Axiomatic Market Making

Frank M. V. Feys
Independent researcher, Antwerp, Belgium
frank@fmvfeys.com
ORCID: 0009-0000-7667-6742

July 12, 2026

Abstract

We develop an axiomatic theory of market making, in which axioms on the quoting rule replace the utility, distributional, and equilibrium-concept commitments of the optimization-based tradition. Nine axioms, together with six conditions on the inventory cost, force a unique closed-form quoting rule in which the mid-quote is linear in inventory and the spread splits additively into inventory and adverse-selection components. The three economic primitives (inventory risk aversion, marginal carrying cost, and adverse-selection cost) are point-identified from three distinct moments of the rule, each from a single regression, and the latent inventory cost function is recoverable from the limit order book depth profile by integration. The spread-cost relationship is a Legendre-Fenchel conjugate pair, a sharp boundary separates functioning from frozen regimes, and the rule coincides with the unique symmetric Nash equilibrium of a Bertrand-competitive market with free entry.

JEL classification. C72 (Noncooperative Games); D47 (Market Design); D82 (Asymmetric and Private Information; Mechanism Design); G12 (Asset Pricing; Trading Volume; Bond Interest Rates); G14 (Information and Market Efficiency; Event Studies; Insider Trading).

Keywords. Market making; bid-ask spread; adverse selection; axiomatic characterization; Legendre-Fenchel duality.

Mathematics Subject Classification (2020). 90C25; 91B24; 91B26; 91G15.

Running title. Axiomatic Market Making.

1 Introduction

1.1 The Problem

A market maker's quote is among the most directly observable objects in a financial market. The primitives the quote encodes, namely her inventory cost, risk preference, and the information she faces, are conventionally recovered only through a specified optimization model. Ho and Stoll (1981) derive the spread from a dealer maximizing expected utility over a finite horizon. Avellaneda and Stoikov (2008) extend the framework to high-frequency settings, and Cartea et al. (2015) refine it through running and terminal inventory penalties.

Glosten and Milgrom (1985) obtains the spread as a Bayesian equilibrium under asymmetric information. Kyle (1985) characterizes equilibrium price impact with a strategic insider trading against competitive makers.

Each model fixes a utility function, a distribution, or an information structure, and computes the implied quoting rule. The map from primitives to spread is model-specific, and so is its inverse. What a quote reveals about costs, preferences, and information depends on which optimization problem one assumes the maker solves. Two analysts positing different problems can read the same quote as evidence for different primitives, with no test to adjudicate between them. The inventory cost function, in particular, is treated as a hidden parameter to be calibrated, with its scalar summary fitted under each chosen specification. The literature accumulates as a catalog of such maps rather than as a structural theory of the quoting rule itself.

This paper pursues an alternative route. We do not specify a utility function, a distribution, or an equilibrium concept. Instead, we impose structural restrictions directly on the quoting rule itself, and ask what functional form, and what identification, those restrictions deliver. The methodological move buys two concrete payoffs. First, the three economic primitives of market making (inventory risk aversion, marginal carrying cost, and adverse-selection cost) are point-identified from three distinct moments of the observed rule, each recovered from a single regression, with no joint estimation, no auxiliary distributional assumption, and no equilibrium computation standing between the data and the parameter. Second, the latent inventory cost function is recoverable from the limit order book depth profile by integration. It is not a free parameter of the theory but a directly observable object.

A note is in order on what is and is not committed, because two choices made along the way, the two-moment summary of beliefs and the treatment of λ as a state variable, are themselves substantive, and neither should pass as free. The axiomatic stance does not eliminate structural commitments. It relocates them. The two-moment summary of beliefs is a substantive restriction, equivalent in content to a mean-variance reduction. The treatment of λ as a state variable rather than a process commits to a steady-state reading of the trading environment. What the axiomatic route replaces is not commitment as such, but the kind of commitment. A specified utility function, a distribution, and an equilibrium concept are replaced by direct restrictions on the rule. The advantages of the relocation are precisely the two payoffs above, namely point identification of structural primitives from observable moments, and recovery of the cost function from the order book.

1.2 Main Results

Let $s = (q, \mu, \sigma^2, \lambda, \pi)$ denote the maker's state, with q inventory, μ belief about the asset value, σ^2 belief variance, λ trade arrival intensity, and π informed-trader fraction. A quoting rule is a function $Q\colon \mathcal{S} \to \mathcal{B}$, with $\mathcal{S} = \mathbb{R} \times \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} \times [0,1]$ and $\mathcal{B} = \{(b,a) \in \mathbb{R}^2 \mid b < a\}$. Nine axioms on the quoting rule, together with six environmental assumptions on the inventory cost function, imply a unique closed-form characterization (Theorem 4.1): every admissible quoting rule takes the form

$$b = \mu - \gamma_0 \sigma^2 q - \sigma^2 \Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}} \pi \Big), \qquad a = \mu - \gamma_0 \sigma^2 q + \sigma^2 \Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}} \pi \Big),$$

for a unique triple $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) \in \mathbb{R}^3_{++}$. The first is the maker's inventory risk aversion, the second her marginal cost per unit of held inventory per unit of price variance, and the third the expected loss per informed trade per unit of price variance. Preference, cost, and information are the three classical primitives of any microstructure model, and the axioms make the quoting rule depend on them through exactly one scalar each. Figure 1 shows the anatomy of the resulting quote.

The closed form is the vehicle through which the structural payoffs of the paper become available. Once the axioms are accepted, the form follows by assembly. What carries the substantive content is what the closed form enables, and we summarize the four principal payoffs in turn.

Identification. Since the axioms force the three primitives to enter through nonoverlapping channels (Theorem 6.1), each is recovered from a distinct moment of the observed rule. The risk aversion γ_0 is identified from the inventory skew of the mid-quote, the marginal inventory cost κ_{inv} from the liquidity scaling of the spread, and the adverse-selection cost κ_{adv} from the informed-flow scaling of the spread. Each primitive is recovered from a single regression, and the decoupling is structural rather than parametric. What the optimization-based tradition conflates in one or two scalars, the axiomatic theory separates into three orthogonal primitives.

Recovery and duality. Identification extends from the parameters to the entire inventory cost function. Theorem 7.3 relaxes the linearity assumption on the cost to general convexity and applies the competitive break-even axiom at each level of the order book; integrating across depth levels recovers the latent κ from the observed quotes. In Ho-Stoll and Avellaneda-Stoikov the cost coefficient is a free parameter to be fit against data, while here the cost function itself is observable in the depth profile. Theorem 8.1 reads the same relation structurally, showing that the order book traces out both members of a Legendre-Fenchel conjugate pair simultaneously.

The phase boundary. The closed form, combined with a bound on the disagreement between maker and counterparties, implies a sharp boundary in parameter space between functioning and frozen regimes. Proposition 9.1 establishes that the market generates positive trade volume if and only if

$$\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi < \frac{\delta_{\max}}{2\sigma^2},$$

where $\delta_{\max}$ is twice the divergence between the maker's and counterparties' mid-beliefs. The functioning region splits at the critical liquidity $\lambda_\times = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$ into an inventory-dominated regime, with spread $\sim 1/\lambda$, and an adverse-selection-dominated regime, with spread saturating at $2\kappa_{\text{adv}}\pi\sigma^2$. A volatility shock, informed-flow surge, or liquidity collapse can drive the market into the frozen regime. The thresholds are explicit in volatility, informed-flow fraction, and liquidity. The result is consistent with the qualitative pattern of liquidity withdrawal observed in stress episodes, though we do not undertake an empirical test here.

