|\, s_i(x;q) \ge k_i^F(x;q) + a_i(x;q,\mathcal{D}) + m_i(x;q) + o_i(x) + f_i(x;q) + \frac{\eta}{2}\Delta V_i(x;q), \tag{2}$$

provided the post-trade inventory satisfies the dealer's hard risk and collateral constraints. Equivalently, for $x \neq 0$,

$$s_i(x;q) \ge \frac{k_i^F(x;q) + a_i(x;q,\mathcal{D}) + m_i(x;q) + o_i(x) + f_i(x;q)}{|x|} + \frac{\eta}{2\,|x|}\left(2x(\Sigma_\varepsilon q)_i + x^2(\Sigma_\varepsilon)_{ii}\right). \tag{3}$$

Proof. Spread revenue is $|x|\, s_i(x;q)$. The right side of (2) is the sum of incremental factor hedge cost, expected adverse-selection loss, model/oracle charge, operating cost, funding/collateral cost, and incremental residual capital charge. If revenue is at least this sum, expected profit net of the capital charge is nonnegative. Equation (3) follows by substituting (1) and dividing by $|x|$. □

Remark 6.3 (Sufficiency, not equilibrium)**.** *Theorem 6.2 is deliberately one-sided. It says that a dealer can cover the modeled mechanical costs of a fill at a given spread. It does not say that the quoted spread is unique, competitive, dynamically optimal, or robust to strategic order splitting. In an equilibrium model, the spread would also depend on dealer entry, order-arrival elasticities, queue position, inventory mean reversion, information acquisition, latency, search, and cross-subsidies across products. The theorem isolates the balance-sheet inequality that those richer models must embed.*

Corollary 6.4 (Marginal residual capital term)**.** *For an infinitesimal inventory change x in claim i, the first-order residual capital term is*

$$\eta x(\Sigma_\varepsilon q)_i.$$

For a finite trade, the per-unit residual capital term is

$$\eta\ \mathrm{sgn}(x)(\Sigma_\varepsilon q)_i + \frac{\eta\,|x|}{2}(\Sigma_\varepsilon)_{ii}.$$

Inventory-increasing trades have positive first-order capital cost; inventory-reducing trades can have negative first-order capital cost and may be quoted more tightly.

Proof. Differentiate $\eta\Delta V_i(x;q)/2$ with respect to x at zero and divide the finite-trade expression by $|x|$. □

Corollary 6.5 (Liquidity without own-flow)**.** *Claim i can be quoted without deep natural two-sided flow if the right side of* (3) *is finite, the post-trade book is feasible, and traders value execution at the resulting spread. Own-flow is sufficient for liquidity but not necessary. What is necessary in this local model is finite hedge cost, financeable residual risk, bounded informed-flow loss, bounded model/oracle risk, and acceptable operating cost. Market-wide depth additionally requires enough dealer capital, trader demand, and competitive or institutional support to sustain such local quotes over time.*

Remark 6.6 (Spread decomposition)**.** *The economically useful decomposition is*

$$\begin{aligned}\textit{half-spread} = {} & \textit{factor hedge cost} + \textit{residual capital charge} + \textit{adverse-selection charge} \\ & + \textit{model/oracle charge} + \textit{operating cost} + \textit{funding/collateral charge}.\end{aligned}$$

Automation attacks operating cost. Deep factor markets attack hedge cost. Diversification attacks residual capital cost. Data attacks beta error, oracle risk, and adverse selection. Capital, collateral, and regulation determine whether the remaining risk can be warehoused.

Proposition 6.7 (Quote capacity at a posted spread)**.** *Fix claim i, starting inventory q, and a maximum half-spread $\bar{s}$ acceptable to traders. Suppose non-capital costs satisfy*

$$k_i^F(x;q) + a_i(x;q,\mathcal{D}) + m_i(x;q) + o_i(x) + f_i(x;q) \le c_i\,|x| + \frac{\chi_i}{2}x^2$$

for $|x| \le \bar{u}$. A sufficient condition for quotes up to size $u \le \bar{u}$ on side $\operatorname{sgn}(x)$ *is*

$$\bar{s} \ge c_i + \eta\,\operatorname{sgn}(x)(\Sigma_\varepsilon q)_i + \frac{u}{2}\left(\chi_i + \eta(\Sigma_\varepsilon)_{ii}\right).$$

Thus locally feasible quote size decreases with own residual variance, same-side inventory covariance, convex hedge cost, and adverse-selection curvature. The result is a sufficient capacity bound, not an equilibrium depth schedule.

Proof. Substitute the assumed cost bound into (2), divide by $|x|$, and set $u = |x|$. □

7 Liquidity Regimes and Flow Ecology

The quote-feasibility condition is a cost stack. Which terms dominate depends on why traders arrive. This paper distinguishes three liquidity regimes.

Regime	Main liquidity source	Main failure mode
Factor-hedged dealer	common risk can be hedged in deep markets	hedge gaps, beta instability, residual capital
Bookmaker / insurer	broad non-toxic residual flow diversifies	sharp flow, one-sided flow, residual correlation
Subsidized information market	sponsor or protocol pays for price discovery	subsidy exhaustion, ambiguous settlement, manipulation

These regimes have different comparative statics. Factor-hedged liquidity improves when claims share hedgeable components. Bookmaker liquidity improves when residuals are independent and flow is non-toxic. Subsidized information-market liquidity improves when the sponsor values the public-good price enough to pay for depth.

Proposition 7.1 (Three-regime local liquidity frontier)**.** *Fix a claim i, inventory q, signed order xe_i, target half-spread s, and the dealer's hard risk and collateral constraints. Let*

$$R_i(x;q) = k_i^F(x;q) + a_i(x;q,\mathcal{D}) + m_i(x;q) + o_i(x) + f_i(x;q) + \frac{\eta}{2}\Delta V_i(x;q)$$

be the mechanical cost stack from Theorem 6.2. A local quote at half-spread s is mechanically feasible whenever constraints hold and at least one of the following closes the deficit:

(a) Factor-hedged dealer liquidity: *hedge markets and residual capital make* $R_i(x;q) \le |x|\, s$.

(b) Bookmaker or insurer residual-diversification liquidity: *factor exposure is negligible over the relevant book, effective residual breadth is high, flow toxicity is bounded, and the sufficient bound in Proposition 5.9 is below* s.