Competitive foundation, bijection, and structural invariants. The axioms constrain a single maker through indifference conditions, yet Theorem 12.1 recovers the same rule as the unique symmetric Nash equilibrium of a Bertrand market with free entry, showing the indifference axioms to be competitive logic in single-maker language. Theorem 5.1 establishes a canonical bijection between the nine-axiom system and $\mathbb{R}^3_{++}$, with Proposition 5.2 giving

a measurable counterexample for each axiom, and Theorem 13.1 identifies the features of the quoting rule that remain invariant as the structural primitives (κ, h, β) vary, locating Ho-Stoll, Avellaneda-Stoikov, Cartea-Jaimungal, Glosten-Milgrom, and Kyle as structural neighbors of our family.

Executive summary. For the reader who wants the bottom line before the details:

- *What is forced.* Nine axioms on the quoting rule, together with six environmental assumptions on the inventory cost, force a unique closed form in which the mid-quote is linear in inventory and the spread splits additively into an inventory component and an adverse-selection component (Theorem 4.1):

$$b = \mu - \gamma_0 \sigma^2 q - \sigma^2 \Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}} \pi \Big), \qquad a = \mu - \gamma_0 \sigma^2 q + \sigma^2 \Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}} \pi \Big).$$

- *What is free.* A triple $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) \in \mathbb{R}^3_{++}$, one scalar for preference (inventory risk aversion), one for cost (marginal carrying cost per unit variance), and one for information (expected loss per informed trade per unit variance). The nine-axiom system stands in canonical bijection with this parameter space (Theorem 5.1), and each axiom is independent, with a measurable counterexample supplied for each (Proposition 5.2).

- *What is identified.* Each of the three primitives is recovered from a distinct moment of the observed rule. The risk aversion γ_0 is identified from the inventory skew of the mid, the marginal inventory cost κ_{inv} from the liquidity scaling of the spread, and the adverse-selection cost κ_{adv} from the informed-flow scaling of the spread. Each requires a single regression, with no joint estimation, no distributional assumption, and no equilibrium computation standing between the data and the parameter.

- *What is recoverable.* The entire latent inventory cost function, not merely a scalar summary, is recoverable from the limit order book depth profile by integration (Theorem 7.3). The cost function is not a free parameter of the theory but a directly observable object.

- *What is dual.* Spread and cost stand in Legendre-Fenchel conjugacy, with the half-spread the convex conjugate of the marginal inventory cost (Theorem 8.1). The limit order book traces out both members of the conjugate pair simultaneously.

- *Where the market freezes.* The market generates positive trade volume if and only if

$$\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}} \pi < \frac{\delta_{\max}}{2\sigma^2},$$

  with $\delta_{\max}$ twice the belief divergence between the maker and her counterparties (Proposition 9.1). Inside the functioning region, a critical liquidity $\lambda_\times = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$ separates the inventory-dominated regime, in which the spread scales as $1/\lambda$, from the adverse-selection-dominated regime, in which the spread saturates at $2\kappa_{\text{adv}}\pi\sigma^2$.

- *What is not covered.* Queue position, latency, order-flow toxicity dynamics, and learning about λ or π over time are outside the framework. The treatment of λ and π as state variables commits the analysis to a steady-state reading of the trading environment, and departures from that reading are a separate object of study.

- *A companion note.* The dual axiomatization, taking the maker's preferences over inventory trajectories as the primitive instead, is developed in the companion paper (Feys, 2026). The two axiomatizations operate at different structural levels and are joined by the Legendre-Fenchel duality of Section 8, which translates between the cost-side and price-side representations.

- *What is competitive.* The axioms are stated on a single maker through indifference conditions alone, yet the same closed-form rule arises as the unique symmetric Nash equilibrium of a Bertrand-competitive market with free entry (Theorem 12.1). The indifference axioms are competitive logic in single-maker language.

1.3 Relation to the Literature

The setup is in the tradition of microstructure models combining inventory effects with adverse selection. The inventory-cost lineage runs through Garman (1976), Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981), Avellaneda and Stoikov (2008), and Guéant et al. (2013). The adverse-selection lineage runs through Glosten and Milgrom (1985), Kyle (1985), Back (1992), Easley and O'Hara (1987), and Easley et al. (1996). The closed-form rule contains the inventory and adverse-selection components separately and additively, in the structural-decomposition tradition of Madhavan and Smidt (1993) and Huang and Stoll (1997). Each of Ho-Stoll, Avellaneda-Stoikov, and Glosten-Milgrom can be recovered from our characterization under additional structural commitments. Our contribution is to derive the closed form from the axioms, identify the three parameters as the unique parameter space, and characterize the structural properties (separability, duality, phase transition) that follow.

Axiomatic methods have a long lineage in economics: the social-choice tradition from Arrow (1951) to Sen (1970); the bargaining tradition from Nash (1950) to Kalai and Smorodinsky (1975) and Roth (1979); the mechanism-design tradition from Vickrey (1961) and Hurwicz (1973) through Myerson (1981) and Maskin (2008); and the risk-measurement tradition in Artzner et al. (1999), with decision-theoretic foundations in von Neumann and Morgenstern (1944) and Savage (1954). None of these traditions has, to our knowledge, been applied to the bid-ask spread. The closest precedent is the axiomatic theory of coherent risk measures and its convex extensions (Föllmer and Schied (2002), Frittelli and Rosazza Gianin (2002), Cheridito et al. (2006), Detlefsen and Scandolo (2005)), with our Legendre-Fenchel duality (Section 8) making the connection precise through the convex conjugacy of the inventory cost and the half-spread. The microstructure literature has been overwhelmingly optimization-based, with canonical foundations in O'Hara (1995) and modern syntheses in Foucault et al. (2013) and Hasbrouck (2007). The over-the-counter strand of Duffie et al. (2005) and Weill (2020) derives spreads from search-theoretic bargaining, with structural commitments distinct from ours.

A companion paper (Feys, 2026) develops the dual axiomatization on the objective side, taking the maker's preferences over inventory trajectories as the primitive and deriving the entropic certainty-equivalent objective with running coefficient $\phi = \gamma\sigma^2/2$, unifying Avellaneda-Stoikov and Cartea-Jaimungal as two presentations of a single object. The two axiomatizations are connected by the Legendre-Fenchel duality of Section 8, translating between cost-side and price-side representations.

1.4 Outline

Section 2 fixes notation and primitives. Section 3 presents the nine axioms. Section 4 states and proves the Full Characterization Theorem. Section 5 establishes the axiom-system–parameter-space bijection and irredundancy. Section 6 states the Structural Separability Theorem and the identification corollary. Section 7 extends to the full limit order book and proves the Recovery Theorem. Section 8 establishes the Legendre-Fenchel duality. Section 9 characterizes the phase boundary and the three liquidity regimes. Section 10 collects comparative statics and an empirical-implementation roadmap. Section 11 states the skew-spread ratio result. Section 12 establishes the Bertrand equivalence. Section 13 states the Invariant Features Theorem and locates the standard models. Section 14 relates the work to the axiomatic tradition. Section 15 concludes.

2 Setup

Throughout the paper, $\mathbb{R}$ denotes the real line, $\mathbb{R}_+ = [0, \infty)$ the nonnegative reals, and $\mathbb{R}_{++} = (0, \infty)$ the strictly positive reals.

2.1 The State Space

The maker quotes a single asset. Her state has five variables: inventory $q \in \mathbb{R}$ (positive means long), belief $\mu \in \mathbb{R}$, belief variance $\sigma^2 \in \mathbb{R}_{++}$, trade arrival intensity $\lambda \in \mathbb{R}_{++}$ (with $1/\lambda$ the mean inter-trade time), and informed fraction $\pi \in [0, 1]$. The state space is

$$\mathcal{S} = \mathbb{R} \times \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} \times [0, 1], \qquad s = (q, \mu, \sigma^2, \lambda, \pi).$$

Remark 2.1 (On the Interpretation of μ and σ^2)**.** The axioms below operate on the function $Q\colon \mathcal{S} \to \mathcal{B}$ directly, treating $(\mu, \sigma^2) \in \mathbb{R} \times \mathbb{R}_{++}$ as scalar inputs in the state space. A reader who prefers to ground the informal language of *public information* measure-theoretically may fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ carrying an asset value v and a public-information filtration $\mathcal{F}^{\text{pub}}$, and set $\mu = \mathbb{E}[v \mid \mathcal{F}^{\text{pub}}]$, $\sigma^2 = \text{Var}(v \mid \mathcal{F}^{\text{pub}})$. No theorem in the paper invokes such a construction. The substantive commitment is the two-moment summary, not any measure-theoretic apparatus underlying it.

2.2 The Quoting Rule

The space of valid bid-ask pairs is $\mathcal{B} = \{(b, a) \in \mathbb{R}^2 \mid b < a\}$.

Definition 2.2. A *quoting rule* is a Borel-measurable function

$$Q\colon \mathcal{S} \to \mathcal{B}, \qquad Q(s) = (b(s), a(s)),$$

where $b(s)$ and $a(s)$ denote the bid and ask at state s. The *mid-quote* is $m(s) = \frac{1}{2}(b(s)+a(s))$ and the *spread* is $\delta(s) = a(s) - b(s)$.

Note that the two parametrizations (b, a) and (m, δ) carry the same information. The map $(b, a) \mapsto ((b + a)/2,\ a - b)$ is a bijection from $\mathcal{B}$ onto $\mathbb{R} \times \mathbb{R}_{++}$. We use whichever is more convenient.

The measurability requirement is the standard regularity condition for functions of the state and rules out pathological constructions such as Hamel-basis solutions to Cauchy's functional equation. Every quoting rule arising in microstructure modeling, in optimization-based theories, or in observed data satisfies measurability by construction. We equip $\mathcal{S}$ with the product topology and $\mathcal{B} \subset \mathbb{R}^2$ with the subspace topology. Borel measurability is taken with respect to the induced Borel σ-algebras.

2.3 The Trading Environment

We adopt three primitives of the trading environment in addition to the maker's state.

Informed and uninformed traders. Trades arrive at total rate λ, with a fraction π informed. Informed traders observe the true asset value v. The remaining $1 - \pi$ are uninformed. Both populations trade whenever profitable at the posted quotes.

Belief structure. The maker's belief about the asset value is summarized by the mean μ and variance σ^2, with no further parametric distributional commitment. The underlying belief may be Gaussian, Student-t, or otherwise, provided the first two moments are well defined. The two-moment reduction is the substantive content of the setup, in the spirit of mean-variance microstructure. Alongside the steady-state treatment of λ, it is the principal structural commitment relocated by the axiomatic stance.

Remark 2.3 (Heavy-Tailed Beliefs and the Two-Moment Scope)**.** The rule takes (μ, σ^2) as inputs, and any projection of a full belief distribution onto its two-moment summary is exogenous to the axiom system. Under a Student-t_ν prior with $\nu > 2$, the variance is finite and the rule is well defined on the projected summary; the information discarded is the higher-moment content, not any part of what the axioms constrain. Under $\nu \leq 2$, the setup fails at the input stage rather than at the axiom stage, since σ^2 is undefined, and the scope of the theory is the admissible-input range rather than a modeling defect within it. A practitioner preferring a distribution-aware rule may replace the two-moment reduction with a richer summary, at the cost of losing the closed-form point identification of Theorem 4.1.

Disagreement bound. Each counterparty holds her own belief about the underlying asset value. We write $\delta_{\max} > 0$ for the maximum admissible spread, defined as

$$\delta_{\max} := 2 \max_{C} |\mu_C - \mu|,$$

twice the maximum belief disagreement between the maker and any counterparty. The factor of two reflects that a trade at the ask requires $\mu_C > a$, equivalent to $\mu_C - \mu > \delta/2$. A spread is admissible if and only if $\delta < \delta_{\max}$. The quantity $\delta_{\max}$ plays no role in the axioms but is the key input to the phase-transition analysis of Section 9.

Inventory cost. The maker bears a cost of holding inventory of the form $C(q, \sigma^2) = \kappa(q)\sigma^2$, where $\kappa\colon \mathbb{R} \to \mathbb{R}_+$ is convex, nondecreasing in $|q|$, with $\kappa(0) = 0$, and symmetric in the sign of inventory. The multiplicative variance dependence reflects the dollar variance of the held position, while κ encodes the maker's cost structure beyond pure variance.

2.4 Environmental Assumptions on the Cost Function

We collect six assumptions on the cost function as environmental primitives, distinct from the axioms on the quoting rule.

(C1, origin)	$\kappa(0) = 0$.
(C2, symmetry)	κ is symmetric, so $\kappa(-q) = \kappa(q)$.
(C3, monotonicity)	κ is nondecreasing in $\lvert q\rvert$.
(C4, autonomy)	κ does not depend on λ or π.
(C5, linearity)	κ is linear in $\lvert q\rvert$, that is, $\kappa(q) = \kappa_{\text{inv}}\lvert q\rvert$ for some $\kappa_{\text{inv}} > 0$.
(C6, convexity)	κ is convex on $\mathbb{R}$ (so one-sided derivatives exist and are finite at every interior point).

C1, C2, C3 are minimal regularity. C6 (convexity) is implied by C5 and is stated separately so that the order-book extension can rely on C6 explicitly when C5 is dropped. C4 and C5 are substantive structural commitments.

Condition C4 mirrors (A5), (A6) on the cost side: neither κ nor the skew depends on λ or π. Weakening C4 admits cost structures in which κ depends on the trading environment, with Avellaneda-Stoikov as the canonical example. We list C4 on the cost side rather than as an axiom because κ is an environmental primitive and the axioms are reserved for the maker's quoting behavior, following the standard division in the axiomatic tradition. There, Arrow (1951) constrains the social welfare function, not preferences, and Artzner et al. (1999) constrains the risk measure, not the probability space.

Condition C5 is the linearity assumption. Section 7 relaxes C5 to general convexity and recovers κ from the order book. Under C5, κ collapses to the scalar κ_{inv}. Under (A4) the half-spread equals $\sigma^2\kappa'_+(|q|)/\lambda$, and under C5 this is constant in q, the simplest nontrivial baseline. The quadratic case is covered by the convex-cost extension, where the Recovery Theorem identifies the full κ.

2.5 Comments on the Setup

The five state variables $(q, \mu, \sigma^2, \lambda, \pi)$ form the canonical decomposition of the maker's state in market microstructure. Inventory q captures her current position, μ and σ^2 summarize

her view of the asset value, λ measures market activity, and π measures the information environment. The axioms in Section 3 impose structural restrictions on how the quoting rule depends on each.

3 The Axiom System

We now present the axioms on the quoting rule. There are nine, organized along two complementary lines.

Conceptually, the axioms fall into four groups by the kind of normative claim each makes: *definitional axioms* (A1, A7) fix what it means to play the market-making role; *inventory-channel axioms* (A2, A5, A6) describe how the maker manages her own position; a single *risk-pricing structural axiom* (A3) selects the denomination in which risk enters the quoting rule; and *competitive break-even axioms* (A4, A8, A9) record what steady-state competitive operation forces on the maker. The competitive content of the break-even axioms admits independent justification. Theorem 12.1 derives the same closed-form rule as the unique symmetric Nash equilibrium of a Bertrand market with free entry, and the axioms are natural under any of three weakly competitive conditions present in real markets: counterparty access to alternative liquidity, the maker as one of many small contributors to a consolidated book, or steady-state solvency in a repeated-interaction market where positive margins attract entry and negative margins force exit.