(c) Subsidized information-market liquidity: *a nonnegative outside budget* A_i *covers any local mechanical deficit,* $R_i(x;q) - |x|\, s \le A_i$*, or, in a scoring-rule implementation, the sponsor accepts the bounded-loss obligation in Proposition 11.1.*

If constraints fail, or if $R_i(x;q) - |x|\, s$ *remains positive after the applicable residual-flow and subsidy terms, the claim is not locally quote-feasible at spread* s *in this reduced-form model.*

Proof. Part (a) is exactly Theorem 6.2. Part (b) is Proposition 5.9 applied to the special case in which the factor hedge term is negligible and residual breadth plus bounded flow toxicity control residual capital and adverse selection. Part (c) adds an outside loss budget to the dealer's local break-even inequality, or uses the LMSR loss bound as the budgeted mechanism for an information market. If hard constraints fail, the feasibility premise of Theorem 6.2 fails; if a positive deficit remains, spread revenue plus subsidy does not cover the modeled cost stack. □

Definition 7.2 (Flow-mix vector)**.** *For a claim* i*, let*

$$\lambda_i = (\lambda_i^H, \lambda_i^I, \lambda_i^N, \lambda_i^S, \lambda_i^M)$$

denote the local intensity of hedging flow, immediacy or convenience flow, noise or entertainment flow, speculative/informed flow, and manipulative flow. The vector is not directly observed, but it determines both the spread and the welfare interpretation of the market.

Remark 7.3 (Why flow mix matters)**.** *Hedging flow can create risk-sharing surplus. Immediacy flow pays for execution service. Noise or entertainment flow can be privately voluntary without producing risk-sharing value. Informed flow creates price discovery when the price is socially useful, but it is also adverse selection for the liquidity provider. Manipulative flow can make the market socially harmful even when volume is high. Liquidity and welfare therefore cannot be inferred from volume alone.*

Proposition 7.4 (Dominant-regime classification)**.** *For a claim* i*, define normalized cost shares at a feasible quote by*

$$\omega_i^F = \frac{k_i^F}{|x|\, s_i}, \quad \omega_i^R = \frac{\eta \Delta V_i/2}{|x|\, s_i}, \quad \omega_i^A = \frac{a_i}{|x|\, s_i}, \quad \omega_i^O = \frac{m_i + o_i + f_i}{|x|\, s_i}.$$

If ω_i^F *is small because* $B^\top q$ *is cheaply hedged and* ω_i^R *is controlled by residual breadth, liquidity is factor-hedged. If* $B^\top q \approx 0$*,* ω_i^R *is controlled by high* N_{eff}*, and* ω_i^A *is small because flow is non-toxic, liquidity is bookmaker-style residual diversification. If spreads are below the mechanical break-even stack because an outside budget absorbs losses, liquidity is subsidized information-market liquidity.*

Proof. The classification is a restatement of (2) by dominant cost share. Factor-hedged liquidity lowers the hedge term; bookmaker liquidity lowers residual capital through breadth and adverse selection through non-toxic flow; subsidized liquidity relaxes the break-even requirement by adding an outside loss budget. □

8 Hedging Residual Information and Adverse Selection

Synthetic liquidity does not eliminate information problems. It separates the part of the claim that can be hedged mechanically from the part about which specialists may know more than the dealer.

Let $\mathcal{D}$ be the dealer's information and $\mathcal{S}$ a specialist trader's information, with $\mathcal{D} \subseteq \mathcal{S}$. The conditional fair value of claim i under information $\mathcal{I}$ is

$$p_i(\mathcal{I}) = \alpha_i + \beta_i^\top \mathbb{E}[F \mid \mathcal{I}] + \mathbb{E}[\varepsilon_i \mid \mathcal{I}].$$

If factor prices are public or hedgeable at quoted prices, informational advantage in the factor component is competed away or hedged. The residual conditional expectation remains.

Theorem 8.1 (Residual information wedge)**.** *Suppose factor values are common knowledge at the time of trade or can be hedged at quoted factor prices. Then the specialist's informational edge over the dealer in claim i is*

$$\delta_i(\mathcal{S}, \mathcal{D}) = \mathbb{E}[\varepsilon_i \mid \mathcal{S}] - \mathbb{E}[\varepsilon_i \mid \mathcal{D}].$$

The same residual wedge is the source of:

(i) specialist alpha from trading claim i;

(ii) dealer adverse-selection loss;

(iii) the informed-flow component of the spread;

(iv) factor-adjusted residual price discovery.

Proof. Subtract the dealer's conditional fair value from the specialist's:

$$p_i(\mathcal{S}) - p_i(\mathcal{D}) = \beta_i^\top (\mathbb{E}[F \mid \mathcal{S}] - \mathbb{E}[F \mid \mathcal{D}]) + (\mathbb{E}[\varepsilon_i \mid \mathcal{S}] - \mathbb{E}[\varepsilon_i \mid \mathcal{D}]).$$

If factor values are common or hedgeable at quoted prices, the first term is not a residual source of claim-specific trading alpha. The remaining term is $\delta_i(\mathcal{S}, \mathcal{D})$. A specialist who observes positive δ_i buys, one who observes negative δ_i sells. The dealer's expected loss from such order selection is the same wedge with opposite sign. Empirical residual price discovery is the subsequent incorporation of this wedge into public residual-value estimates. □

Proposition 8.2 (Adverse-selection lower bound)**.** *Consider orders of fixed size $u > 0$ in claim i. Suppose informed traders choose the trade side to exploit residual signal δ_i: they buy from the dealer when $\delta_i > 0$ and sell to the dealer when $\delta_i < 0$. Conditional on trade event T, the dealer's expected residual loss at the mid is*

$$u \, \mathbb{E}\left[|\delta_i| \mid T\right].$$

Therefore any feasible half-spread must satisfy

$$s_i \geq \mathbb{E}\left[|\delta_i| \mid T\right]$$

before factor hedge cost, residual capital, model/oracle risk, operating cost, and funding cost are added. Factor hedging does not reduce this lower bound.