Methodologically, the axioms partition into three nested layers by the kind of denial each admits: an *indispensable core* (A1, A4, A7, A8, A9) whose denial changes the object of study; a *structural choice* axiom (A3) whose denial yields neighbors such as the σ-scaling Glosten-Milgrom family; and three *modularity extensions* (A2, A5, A6) yielding qualitatively distinct quoting rules within the same framework. The layered partition identifies which subset suffices for each conclusion: Theorem 4.1 requires the full theory; the quantitative spread formula and phase boundary require the core plus the structural choice; additive decomposition and qualitative skew require only the core. Section 14.1 traces the layered structure in detail, and irredundancy is established formally in Section 5.

The axioms are stated below in numerical order, each leading with the formal mathematical condition followed by an interpretation paragraph.

3.1 The Nine Axioms

(A1) Inventory aversion. The mid-quote is strictly decreasing in q. That is, for every $s = (q, \mu, \sigma^2, \lambda, \pi)$ and every $q' > q$,

$$m(q', \mu, \sigma^2, \lambda, \pi) < m(q, \mu, \sigma^2, \lambda, \pi).$$

Interpretation. The axiom captures the qualitative content of inventory aversion. A market maker is paid to provide liquidity rather than to take directional positions, and she prefers to keep inventory close to neutral. When holding large positive inventory, she lowers her mid-quote to discourage further buys and encourage sells, reducing exposure. A trader who preferred more inventory to less would be a directional speculator, not a market maker.

(A2) Translation-invariant inventory attitude. A function $\gamma\colon \mathbb{R}_{++}\times\mathbb{R}_{++}\times[0,1]\to\mathbb{R}$ exists such that for every state $s=(q,\mu,\sigma^2,\lambda,\pi)\in\mathcal{S}$,

$$m(s)-\mu=-\gamma(\sigma^2,\lambda,\pi)\,q.$$

Interpretation. Two principles are at work. First, the per-unit adjustment is a property of who the maker is, not where her position sits. A maker who shifts her mid-quote by Δ when inventory moves from 0 to 1 shifts it by the same Δ when inventory moves from 7 to 8 or from -3 to -2. Second, the attitude responds to position, not absolute price level. A shift in μ moves quotes uniformly and leaves the marginal response unchanged. Both principles have direct economic grounding. A market maker operates a *flow* business, and under CARA utility over Gaussian P&L the marginal cost of holding one additional unit of inventory is constant in the inventory level (exponential utility exhibits no wealth effects, the standard foundation for mean-variance analysis). Translation invariance in μ is the price-theoretic counterpart. Inventory risk depends on position and σ^2, not on μ. A maker whose marginal attitude varied with position would have a family of local attitudes. A maker whose attitude depended on μ would adjust inventory pricing in response to nominal price levels, with no bearing on the risk she faces.

We have not included a separate continuity axiom. Linearity in (A2) delivers continuity in q for free, and continuity in each state variable is established as a derived property in Theorem 4.11.

(A3) Variance-proportional risk pricing. There exist functions $f\colon\mathbb{R}\times\mathbb{R}_{++}\times[0,1]\to\mathbb{R}$ and $g\colon\mathbb{R}\times\mathbb{R}_{++}\times[0,1]\to\mathbb{R}_{++}$ such that for every state $s\in\mathcal{S}$,

$$m(s)-\mu=\sigma^2\cdot f(q,\lambda,\pi),\qquad \delta(s)=\sigma^2\cdot g(q,\lambda,\pi).$$

Interpretation. The maker's inventory exposes her to P&L variation over the holding period. Under mean-variance preferences, the cost of carrying one unit per unit time is proportional to the variance σ^2 of the price process, not the standard deviation σ. The quoting rule inherits this σ^2-scaling from the cost it prices, and the same scaling extends to the adverse-selection portion, where information rents are computed against the variance-denominated belief. The mean-variance commitment has decision-theoretic grounding. Under CARA utility with risk-aversion γ_0 and Gaussian P&L, the certainty-equivalent cost of holding inventory q for one unit of time is $\gamma_0\sigma^2q^2/2$, manifestly proportional to σ^2, and the cost per unit time inherits this scaling because the P&L over dt is Gaussian with variance σ^2dt. The contrast with the σ-scaling of Glosten-Milgrom makes the choice explicit. There, the informed trader's profit per trade scales as $\mathbb{E}[|v-\mu|]$, linear in σ under symmetric distributions. Here, information rents are computed against the same variance-denominated belief that governs inventory risk, so both skew and spread price risk in σ^2 units. Section 13 places both scalings within a common axiomatic family parametrized by the exponent.

(A3s) Mixed-scaling alternative. The Glosten-Milgrom σ-scaling alternative replaces uniform σ^2-scaling by a two-component scaling. There exist functions $f, g_{\text{inv}}, g_{\text{adv}}$ and $\beta\in(0,1]$ such that

$$m(s)-\mu=\sigma^2 f(q,\lambda,\pi),\qquad \delta(s)=\sigma^2 g_{\text{inv}}(q,\lambda,\pi)+\sigma^{2\beta}g_{\text{adv}}(q,\lambda,\pi).$$

The Glosten-Milgrom rule corresponds to $\beta = 1/2$. We treat (A3s) as the alternative structural primitive discussed in Section 13.

(A4) Competitive break-even on carry. For each state $s = (q, \mu, \sigma^2, \lambda, \pi) \in \mathcal{S}$,

$$\frac{\delta(s)|_{\pi=0}}{2} = \frac{\sigma^2 \kappa'_+(|q|)}{\lambda},$$

where κ'_+ denotes the right derivative of κ (finite on $\mathbb{R}_+$ by C6).

Interpretation. The maker incurs a real economic cost for holding inventory. Capital is tied up, exposure to price movements must be hedged or borne, and balance-sheet capacity is consumed. She must recover this cost through her quoting revenue or lose money on every trade. At the same time, any margin she charges above the marginal carrying cost is a margin a competitor can undercut. In competitive equilibrium the inventory portion of the half-spread settles at the marginal carrying cost. Above it invites entry, and below it forces exit. The marginal cost per unit time is $\sigma^2 \kappa'_+(|q|)$, and the trade-arrival timescale $1/\lambda$ converts this into a per-trade cost. Under linear cost (C5), the per-trade cost equals $\kappa_{\text{inv}} \sigma^2 / \lambda$ exactly.

(A5) Flow-independent inventory pricing. There exists $\tilde{f} \colon \mathbb{R} \times \mathbb{R}_{++} \times [0,1] \to \mathbb{R}$ such that for every state $s \in \mathcal{S}$,

$$m(s) - \mu = \tilde{f}(q, \sigma^2, \pi).$$

Interpretation. A maker who has accumulated inventory q faces the same risk exposure on it whether the next trade arrives in one second or one minute. Her skew, which prices that exposure, is therefore a function of position and price environment alone. The arrival intensity λ enters the quoting rule through the per-trade carrying cost in (A4), but separately from the skew. The skew measures how much the maker adjusts per unit of inventory, a statement about her preferences over holding risk, while λ governs how quickly she expects to turn over her position, captured by the $1/\lambda$ timescale in the spread. Conflating the two would mean risk aversion depends on how busy the market happens to be, which has no foundation in preferences over inventory risk. The separation has a sharp identification payoff. The intensity λ enters the spread but not the skew, so κ_{inv} is identified from the liquidity scaling of the spread rather than from the inventory slope of the mid-quote.

(A6) Informed-flow-independent inventory pricing. There is $\hat{f} \colon \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} \to \mathbb{R}$ such that for every state $s \in \mathcal{S}$,

$$m(s) - \mu = \hat{f}(q, \sigma^2, \lambda).$$

Interpretation. A maker who holds inventory q faces the same price risk on it whether the next trader is informed or uninformed. The counterparty's informational character enters through the spread (the expected information rent per trade), not the inventory pricing. Position risk is the variance of unrealized P&L attributable to inventory, determined by q and σ^2. Adverse-selection risk is the expected loss per trade due to informational advantage. The skew responds to a stock (inventory held), the spread to a flow (the counterparty

population), and a given unit of inventory generates the same dollar variance regardless of its originating trade. A sophisticated maker might in principle use her quotes to signal beliefs about counterparty type, and (A6) rules out such signaling in the inventory response. As with (A5), the identification payoff is sharp. The informed fraction π enters the spread but not the skew, so κ_{adv} is identified from the informed-flow scaling.

(A7) Spread monotonicity in information. The spread $\delta(s)$ is strictly increasing in π for $\pi \in [0, 1]$.

Interpretation. The axiom captures the qualitative content of adverse selection. The maker quotes on public information summarized by μ. An informed trader, who has observed the true value v, buys at the ask precisely when $v > a$ and sells at the bid precisely when $v < b$. Every informed trade is therefore a loss in expectation. The maker pays an information rent $\mathbb{E}[v - a \mid v > a] > 0$ on every informed buy and $\mathbb{E}[b - v \mid v < b] > 0$ on every informed sell. A higher fraction π of informed flow means more of these losing trades per unit time, and a maker who failed to widen her spread in response would bleed capital at a rate increasing in π. The spread widening is the structural mechanism by which the maker remains solvent in the presence of better-informed counterparties. A participant who did not widen in response to informed flow would not survive as a market maker. The qualitative direction is fixed here, and the quantitative form is determined by (A8).

(A8) Cross-subsidization break-even. Define the adverse-selection increment

$$h(q, \mu, \sigma^2, \lambda, \pi) := \delta(q, \mu, \sigma^2, \lambda, \pi) - \delta(q, \mu, \sigma^2, \lambda, 0).$$

There exists a function $\tilde{h} \colon \mathbb{R}_{++} \times [0, 1] \to \mathbb{R}_+$ such that $h(q, \mu, \sigma^2, \lambda, \pi) = \tilde{h}(\sigma^2, \pi)$ for every state $s \in \mathcal{S}$, matching the expected information rent per trade.

Interpretation. The cross-subsidization mechanism is the structural heart of adverse-selection-based market making. Every trade the maker executes is priced at the same spread, but only a fraction π of trades cost her the information rent of (A7). The spread she collects from *uninformed* flow must finance the expected loss she incurs on *informed* flow, with the spread set at the level that exactly balances the two over the population. Competitive entry forces the balance. A wider spread invites entry by a competitor willing to undercut on the adverse-selection channel, and a narrower spread leaves the maker losing capital to informed traders at a rate she cannot sustain. The widening is a property of the informational environment alone, namely the price volatility on which information rents are computed and the informed fraction of the flow. Inventory and arrival rate enter the maker's economics through the carrying-cost channel of (A4), not this one. We write $h(\sigma^2, \pi)$ for $\tilde{h}(\sigma^2, \pi)$ where no confusion arises.

(A9) Non-interaction among informed traders. For every (σ^2, π) and decomposition $\pi = \pi_1 + \pi_2$ with $\pi_i \geq 0$ and $\pi_1 + \pi_2 \leq 1$,

$$h(\sigma^2, \pi) = h(\sigma^2, \pi_1) + h(\sigma^2, \pi_2).$$

Interpretation. The economic claim is that informed traders do not crowd one another out, so that splitting the informed population into any two subgroups leaves each subgroup's rent contribution exactly as it would be on its own. The classical model of informed trading takes each informed participant to have her own private signal about the asset value, drawn independently across traders. Two informed traders acting on independent signals each extract their full information rent on their respective trades. Neither's expected profit per trade is reduced by the other's presence, and the aggregate rent the maker pays is the sum of individual rents. This is the structural environment underlying the Glosten-Milgrom tradition. Information is dispersed, each informed participant trades on her own piece, and the maker faces a population of independent extractors rather than a coordinated group. The contrast clarifies the content. When informed traders share a common signal and trade strategically, as in Kyle, each additional informed participant competes for the same finite rent and the aggregate impact is sublinear in the population size. In that case h becomes concave in π rather than linear, capturing the rent dissipation among strategic informed traders. The present axiom commits to the dispersed-information case.

The combined effect of (A7), (A8), (A9) with measurability and (A3) is to make h linear in π with σ^2-coefficient, $h(\sigma^2, \pi) = 2\kappa_{\text{adv}}\pi\sigma^2$, as established in Lemma 4.6.

3.2 Summary of the Axiom System

Table 1 summarizes the four conceptual groups. The definitional axioms (A1), (A7) record the qualitative monotonicity content of the market-making role. The inventory-channel axioms (A2), (A5), (A6) describe the structure of the maker's response to her own position. The risk-pricing axiom (A3) selects the variance denomination of the family. The competitive break-even axioms (A4), (A8), (A9) record what steady-state competitive operation forces on each channel. Translation equivariance in μ, oddness of the skew in q, and continuity in each state variable emerge as derived theorems in Section 4.1.

4 The Full Characterization

Together, the nine axioms reduce the quoting rule to a three-parameter family. The axioms are not minimal in the Cauchy-additivity sense. Once the commitments in (A2), (A4), (A8), and (A9) are in place, the closed form follows largely by assembly, and the substantive deliverables are what the closed form makes possible, namely the additive spread decomposition, the three-parameter bijection, the Legendre-Fenchel duality reading, and the sharp phase boundary. The closed-form rule established here is the bridge between the axiom system and those structural payoffs. Identification (Theorem 6.1, Corollary 6.3), recovery of the latent cost function (Theorem 7.3), Legendre-Fenchel duality (Theorem 8.1), and the phase transition with explicit thresholds (Proposition 9.1) all rest on the form derived below.

Axiom	Layer	Content	Role
(A1)	I	$\partial m/\partial q < 0$	Inventory aversion
(A7)	I	$\partial \delta/\partial \pi > 0$	Adverse-selection monotonicity
(A2)	III	Skew linear in q, γ independent of μ	Single inventory-risk attitude with translation invariance
(A5)	III	Skew independent of λ	Inventory-pricing autonomy under flow
(A6)	III	Skew independent of π	Inventory-pricing autonomy under information
(A3)	II	Q homogeneous in σ^2	Variance-proportional risk pricing
(A4)	I	Inventory half-spread equals marginal carrying cost	Competitive break-even on inventory
(A8)	I	Adverse-selection increment matches information rent	Competitive break-even on information rent
(A9)	I	Adverse-selection increment additive in π	Non-interaction among informed traders

Table 1: The nine axioms, with conceptual role and three-layer assignment. Horizontal rules separate the four conceptual groups. Layer I (five core axioms) is forced by the market-making role and competitive operation. Layer II selects the variance-denominated family. Layer III records the inventory channel's autonomy from the trading-environment variables. Translation equivariance, skew oddness, and continuity are derived theorems (Section 4.1). Strict spread positivity follows from (A4) and the cost structure, and is built into $\mathcal{B}$.