Proof. When $\delta_i > 0$, the trader buys from the dealer; the dealer sells a claim whose residual value is higher than the dealer's estimate by δ_i, losing $u\delta_i$ at the mid. When $\delta_i < 0$, the trader sells to the dealer; the dealer buys a claim whose residual value is lower than the dealer's estimate by $|\delta_i|$, losing $u\,|\delta_i|$. Conditional on trade, the expected mid loss is $u\mathbb{E}[|\delta_i| \mid T]$. A half-spread below this amount per unit loses money against this order-selection rule even if factor risk is perfectly hedged. □

Remark 8.3 (Endogenous trade sets). *In equilibrium, T depends on the quoted spread. For example, if informed traders trade only when $|\delta_i| > s_i + c$, the break-even spread solves a fixed-point condition involving $\mathbb{E}[|\delta_i| \mid |\delta_i| > s_i + c]$. The proposition states the conditional lower bound for any realized informed-flow set.*

Corollary 8.4 (Residual price discovery test). *Let $\Delta p_{i,t+1}$ be the post-trade price change of claim i, and let $\beta_i^\top \Delta p_{F,t+1}$ be the contemporaneous factor-implied price change. Under the residual-information model, signed order flow should forecast the factor-adjusted price change*

$$\Delta p_{i,t+1} - \beta_i^\top \Delta p_{F,t+1}$$

when order flow contains residual information. The predictive component should be stronger in claims with high information asymmetry and weaker after disclosure, oracle, or verification improvements.

Proof. By Theorem 8.1, informed order flow is informative about $\mathbb{E}[\varepsilon_i \mid \mathcal{S}] - \mathbb{E}[\varepsilon_i \mid \mathcal{D}]$, not about already-hedged factor value. Public learning from order flow therefore appears in the residual component after controlling for factor moves. □

Remark 8.5 (Specialists still matter). *The dealer can quote the factor component without knowing every microclimate, product cohort, creator-audience curve, restaurant lease, SKU substitution pattern, or idiosyncratic borrower detail. Specialists who know those residuals can still earn alpha. Synthetic liquidity creates a surface on which residual information can trade; it does not make that information irrelevant.*

9 Liquidity Coordination and Market Birth

The quote-feasibility condition is local. A market also needs enough traders to arrive, enough market makers to compete, and enough volume to justify fixed operating cost. Liquidity can therefore be a coordination equilibrium: traders route to markets that are liquid, and markets become liquid because traders route there.

Proposition 9.1 (Self-fulfilling liquidity threshold). *Let $n \geq 0$ be expected participant mass in a claim market and let quoted half-spread be $s(n)$, with $s'(n) < 0$ over the relevant range because order-flow breadth, competition, and inventory netting improve with participation. Suppose a representative marginal trader participates when*

$$b - s(n) \geq 0,$$

where b is the trader's execution, hedging, information, or consumption benefit. If

$$b - s(0) < 0 \qquad \text{and} \qquad b - s(\bar{n}) > 0$$

for some $\bar{n} > 0$, then both a low-liquidity nonparticipation outcome and a high-liquidity participation outcome can be self-consistent. A subsidy, designated market maker, benchmark mandate, or platform routing rule can select the high-liquidity outcome by temporarily raising n or lowering $s(n)$.

Proof. At $n = 0$, the marginal participation condition fails, so nonparticipation is self-consistent. At $\bar{n}$, the condition holds, so participation at that scale is self-consistent if the venue can coordinate that participant mass. The comparative-static interventions change either the spread schedule or the participant mass, moving the market across the threshold. □

Remark 9.2 (Fixed costs versus coordination). *Infrastructure under-entry is a fixed-cost problem. Liquidity coordination is an expectations problem. Both can justify platform or public intervention, but the intervention differs: fixed costs require building shared rails; coordination requires selecting the liquid equilibrium through seed flow, maker obligations, routing, or temporary subsidy.*

10 Stress Fragility Results

Synthetic liquidity is normal-time liquidity. Its failure modes are exactly the state variables of the model: residual covariance, hedge cost, order-flow balance, beta stability, collateral, and data/oracle reliability.

10.1 Residual correlation shocks

Proposition 10.1 (Correlation shock). *Suppose $q_i = Q/N$, residuals have common variance σ^2, and residual correlation jumps from ρ_0 to $\rho_1 > \rho_0$. Residual variance increases by*

$$\Delta \operatorname{Var} = \sigma^2 Q^2 (\rho_1 - \rho_0) \left(1 - \frac{1}{N}\right).$$

With quadratic residual capital charge $\eta/2$, capital required increases by

$$\frac{\eta}{2} \sigma^2 Q^2 (\rho_1 - \rho_0) \left(1 - \frac{1}{N}\right).$$

Proof. Subtract the equicorrelation formula in Corollary 5.3 at ρ_0 from the same formula at ρ_1. Multiply by $\eta/2$ for capital. □

Remark 10.2. *Diversification is an assumption about covariance, not a law of nature. If residuals become correlated in stress, the law of large numbers stops protecting the dealer precisely when capital is most expensive.*

10.2 Hedge-gap stress

Proposition 10.3 (Hedge-gap stress). *Let normal-time factor hedge cost for order xe_i be $k_i^F(x;q)$, and stress-time cost be $k_{i,s}^F(x;q) > k_i^F(x;q)$. Holding other terms fixed, the required half-spread in Theorem 6.2 increases by at least*

$$\frac{k_{i,s}^F(x;q) - k_i^F(x;q)}{|x|}.$$

If stress also changes factor-market depth so that the hedge cannot be implemented within the dealer's risk limit, the quote is infeasible at any finite spread acceptable before the constraint binds.

Proof. The hedge-cost term enters (3) additively and linearly after division by $|x|$. If the hedge cannot be implemented within hard constraints, the feasibility premise of Theorem 6.2 fails. □

10.3 One-sided flow and inventory concentration

Proposition 10.4 (One-sided flow capacity). *Suppose a dealer repeatedly absorbs same-side orders of size u in a single claim i, residual variance is $\sigma_i^2 = (\Sigma_\varepsilon)_{ii}$, and there are no offsetting residual* *covariances. With residual-risk budget K and quadratic capital charge $\eta/2$, the maximum number of orders that can be absorbed before the residual budget binds is*

$$n_i^{\max} = \left\lfloor \frac{1}{u}\sqrt{\frac{2K}{\eta\sigma_i^2}} \right\rfloor .$$

If the same aggregate signed inventory $Q = nu$ is instead spread equally across N independent claims with common variance σ^2, the feasible aggregate inventory is $Q \leq \sqrt{2KN/(\eta\sigma^2)}$. Concentration removes the diversification benefit.