Theorem 4.1 (Full Characterization)**.** *A measurable quoting rule $Q\colon \mathcal{S} \to \mathcal{B}$ satisfies (A1) through (A9) with environmental assumptions C1 through C6 if and only if there exist constants $\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}} \in \mathbb{R}_{++}$ such that*

$$b(s) = \mu - \gamma_0\sigma^2 q - \sigma^2\Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi\Big), \quad a(s) = \mu - \gamma_0\sigma^2 q + \sigma^2\Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi\Big). \tag{1}$$

Figure 1 shows the anatomy of the characterized quote.

Example 4.2 (The Quote in Numbers)**.** Fix $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) = (0.5, 0.2, 0.1)$ and consider a maker at the state $(q, \mu, \sigma^2, \lambda, \pi) = (5, 100, 0.5, 1, 0.2)$. The characterization (1) produces an inventory skew $\gamma_0\sigma^2 q = 1.25$, an inventory half-spread component $\sigma^2\kappa_{\text{inv}}/\lambda = 0.10$, and an adverse-selection half-spread component $\sigma^2\kappa_{\text{adv}}\pi = 0.01$. The mid-quote is therefore $m = \mu - 1.25 = 98.75$, the total half-spread is 0.11, and the quote is $(b, a) = (98.64, 98.86)$. Since the maker holds five units, the skew shifts both bid and ask by 1.25 below belief, attracting buying counterparties and reducing her long position. In this state the inventory piece dominates the adverse-selection piece ten to one, though the ratio depends on the arrival rate. Doubling λ to 2 halves the inventory component to 0.05 while leaving the adverse-selection component fixed at 0.01, an instance of the $1/\lambda$ scaling that determines the phase boundary of Section 9.

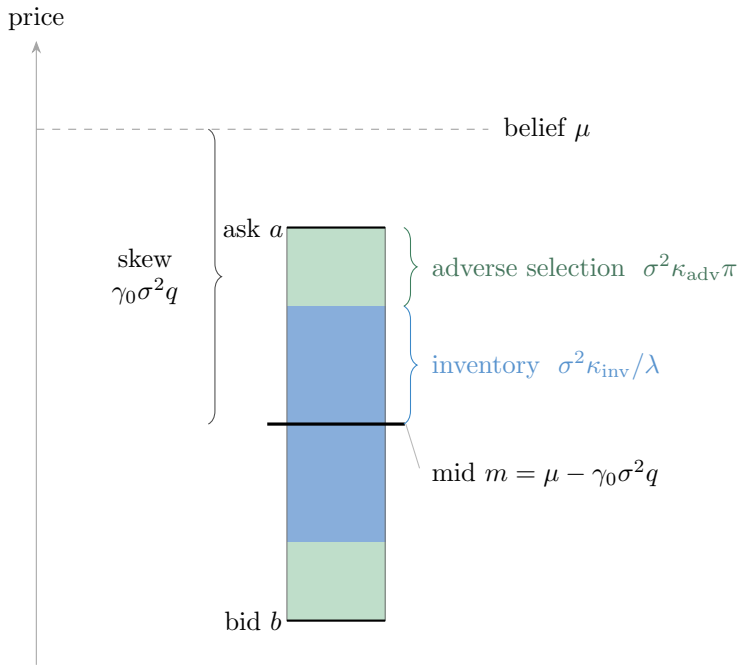


Figure 1: Anatomy of the characterized quote (1) for a maker at positive inventory $q > 0$. The mid-quote $m = \mu - \gamma_0\sigma^2 q$ sits below the belief μ by the inventory skew $\gamma_0\sigma^2 q$, with bid and ask one half-spread on either side. The half-spread decomposes additively into an inventory component $\sigma^2\kappa_{\text{inv}}/\lambda$ (blue) and an adverse-selection component $\sigma^2\kappa_{\text{adv}}\pi$ (green). Each of the three parameters occupies a geometrically separate role, namely γ_0 the skew, κ_{inv} the inventory band, and κ_{adv} the adverse-selection band. The structural separability of Theorem 6.1 is visible directly in the quote. The color convention matches Figure 3.

The proof has two directions. Sufficiency, that the rule (1) satisfies each axiom, is a direct verification. The substantive content is the necessity direction: the axioms force the form (1). We prove necessity through three lemmas, each isolating a structural piece of the derivation.

Lemma 4.3 (The Skew)**.** *Under (A1), (A2), (A3), (A5), and (A6), there exists $\gamma_0 > 0$ such that*

$$m(s) - \mu = -\gamma_0\sigma^2 q.$$

Proof. By (A2), the skew has the form $m(s) - \mu = -\gamma(\sigma^2, \lambda, \pi) \cdot q$ for some function $\gamma\colon \mathbb{R}_{++} \times \mathbb{R}_{++} \times [0,1] \to \mathbb{R}$ that does not depend on μ. By (A5), γ does not depend on λ. By (A6), γ does not depend on π. Thus, $\gamma = \gamma(\sigma^2)$.

By (A3), the skew is homogeneous of degree one in σ^2, so $\gamma(t\sigma^2) \cdot q = t\,\gamma(\sigma^2) \cdot q$ for every $t > 0$ and every q, hence $\gamma(\sigma^2) = \sigma^2 \cdot \gamma_0$ for a constant $\gamma_0 \in \mathbb{R}$.

Finally, by (A1), $\partial m/\partial q < 0$, which gives $-\gamma_0\sigma^2 < 0$, hence $\gamma_0 > 0$. □

Remark 4.4 (Linearity from (A2) versus Cauchy plus Measurability)**.** Axiom (A2) directly asserts $\mathbb{R}$-linearity of the skew in q, not merely Cauchy additivity, so no measurability-plus-Cauchy argument is required in Lemma 4.3. The measurability-plus-monotonicity argument

enters substantively only in Lemma 4.6, where the adverse-selection function on $[0,1]$ is treated by combining rational additivity from (A9) with monotonicity from (A7).

Lemma 4.5 (The Inventory Spread)**.** *Under (A4) and the environmental assumptions C1 through C6,*

$$\delta(s)|_{\pi=0} = \frac{2\kappa_{\text{inv}}\sigma^2}{\lambda}.$$

Proof. By (A4) at $\pi = 0$, $\delta(s)|_{\pi=0}/2 = \sigma^2\,\kappa'_+(|q|)/\lambda$. Under C5, $\kappa(q) = \kappa_{\text{inv}}|q|$ for some $\kappa_{\text{inv}} > 0$, and by C2 the right derivative $\kappa'_+(|q|) = \kappa_{\text{inv}}$ for every $q \in \mathbb{R}$. Substituting gives $\delta(s)|_{\pi=0} = 2\kappa_{\text{inv}}\sigma^2/\lambda$. Section 7 relaxes C5 to general convexity and obtains a richer order-book structure. □

Lemma 4.6 (The Adverse Selection Spread)**.** *Under axioms (A3), (A7), (A8), (A9), and measurability,* $h(\sigma^2,\pi) := \delta(q,\mu,\sigma^2,\lambda,\pi) - \delta(q,\mu,\sigma^2,\lambda,0) = 2\kappa_{\text{adv}}\pi\sigma^2$ *for some* $\kappa_{\text{adv}} > 0$.

Proof. By (A8), h depends only on (σ^2,π). By (A3), h is homogeneous of degree one in σ^2, so $h(\sigma^2,\pi) = \sigma^2\,\tilde{h}(\pi)$ for some $\tilde{h}\colon [0,1] \to \mathbb{R}$. The measurability of Q gives the measurability of $\tilde{h}$ since $\tilde{h}(\pi) = h(\sigma_0^2,\pi)/\sigma_0^2$ for any fixed $\sigma_0^2 > 0$.