Proof. A concentrated position $q = nue_i$ has residual variance $n^2u^2\sigma_i^2$. Feasibility requires $\eta n^2u^2\sigma_i^2/2 \leq K$. Solve for n. The diversified formula is Proposition 5.5. □

10.4 Beta and model-error stress

Proposition 10.5 (Beta-error stress)**.** *Let beta-estimation error rise from $\|\Delta B\|_{op} \leq \delta_0$ to $\|\Delta B\|_{op} \leq \delta_1$, with $\delta_1 > \delta_0$. Under the conditions of Proposition 4.3, the model-error component of residual variance can increase by at most*

$$\lambda_{\max}(\Omega_F)(\delta_1^2 - \delta_0^2)\,\|q\|_2^2 .$$

The resulting increase in quadratic capital charge is at most

$$\frac{\eta}{2}\lambda_{\max}(\Omega_F)(\delta_1^2 - \delta_0^2)\,\|q\|_2^2 .$$

Proof. Apply the bound in Proposition 4.3 before and after the beta-error shock and subtract. □

10.5 Collateral haircut shocks

Proposition 10.6 (Collateral haircut shock)**.** *Suppose a dealer must post collateral*

$$M(q) = h_F \left\| B^\top q \right\|_2 + h_R\sqrt{q^\top \Sigma_\varepsilon q}$$

against factor hedges and residual risk, with collateral budget $C^{\max}$. If haircuts rise from (h_F, h_R) to $(h'_F, h'_R) \geq (h_F, h_R)$, the feasible inventory set weakly shrinks:

$$\{q : M'(q) \leq C^{\max}\} \subseteq \{q : M(q) \leq C^{\max}\} .$$

For any ray $q = tq_0$, feasible scale falls from

$$t \leq \frac{C^{\max}}{h_F \|B^\top q_0\|_2 + h_R\sqrt{q_0^\top \Sigma_\varepsilon q_0}}$$

to the same expression with (h'_F, h'_R).

Proof. Since $(h'_F, h'_R) \geq (h_F, h_R)$, $M'(q) \geq M(q)$ for every q, so the feasible set shrinks. Homogeneity of $M(tq_0) = tM(q_0)$ gives the ray formula. □

10.6 Oracle and data failure

Proposition 10.7 (Oracle-error residual)**.** *Suppose claim i is settled not on X_i but on observed payoff*

$$\widetilde{X}_i = X_i + \zeta_i,$$

where ζ_i is oracle, measurement, or settlement error with $\mathbb{E}[\zeta_i] = 0$, $\mathrm{Cov}(\zeta, F) = 0$, and covariance Σ_ζ. If $\mathrm{Cov}(\varepsilon, \zeta) = 0$, then post-hedge residual covariance becomes

$$\Sigma_\varepsilon + \Sigma_\zeta.$$

The incremental residual capital term for order xe_i increases by

$$\frac{\eta}{2}\left(2x(\Sigma_\zeta q)_i + x^2(\Sigma_\zeta)_{ii}\right).$$

If oracle failure creates nonzero expected settlement bias conditional on trader information, that bias also enters the adverse-selection charge.

Proof. The hedged residual is $q^\top \varepsilon + q^\top \zeta$. Under the stated covariance assumptions, its variance is $q^\top(\Sigma_\varepsilon + \Sigma_\zeta)q$. The incremental capital formula follows from (1) with Σ_ζ. Conditional expected bias is an information wedge and therefore enters a_i. □

Proposition 10.8 (Stress spread decomposition)**.** *Let $s_i^0(x; q)$ be a normal-time feasible half-spread and $s_i^s(x; q)$ a stress-time feasible half-spread. Holding order size and starting inventory fixed, any stress changes in hedge cost, residual covariance, adverse selection, model/oracle risk, operating cost, and funding/collateral cost imply*

$$s_i^s(x;q) - s_i^0(x;q) \geq \frac{1}{|x|}\Big(\Delta k_i^F + \Delta a_i + \Delta m_i + \Delta o_i + \Delta f_i + \frac{\eta}{2}\left[\Delta V_i^s(x;q) - \Delta V_i^0(x;q)\right]\Big),$$

whenever both quotes are feasible. If stress violates hard hedge, collateral, or residual-risk constraints, the quote is infeasible rather than merely wider.

Proof. Write the feasibility inequality (2) in normal time and stress time and subtract. If constraints fail, the sufficient condition cannot be satisfied by any finite spread within the model because the dealer cannot carry the post-trade book. □

Remark 10.9 (Normal-time liquidity)**.** *Synthetic liquidity is not a promise of continuous depth in every state. It is a mechanism for producing normal-time quotes when common risk is hedgeable and residual risk is financeable. Stress states are precisely the states in which hedge costs, correlations, model errors, adverse selection, and haircuts move together.*

10.7 Residual warehouses and tail concentration

Long-tail liquidity can diversify everyday risk while concentrating residual tail risk on a small number of balance sheets. This is the systemic dark side of synthetic liquidity and bookmaker liquidity.

Proposition 10.10 (Residual warehouse concentration). *Suppose aggregate residual inventory q is warehoused by L liquidity providers, with provider ℓ holding q^ℓ and $\sum_{\ell=1}^{L} q^\ell = q$. If each provider faces quadratic residual capital $\eta(q^\ell)^\top \Sigma_\varepsilon q^\ell/2$, then for an equal split $q^\ell = q/L$, total private capital charge is*

$$\sum_{\ell=1}^{L} \frac{\eta}{2}(q^\ell)^\top \Sigma_\varepsilon q^\ell = \frac{1}{L}\frac{\eta}{2} q^\top \Sigma_\varepsilon q.$$

However, if stress losses have a fixed distress cost D_ℓ when provider ℓ's loss exceeds its capital buffer, the social cost of concentration includes the tail event

$$\sum_{\ell=1}^{L} D_\ell \Pr\left(\left|(q^\ell)^\top \varepsilon\right| > \mathit{buffer}_\ell\right),$$

which need not fall with normal-time variance charges. Normal-time quote feasibility therefore does not imply systemic robustness.

Proof. The quadratic capital identity follows by substituting $q^\ell = q/L$. The distress term is not implied by the quadratic charge; it is an additional externality from balance-sheet failure. Its probability depends on buffers, tail dependence, liquidation dynamics, and common shocks, not only on the normal-time covariance matrix. □

Remark 10.11 (The 2008 warning). *The attractive normal-time result is that many long-tail claims can be quoted by consolidating risk into hedge and residual books. The dangerous version is that many markets depend on a few residual warehouses whose failure forces quote withdrawal or liquidation across otherwise distinct claims. This is a Paper 3 admissibility and systemic-risk issue, not just a Paper 2 spread issue.*

11 Subsidized Liquidity for Pure Information Claims

Some claims have no useful hedgeable underlying. A one-off event claim, governance question, scientific replication question, or policy forecast may be almost pure information. For such claims, liquidity can be supplied by subsidy rather than factor hedging. Scoring-rule market makers are the canonical mechanism.

Consider K mutually exclusive and exhaustive outcomes. Let q_k be the number of shares outstanding for outcome k. The logarithmic market scoring rule (LMSR) cost function is

$$C(q) = b \log\left(\sum_{k=1}^{K} e^{q_k/b}\right),$$

where $b > 0$ is a liquidity parameter. The price of outcome k is

$$p_k(q) = \frac{\partial C}{\partial q_k} = \frac{e^{q_k/b}}{\sum_{\ell=1}^{K} e^{q_\ell/b}}.$$

Prices are positive and sum to one.

Proposition 11.1 (LMSR bounded-loss liquidity). *Starting from $q = 0$, the LMSR market maker's worst-case loss is bounded by*

$$b \log K.$$

Increasing b increases depth and increases the worst-case subsidy proportionally.

Proof. If outcome j occurs, the market maker pays q_j. It has collected $C(q) - C(0)$. Its loss is

$$q_j - [C(q) - C(0)] = q_j - C(q) + b \log K.$$

Because $C(q) = b \log \sum_k e^{q_k/b} \geq b \log e^{q_j/b} = q_j$, the loss is at most $b \log K$. The Hessian is

$$\nabla^2 C(q) = \frac{1}{b} \left(\mathrm{diag}(p(q)) - p(q)p(q)^\top \right),$$

so larger b lowers instantaneous price impact and raises depth. The loss bound scales linearly with b. □

Proposition 11.2 (Subsidy liquidity is not factor or bookmaker liquidity)**.** *For a pure information claim with no hedgeable factor representation useful to the dealer and no broad non-toxic residual-flow book to diversify, an LMSR-style market supplies depth by accepting a bounded loss subsidy. The sponsor pays for price discovery. In contrast, factor-hedged liquidity supplies depth by offsetting common risk and charging for residual risk, while bookmaker-style liquidity supplies depth by diversifying broad residual flow. These mechanisms are economically distinct even when they produce similar user-facing quotes.*

Proof. In the factor-hedged mechanism, the dealer's risk is reduced by trading hedge instruments and residual risk is priced through a capital charge. In the bookmaker mechanism, residual risk is made tolerable by breadth, limits, and low flow toxicity. In LMSR, the market maker's state-dependent liability is not offset by hedge instruments or diversified away by customer flow. The cost function bounds worst-case loss by design. The source of liquidity is therefore an ex ante subsidy budget rather than a hedge-and-capital book or residual-flow warehouse. □

Remark 11.3 (When subsidy is natural)**.** *Pure information prices can be public goods. A sponsor may rationally pay $b \log K$ to produce a probability estimate that improves decisions elsewhere. This is not a refutation of market liquidity; it is a different source of liquidity.*

12 Empirical Predictions and Test Plan

The model predicts that long-tail liquidity appears first where at least one liquidity regime is favorable: factor hedgeability is high, residual covariance is low or diversifiable with non-toxic flow, or an information sponsor is willing to subsidize depth. Its most distinctive empirical prediction is not merely that hedgeable claims have lower spreads. It is that order flow in synthetically liquid claims should reveal information in the factor-adjusted residual: signed trades forecast residual price changes when traders know something about ε_i, and that predictability should weaken when disclosure, oracle quality, or verification improves.

12.1 Spread decomposition tests

The quote-feasibility theorem suggests the reduced-form spread regression

$$\begin{aligned}\mathrm{Spread}_{i,t} = \gamma_0 &+ \gamma_1 \mathrm{HedgeCost}_{i,t} + \gamma_2 \mathrm{MarginalResidualCapital}_{i,t} + \gamma_3 \mathrm{AdverseSelection}_{i,t} \\ &+ \gamma_4 \mathrm{ModelOracleRisk}_{i,t} + \gamma_5 \mathrm{OperatingCost}_{i,t} + \gamma_6 \mathrm{FundingCollateral}_{i,t} + \mu_i + \tau_t + u_{i,t}.\end{aligned}$$

This regression should be read as a decomposition test, not as structural estimation of an equilibrium quoting game. The theory predicts which mechanical cost and capacity variables should move spreads, holding the market institution fixed; it does not by itself identify dealer markups, strategic quote shading, or entry margins.

The distinctive term is marginal residual capital:

$$\eta(\Sigma_\varepsilon q)_i + \frac{\eta u}{2}(\Sigma_\varepsilon)_{ii},$$

which varies with dealer inventory and residual covariance, not merely claim-level volatility.

12.2 Predictions

1. **Factor hedgeability.** Claims with higher factor R^2 and deeper hedge instruments should have lower spreads and greater quoted depth, holding residual information risk fixed.
2. **Benchmark creation.** Creation of a new hedge benchmark, index, futures contract, ETF, swap, or reinsurance instrument should reduce spreads most for claims with high loading on the new factor.
3. **Residual breadth.** Dealer capacity should increase with effective residual breadth N_{eff}, not with raw claim count. Equal-weight independent residual books should support more aggregate inventory than concentrated or correlated books.
4. **Bookmaker liquidity.** Markets with high independent event breadth, balanced small-stake flow, and weak post-trade drift should support liquidity even with low factor R^2.
5. **Flow toxicity.** Sharp or informed participation should predict lower limits, wider spreads, account restrictions, delayed settlement, or quote withdrawal.
6. **Stress correlation.** Stress episodes should widen spreads most in claims whose residuals become correlated after factor controls.
7. **One-sided flow.** Same-side order imbalances should predict spread widening and quote withdrawal through the term $(\Sigma_\varepsilon q)_i$.
8. **Beta instability.** Claims with unstable factor loadings should have wider spreads, especially when factor volatility rises.
9. **Oracle and data improvements.** Launches of reliable oracles, standardized data feeds, audited APIs, or settlement standards should reduce model/oracle charges and adverse-selection charges.
10. **Residual price discovery.** Signed order flow should forecast factor-adjusted price changes in claims with high residual information asymmetry.
11. **Liquidity coordination.** Temporary subsidies, designated market makers, or platform routing should have persistent effects when they move a market across the participation threshold.
12. **Warehouse concentration.** Markets dependent on a small number of residual-risk warehouses should show abrupt quote withdrawal or cross-claim spread widening when those balance sheets are stressed.
13. **Subsidized information markets.** In LMSR-style markets, depth should scale with b, and realized sponsor loss should respect the $b \log K$ bound up to fees and implementation details.