By (A9), $\tilde{h}$ satisfies the restricted Cauchy equation

$$\tilde{h}(\pi_1+\pi_2) = \tilde{h}(\pi_1) + \tilde{h}(\pi_2) \quad \text{for } \pi_1,\pi_2 \geq 0,\ \pi_1+\pi_2 \leq 1.$$

Setting $\pi_1 = \pi_2 = 0$ gives $\tilde{h}(0) = 0$. By induction, $\tilde{h}(n\pi) = n\tilde{h}(\pi)$ whenever $n\pi \in [0,1]$, and substituting $\pi = 1/n$ yields $\tilde{h}(1/n) = \tilde{h}(1)/n$. Combining gives $\tilde{h}(r) = r\tilde{h}(1)$ for every rational $r \in [0,1]$.

By (A7), $\tilde{h}$ is strictly increasing on $[0,1]$. Set $c := \tilde{h}(1)$. Suppose $\tilde{h}(\pi_0) > c\pi_0$ for some $\pi_0 \in (0,1)$. Choose rational $r > \pi_0$ close enough to π_0 that $cr < \tilde{h}(\pi_0)$. Then we have $\tilde{h}(r) > \tilde{h}(\pi_0) > cr$ by monotonicity, contradicting $\tilde{h}(r) = cr$. The case $\tilde{h}(\pi_0) < c\pi_0$ yields the symmetric contradiction with rational $r < \pi_0$. Hence, $\tilde{h}(\pi) = c\pi$ on $[0,1]$. By (A7), $c = \tilde{h}(1) > 0$, and setting $\kappa_{\text{adv}} := c/2$ gives $h(\sigma^2,\pi) = 2\kappa_{\text{adv}}\pi\sigma^2$. □

Remark 4.7 (The Role of Measurability and Monotonicity)**.** The classical Cauchy functional equation $f(x+y) = f(x) + f(y)$ on $\mathbb{R}$ admits, without regularity, uncountably many nonmeasurable solutions constructed via Hamel bases, see Aczél (1966). The proof of Lemma 4.6 uses monotonicity (from (A7)) as the regularity input, keeping the argument elementary on the restricted domain $\pi_1 + \pi_2 \leq 1$ where the Cauchy additivity of (A9) holds for $\tilde{h}$. The Banach-Sierpiński route yields the same conclusion but requires extending $\tilde{h}$ to all of $\mathbb{R}$, which is awkward on a bounded domain. Monotonicity is therefore preferred.

We can now assemble the proof of Theorem 4.1.

Proof of Theorem 4.1. Sufficiency. We first verify that the rule (1) satisfies each axiom. (A1): $\partial m/\partial q = -\gamma_0\sigma^2 < 0$. (A2): $m-\mu$ is linear in q with coefficient $-\gamma_0\sigma^2$ independent of μ. (A3): both $m-\mu$ and δ are proportional to σ^2. (A4) at $\pi = 0$: $\delta/2 = \kappa_{\text{inv}}\sigma^2/\lambda = \sigma^2\,\kappa'_+(|q|)/\lambda$ under C5. (A5): $m - \mu = -\gamma_0\sigma^2 q$ does not depend on λ. (A6): $m - \mu$ does not depend on π. (A7): $\partial\delta/\partial\pi = 2\kappa_{\text{adv}}\sigma^2 > 0$ trivially. (A8): $h(\sigma^2,\pi) = 2\kappa_{\text{adv}}\pi\sigma^2$ depends only on (σ^2,π). (A9): $h(\sigma^2,\pi_1+\pi_2) = 2\kappa_{\text{adv}}(\pi_1+\pi_2)\sigma^2 = h(\sigma^2,\pi_1) + h(\sigma^2,\pi_2)$.

Necessity. Lemma 4.3 gives $m - \mu = -\gamma_0 \sigma^2 q$ for some $\gamma_0 > 0$. Lemma 4.5 implies that $\delta|_{\pi=0} = 2\kappa_{\text{inv}}\sigma^2/\lambda$. Lemma 4.6 gives $h(\sigma^2, \pi) = 2\kappa_{\text{adv}}\pi\sigma^2$. Combining, we have $\delta = 2\sigma^2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi) > 0$ and $m = \mu - \gamma_0\sigma^2 q$, so $(b, a) = (m \mp \delta/2)$ takes the form (1) and respects $b < a$. □

Remark 4.8 (Parameter Interpretations)**.** The three parameters in (1) admit transparent economic interpretations. The risk aversion $\gamma_0 > 0$ is the maker's inventory risk aversion per unit of price variance per unit of inventory. The carrying cost $\kappa_{\text{inv}} > 0$ is the marginal cost per unit of held inventory per unit of price variance. The adverse-selection cost $\kappa_{\text{adv}} > 0$ is the expected loss per informed trade per unit of price variance. Concretely, γ_0 is a preference primitive, κ_{inv} a cost primitive, and κ_{adv} an information-environment primitive.

The spread admits an exact additive decomposition,

$$\delta(s) = \underbrace{\frac{2\kappa_{\text{inv}}\sigma^2}{\lambda}}_{\text{inventory}} + \underbrace{2\kappa_{\text{adv}}\pi\sigma^2}_{\text{adverse selection}} . \tag{2}$$

Both components scale with σ^2, a consequence of (A3). The inventory component is inversely proportional to λ, because more frequent trades dilute the per-trade carrying cost. The adverse-selection component is proportional to π. The structural consequences are developed in Sections 6 and 9.

4.1 Derived Properties of the Quoting Rule

Three structural properties emerge as direct theorems rather than imposed separately: translation equivariance in μ (relying on (A2), (A4), (A8)), oddness of the skew in q (relying on (A2) alone), and continuity in each state variable (relying on (A2)–(A9) with measurability).

Theorem 4.9 (Translation Equivariance)**.** *Under (A2), (A4), and (A8), the quoting rule satisfies, for every $s = (q, \mu, \sigma^2, \lambda, \pi) \in \mathcal{S}$ and every $\tau \in \mathbb{R}$,*

$$Q(q, \mu + \tau, \sigma^2, \lambda, \pi) = Q(q, \mu, \sigma^2, \lambda, \pi) + (\tau, \tau).$$

Proof. By (A2), $m(s) = \mu - \gamma(\sigma^2, \lambda, \pi)q$ shifts uniformly with μ. By (A4), the inventory half-spread depends only on (q, σ^2, λ), and by (A8) the adverse-selection increment depends only on (σ^2, π). The full spread $\delta(s)$ is therefore μ-independent, so $b = m - \delta/2$ and $a = m + \delta/2$ each shift uniformly with μ. □

Corollary 4.10 (Skew Symmetry)**.** *Under (A2), the mid-quote has $m(0, \mu, \sigma^2, \lambda, \pi) = \mu$ for every $(\mu, \sigma^2, \lambda, \pi)$, and the skew is odd in q, that is,*

$$m(-q, \mu, \sigma^2, \lambda, \pi) - \mu = -\big(m(q, \mu, \sigma^2, \lambda, \pi) - \mu\big)$$

for every $(q, \mu, \sigma^2, \lambda, \pi) \in \mathcal{S}$.

Proof. By (A2), $m(s) - \mu = -\gamma(\sigma^2, \lambda, \pi)q$, so setting $q = 0$ gives $m(0, \ldots) = \mu$, and replacing q by $-q$ flips the sign. □

Theorem 4.11 (Continuity of the Quoting Rule)**.** *Under (A2), (A3), (A4), (A5), (A6), (A7), (A8), and (A9), together with the Borel-measurability assumption of Section 2 and the cost-structure assumptions C1 through C6, the quoting rule $Q\colon \mathcal{S} \to \mathcal{B}$ is continuous in each of the five state variables separately, and moreover is jointly continuous.*

Proof. We establish continuity coordinate by coordinate.