12.3 Empirical proxies

Theoretical object	Empirical proxy
Factor hedgeability	R^2 against traded factors; availability and depth of ETFs, futures, swaps, options, reinsurance, indexes, or baskets
Factor hedge cost	bid-ask spread, market impact, borrow cost, margin, funding spread, hedge slippage
Residual covariance	covariance of factor-adjusted returns or payoffs; stress residual correlations; covariance forecast error
Effective residual breadth	inverse Herfindahl of residual exposures; N_{eff}; concentration of dealer inventory
Marginal residual capital	$(\Sigma_\varepsilon q)_i$, $(\Sigma_\varepsilon)_{ii}$, VaR/ES utilization, capital budget usage
Adverse selection	factor-adjusted post-trade drift, informed-flow proxies, insider participation, disclosure asymmetry
Flow toxicity	post-trade drift by customer class, limit cuts, account restrictions, quote rejection, sharp-versus-recreational share
Natural flow breadth	number of small independent counterparties, bet-size dispersion, event independence, balanced side flow
Liquidity coordination	designated-maker obligations, seed subsidies, routing changes, venue migration, launch promotions
Residual warehouse concentration	market-maker concentration, clearing-member exposure, capital utilization, cross-claim quote withdrawal
Model/oracle risk	beta instability, forecast revisions, data latency, oracle downtime, dispute frequency
Operating cost	contract standardization, settlement automation, custody cost, compliance cost, dispute cost
Funding/collateral	haircuts, margin requirements, collateral eligibility, rehypothecation constraints

12.4 Research designs

Difference-in-differences around hedge-basis creation. Compare claims with high versus low loadings on a newly created hedge factor before and after the launch of a benchmark, index, futures contract, ETF, swap, or reinsurance facility. The model predicts spread compression and depth expansion for high-loading claims.

Worked design: hedge-basis shock. Let event date t_0 be the launch of a new traded hedge instrument $f^\star$, such as an index future, ETF, swap, weather index, catastrophe-risk facility, or standardized reinsurance layer. Estimate each claim's pre-event exposure $\hat{\beta}_i^\star$ using only pre-t_0 data, and define treated claims as those in the top quantile of $\left|\hat{\beta}_i^\star\right|$. For spread, quoted depth, or quote-withdrawal outcome $Y_{i,t}$, estimate

$$Y_{i,t} = \mu_i + \tau_t + \theta\, \mathbf{1}\{t \geq t_0\} \left|\hat{\beta}_i^\star\right| + \Gamma Z_{i,t} + u_{i,t},$$

where $Z_{i,t}$ controls for claim volatility, own-flow, residual information proxies, and venue effects. The synthetic-liquidity prediction is $\theta < 0$ for spreads and quote withdrawals and $\theta > 0$ for depth. The falsification test is that the effect should be weak for low-loading claims and should not appear before t_0. The residual-capacity version of the test asks whether the event reduces factor-hedge cost while leaving factor-adjusted adverse-selection and oracle-risk terms as remaining spread components.

Stress event studies. Estimate residual correlations before and during stress episodes. The model predicts that spread widening is increasing in the change in residual covariance after controlling for factor volatility and hedge-market depth.

Inventory-sensitive spread tests. Use dealer or market-maker inventory proxies to test whether quotes widen in the direction that increases $(\Sigma_\varepsilon q)_i$ and tighten for inventory-reducing trades.

Oracle and disclosure shocks. Study data-standard adoption, audited API launches, oracle upgrades, legal-template adoption, or reporting-rule changes. The model predicts lower model/oracle charges and weaker factor-adjusted adverse price movement after trades.

Residual alpha tests. Decompose returns into factor and residual components. Specialist returns should load on residual forecasting skill, not merely on factor exposure, after controlling for hedge costs and spreads.

12.5 Candidate empirical surfaces

The theory can be tested in listed options across sparse strikes and maturities; single-name credit and CDS index markets; insurance-linked securities and catastrophe risk; parametric weather and crop-risk markets; energy, congestion, and demand-response claims; revenue-share and royalty markets; tokenized private-asset claims; prediction markets; and automated market makers for event claims.

13 Relation to the Broader Research Program

This paper is Paper 2 of the *Trillion Markets* research program. Its role is mechanical and deliberately bounded.

Paper 1: valuation-residual surplus. Paper 1 asks which payoff subspaces are worth creating. Its central object is valuation dispersion projected onto the payoff residual that existing markets fail to span. Its question is: when does a missing payoff basis have enough value, net of implementation cost, to become a market?

Paper 2: hedging residual risk. This paper asks how claims can be quoted once they exist or are worth creating. Its central object is the residual risk that factor hedges fail to eliminate. Its question is: when can a dealer provide liquidity by hedging common risk and warehousing the hedging residual?

Paper 3: admissibility. Paper 3 should ask which markets should be allowed, restricted, subsidized, or prohibited. This paper does not solve the normative problem. It includes oracle, manipulation, collateral, and adverse-selection terms only insofar as they affect quote feasibility. Claim-centric regulation belongs in Paper 3.

Papers 4 and 5. Paper 4 should develop the shared infrastructure of state-contingent claims across prediction, insurance, derivatives, and parametric products. Paper 5 should develop the agentic transaction-cost theory. This paper uses infrastructure and automation only as inputs to hedge cost, operating cost, oracle risk, and quote feasibility.

Program bridge. The shared language is projection:

- Paper 1 valuation-residual surplus is the value of payoff residuals after projecting candidate claims onto the existing market span and weighting the remaining direction by valuation dispersion.
- Paper 2 hedging residual risk is what remains after projecting claim inventory onto hedgeable factor markets.
- Paper 2 flow ecology explains whether liquidity comes from hedging demand, non-toxic residual diversification, immediacy, information production, entertainment, or speculation.
- Paper 3 admissibility should govern which residual claims may be created or quoted once costs fall enough to make them possible.
- Papers 4 and 5 explain the institutional and agentic rails that may lower those costs.

In the computational notation of Paper 5, this paper supplies the liquidity state L_t: hedge bases, residual covariance, effective breadth, flow toxicity, market-maker balance sheets, collateral, and subsidy budgets. A market-search platform that estimates $V(B \mid \mathcal{H})$ but not L_t can identify valuable claims that no one can quote.

The public essay can say that liquidity sources scale with risk and flow ecology rather than claims. This paper states the disciplined version: liquidity costs are functions of factor hedge cost, residual effective rank, inventory concentration, adverse selection, flow toxicity, model/oracle risk, operating cost, and capital constraints. Claim count alone is not the state variable.

14 Conclusion

Long-tail market creation does not require a separate crowd for every claim. It requires at least one viable liquidity mechanism: a mapping from claims to hedgeable factors, a broad book of non-toxic independent residual flow, or a sponsor willing to pay for information-market depth. It also requires a capital model for residual covariance, mechanisms for bounding informed-flow losses, and operating infrastructure cheap enough to quote at scale. Fragmentation at the level of claims is compatible with consolidation at the level of risk and flow.

The formal message is narrow but powerful. A dealer holding inventory q in N claims does not face N independent liquidity problems if payoffs decompose as $X = \alpha + BF + \varepsilon$. It faces a d-dimensional factor hedge $B^\top q$ and a residual covariance exposure $q^\top \Sigma_\varepsilon q$. Residual diversification can expand capacity, but only through effective breadth and only when flow is not too toxic. Correlation, one-sided flow, sharp participation, beta error, oracle failure, haircut shocks, and capital constraints can destroy that capacity. Adverse selection remains in the residual and must be charged, reduced, or avoided. The spread condition in the paper is correspondingly modest: it is a sufficient break-even inequality for a dealer book, not a complete theory of market-microstructure equilibrium.

This is the mechanical answer to the liquidity objection. The number of claims an economy can quote is not the number of natural order books it can sustain. It is the number of claims whose common risk can be hedged, whose residual risk can be financed or diversified, whose flow toxicity can be bounded, whose information value can justify subsidy where needed, and whose operating costs can be compressed.

A Appendix A: Projection Derivation of the Factor Model

Let $F \in L^2(\mathbb{P})^d$ have mean μ_F and nonsingular covariance Ω_F. For a payoff X_i, the best linear projection coefficient is

$$\beta_i = \Omega_F^{-1} \operatorname{Cov}(F, X_i).$$

Define

$$\alpha_i = \mathbb{E}[X_i] - \beta_i^\top \mathbb{E}[F], \qquad \varepsilon_i = X_i - \alpha_i - \beta_i^\top F.$$

Then $\mathbb{E}[\varepsilon_i] = 0$ and $\operatorname{Cov}(\varepsilon_i, F) = 0$. This is the projection residual used throughout the paper. If F is a traded factor return vector with zero cost to enter at the margin, β_i is also the minimum-variance hedge loading. If trading costs or constraints matter, the optimal hedge loading may differ from the statistical projection; the difference becomes part of model or basis risk.

B Appendix B: Residual Bounds and Effective Breadth

B.1 Herfindahl form under independent residuals

If residuals are independent with common variance σ^2, then

$$q^\top \Sigma_\varepsilon q = \sigma^2 \|q\|_2^2.$$

Let $Q_1 = \|q\|_1$ and define the inventory Herfindahl index

$$H(q) = \frac{\|q\|_2^2}{\|q\|_1^2}.$$

Then

$$q^\top \Sigma_\varepsilon q = \sigma^2 Q_1^2 H(q).$$

The effective number of equally sized independent residual positions is $1/H(q)$. Equal inventory has $H(q) = 1/N$. A concentrated position has $H(q) = 1$.

B.2 Equicorrelation effective breadth

For equal weights and equicorrelation ρ, residual variance is

$$\sigma^2 Q^2 \left(\rho + \frac{1-\rho}{N}\right) = \sigma^2 \frac{Q^2}{N_{\text{eff}}},$$

where

$$N_{\text{eff}} = \frac{N}{1 + (N-1)\rho}.$$

Thus the effective breadth converges to $1/\rho$ when $\rho > 0$. A book with one thousand labels and residual correlation $\rho = 0.05$ has effective breadth of approximately 20, not 1000, for equal-weight residual risk.

C Appendix C: Expected Profit of a One-Period Quote

Let $\Pi(x;q)$ be the dealer's expected net profit from filling order xe_i around its mid. Suppressing terms unrelated to the order,

$$\begin{aligned}\Pi(x;q) = & \ |x|\, s_i(x;q) - k_i^F(x;q) - a_i(x;q,\mathcal{D}) - m_i(x;q) - o_i(x) - f_i(x;q) \\ & - \left[\mathcal{K}_R(q+xe_i) - \mathcal{K}_R(q)\right].\end{aligned}$$

With $\mathcal{K}_R(q) = \eta q^\top \Sigma_\varepsilon q/2$,

$$\mathcal{K}_R(q+xe_i) - \mathcal{K}_R(q) = \frac{\eta}{2}\left(2x(\Sigma_\varepsilon q)_i + x^2(\Sigma_\varepsilon)_{ii}\right).$$

The quote-feasibility theorem is the condition $\Pi(x;q) \geq 0$, together with hard feasibility of the post-trade book. This is a sufficient one-period break-even condition. Equilibrium spreads would require additional primitives for dealer competition, order flow, dynamic inventory choice, trader participation, and information production.

D Appendix D: Adverse-Selection Fixed Point

A simple Glosten-Milgrom-style residual selection problem illustrates the fixed point. Let δ be a specialist's residual-value advantage, symmetrically distributed with continuous density. The specialist trades size u only when expected residual profit exceeds spread and private trading cost c:

$$|\delta| > s + c.$$

Conditional on trade, the dealer's residual loss per unit at the mid is

$$\mathbb{E}[|\delta| \mid |\delta| > s + c].$$

Ignoring other costs, a competitive break-even spread solves

$$s = \mathbb{E}[|\delta| \mid |\delta| > s + c].$$

This equation may have no finite solution for sufficiently heavy-tailed residual information or sufficiently selective informed flow. In that case synthetic factor liquidity cannot create a viable quote without disclosure, participation restrictions, batch auctions, retention, collateral, or some other adverse-selection control.

E Appendix E: Convex Risk Charges

The quadratic risk charge is analytically convenient but not essential. Let

$$\mathcal{K}_R(q) = \psi(q^\top \Sigma_\varepsilon q),$$

where ψ is increasing and differentiable. Then

$$\mathcal{K}_R(q+xe_i) - \mathcal{K}_R(q) = \psi\left(q^\top \Sigma_\varepsilon q + \Delta V_i(x;q)\right) - \psi(q^\top \Sigma_\varepsilon q).$$

For small x,

$$\mathcal{K}_R(q+xe_i) - \mathcal{K}_R(q) = 2x\psi'(q^\top \Sigma_\varepsilon q)(\Sigma_\varepsilon q)_i + o(x).$$

Thus the marginal inventory term remains proportional to $(\Sigma_\varepsilon q)_i$, but the shadow price of variance becomes state-dependent.

For expected-shortfall or value-at-risk constraints, one can replace $q^\top \Sigma_\varepsilon q$ by the relevant risk measure of $q^\top \varepsilon$. Under elliptically distributed residuals, VaR and ES are proportional to $\sqrt{q^\top \Sigma_\varepsilon q}$, so the same covariance geometry remains. Under non-elliptical or jump-tailed residuals, higher moments and tail dependence become additional state variables.

F Appendix F: LMSR Details

The LMSR price impact matrix is

$$\frac{\partial p}{\partial q} = \nabla^2 C(q) = \frac{1}{b}\left(\mathrm{diag}(p) - pp^\top\right).$$

For a binary market with price p, buying Δq shares of outcome 1 changes the log odds by

$$\log \frac{p'}{1-p'} - \log \frac{p}{1-p} = \frac{\Delta q}{b}.$$

Thus b is the number of shares required to move log odds by one. Larger b means deeper liquidity and a larger subsidy. The bounded-loss result relies on mutually exclusive and exhaustive outcomes and correct settlement. Oracle failure, ambiguous resolution, or non-exhaustive outcome definitions add dispute and model costs outside the LMSR bound.

References

Abernethy, Jacob, Yiling Chen, and Jennifer Wortman Vaughan. 2013. "Efficient Market Making via Convex Optimization, and a Connection to Online Learning." *ACM Transactions on Economics and Computation* 1(2). Citation candidate; verify exact issue and article number before submission.

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84(3): 488–500.

Allen, Franklin, and Douglas Gale. 1994. *Financial Innovation and Risk Sharing.* Cambridge, MA: MIT Press.

Avellaneda, Marco, and Sasha Stoikov. 2008. "High-Frequency Trading in a Limit Order Book." *Quantitative Finance* 8(3): 217–224.

Bisin, Alberto. 1998. "General Equilibrium with Endogenously Incomplete Financial Markets." *Journal of Economic Theory* 82(1): 19–45.

Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81(3): 637–654.

Demange, Gabrielle, and Guy Laroque. 1995. "Private Information and the Design of Securities." *Journal of Economic Theory* 65(1): 233–257.

DeMarzo, Peter, and Darrell Duffie. 1999. "A Liquidity-Based Model of Security Design." *Econometrica* 67(1): 65–99.

Demsetz, Harold. 1968. "The Cost of Transacting." *Quarterly Journal of Economics* 82(1): 33–53.

Duffie, Darrell, and Rohit Rahi. 1995. "Financial Market Innovation and Security Design: An Introduction." *Journal of Economic Theory* 65(1): 1–42.

Garman, Mark B. 1976. "Market Microstructure." *Journal of Financial Economics* 3(3): 257–275.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14(1): 71–100.

Grossman, Sanford J., and Merton H. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43(3): 617–633.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70(3): 393–408.

Hanson, Robin. 2003. "Combinatorial Information Market Design." *Information Systems Frontiers* 5(1): 107–119.

Hanson, Robin. 2007. "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation." *Journal of Prediction Markets* 1(1): 3–15.

Hara, Chiaki. 1995. "Commission-Revenue Maximization in a General Equilibrium Model of Asset Creation." *Journal of Economic Theory* 65(1): 258–298.

Ho, Thomas, and Hans R. Stoll. 1981. “Optimal Dealer Pricing under Transactions and Return Uncertainty.” *Journal of Financial Economics* 9(1): 47–73.

Kyle, Albert S. 1985. “Continuous Auctions and Insider Trading.” *Econometrica* 53(6): 1315–1335.

Madhavan, Ananth. 2000. “Market Microstructure: A Survey.” *Journal of Financial Markets* 3(3): 205–258.

Merton, Robert C. 1973. “Theory of Rational Option Pricing.” *Bell Journal of Economics and Management Science* 4(1): 141–183.

O’Hara, Maureen. 1995. *Market Microstructure Theory.* Cambridge, MA: Blackwell.

Ohashi, Kazuhiko. 1995. “Endogenous Determination of the Degree of Market-Incompleteness in Futures Innovation.” *Journal of Economic Theory* 65(1): 198–217.

Othman, Abraham, Tuomas Sandholm, David M. Pennock, and Daniel M. Reeves. 2013. “A Practical Liquidity-Sensitive Automated Market Maker.” *ACM Transactions on Economics and Computation.* Citation candidate; verify exact volume, issue, and article number before submission.

Pagano, Marco. 1989. “Trading Volume and Asset Liquidity.” *Quarterly Journal of Economics* 104(2): 255–274.

Pesendorfer, Wolfgang. 1995. “Financial Innovation in a General Equilibrium Model.” *Journal of Economic Theory* 65(1): 79–116.

Ross, Stephen A. 1976. “The Arbitrage Theory of Capital Asset Pricing.” *Journal of Economic Theory* 13(3): 341–360.

Stoll, Hans R. 1978. “The Supply of Dealer Services in Securities Markets.” *Journal of Finance* 33(4): 1133–1151.