By (A2), the skew $-\gamma(\sigma^2, \lambda, \pi)q$ is $\mathbb{R}$-linear in q and shifts additively in μ. By (A4) with C5 and (A8), the spread is constant in q and μ. By (A3), $m - \mu$ and δ are $\mathbb{R}$-linear in σ^2 on $\mathbb{R}_{++}$. By (A5), the skew is λ-independent, and the inventory half-spread $\kappa_{\text{inv}}\sigma^2/\lambda$ is continuous in λ on $\mathbb{R}_{++}$. For the informed fraction π, axiom (A6) makes the skew π-independent. By (A9) the adverse-selection increment satisfies the Cauchy equation on $[0, 1]$ with $h(\sigma^2, 0) = 0$, and by (A7) is monotone. A Borel-measurable monotone solution of the Cauchy equation is necessarily $\mathbb{R}$-linear (Aczél (1966, Chapter 2)), ruling out Hamel-basis pathologies.

Joint continuity follows from Theorem 4.1. Each component of Q is a polynomial in (q, μ) and rational in (σ^2, λ, π). □

The three theorems together identify the maker's quote as a nondegenerate book centered on her belief at zero inventory, equivariant under translations of that belief, and continuous in each state variable. The symmetry theorems rely on (A2). The continuity theorem additionally uses (A3), (A5), (A6), (A4), (A7), (A8), (A9) and measurability. The only axiom not invoked is the inventory-aversion axiom (A1).

5 The Bijection and Irredundancy of the Axioms

The Full Characterization identifies a three-parameter family of admissible quoting rules. We now show this family is in canonical bijection with $\mathbb{R}^3_{++}$, and that no axiom is redundant.

5.1 The Bijection Theorem

Let $\mathcal{A} = \{(\text{A1}), \ldots, (\text{A9})\}$ denote the nine-axiom system, and let $\mathcal{Q}$ denote the set of measurable quoting rules satisfying $\mathcal{A}$ together with C1–C6.

Theorem 5.1 (Bijection between Axioms and Parameters)**.** *The set $\mathcal{Q}$ is in canonical bijection with $\mathbb{R}^3_{++}$. The bijection is given by the parameter map*

$$Q \longmapsto (\gamma_0(Q),\, \kappa_{\text{inv}}(Q),\, \kappa_{\text{adv}}(Q)),$$

where

$$\gamma_0(Q) = -\frac{1}{\sigma^2}\frac{\partial m_Q}{\partial q}, \qquad \kappa_{\text{inv}}(Q) = \frac{\lambda\, \delta_Q|_{\pi=0}}{2\sigma^2}, \qquad \kappa_{\text{adv}}(Q) = \frac{1}{2\sigma^2}\frac{\partial \delta_Q}{\partial \pi}.$$

The inverse map sends $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) \in \mathbb{R}^3_{++}$ to the quoting rule defined by (1).

Axiom dependency. Theorem 5.1 relies on the full nine-axiom system.

Proof. Injectivity: distinct parameter values give distinct quoting rules, since the three partial derivatives separate the parameters. Surjectivity: by Theorem 4.1, every $Q \in \mathcal{Q}$ takes the form (1) for some $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) \in \mathbb{R}^3_{++}$. The three quantities defined in the parameter map are precisely those constants for the corresponding Q. □

The Bijection Theorem is the formal statement of forced uniqueness. Nine axioms reduce the infinite-dimensional space of measurable functions $\mathcal{S} \to \mathcal{B}$ to a three-dimensional manifold parametrized by $\mathbb{R}^3_{++}$. Each parameter corresponds to one axiom group, the parameter map is a bijection onto $\mathbb{R}^3_{++}$, and no axiom contributes redundantly to the same parameter as another.

5.2 Irredundancy of the Axioms

The nine axioms are irredundant: removing any one enlarges the solution set $\mathcal{Q}$.

Proposition 5.2 (Irredundancy)**.** *For each axiom $A_i \in \mathcal{A}$, there exists a measurable quoting rule satisfying all axioms in $\mathcal{A} \setminus \{A_i\}$ but violating A_i.*

Proof. We exhibit a measurable counterexample for each axiom. In each case, the perturbation modifies only the channel governed by the targeted axiom (skew, inventory spread, or adverse-selection increment), so the eight other axioms are immediate from the unchanged components of the canonical rule. We give the counterexample and indicate the local verification.

Violating (A1): the rule with skew $+\gamma_0\sigma^2 q$ (sign reversed) makes the mid-quote increasing in q, violating inventory aversion. The spread is unchanged at $\delta = 2\sigma^2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi) > 0$, so $b < a$, and the sign-reversed skew remains linear in q with coefficient independent of μ, λ, π and degree-one in σ^2, satisfying (A2), (A3), (A5), (A6). Axioms (A4), (A7), (A8), (A9) hold since neither the inventory nor adverse-selection spread component is touched.

Violating (A2): the rule with cubic skew $-\gamma_0\sigma^2 q^3$ violates linearity. The cubic skew is strictly decreasing on $\mathbb{R}$, so (A1) is satisfied. It is degree-one in σ^2 and has no λ- or π-dependence, and the spread is unchanged.

Violating (A3): the rule $m(s) - \mu = -\gamma_0\sigma^4 q$, $\delta(s) = 2\kappa_{\text{inv}}\sigma^2/\lambda + 2\kappa_{\text{adv}}\pi\sigma^2$ violates uniform σ^2-homogeneity of the skew, which is now degree two in σ^2 rather than degree one, while leaving every other channel as in the canonical rule, so (A1), (A2), (A4)–(A9) continue to hold.

Violating (A4): the rule with $\delta|_{\pi=0}/2 = 2\kappa_{\text{inv}}\sigma^2/\lambda$ (twice the correct value) violates competitive break-even on carrying cost. The skew is unchanged so (A1), (A2), (A5), (A6) hold. The modified spread $4\kappa_{\text{inv}}\sigma^2/\lambda + 2\kappa_{\text{adv}}\pi\sigma^2$ remains degree-one in σ^2, strictly increasing in π, and the canonical adverse-selection increment is preserved, so (A3), (A7), (A8), (A9) hold.

Violating (A5): the rule with skew $-\gamma_0\sigma^2 q/(1+\lambda)$ makes the skew depend on λ. The modified skew is linear in q, degree-one in σ^2, μ- and π-independent, and strictly decreasing in q since $1 + \lambda > 0$, so (A1), (A2), (A3), (A6) hold. The spread is unchanged, so (A4), (A7), (A8), (A9) hold.

Violating (A6): the rule with skew $-\gamma_0\sigma^2 q(1 - \pi/2)$ makes the skew depend on π. The modified skew is linear in q, degree-one in σ^2, μ- and λ-independent, and strictly decreasing in q since $1 - \pi/2 \geq 1/2 > 0$ on $[0, 1]$, so (A1), (A2), (A3), (A5) hold. The spread is unchanged, so (A4), (A7), (A8), (A9) hold.

Violating (A7): the rule with $h(\sigma^2, \pi) \equiv 0$ produces a constant spread $\delta = 2\kappa_{\text{inv}}\sigma^2/\lambda$ that does not depend on π, violating strict monotonicity of the spread in π. The skew is unchanged so (A1)–(A6) hold, and the zero increment depends only on (σ^2, π) (in fact on nothing) and is additive $(0 = 0 + 0)$, so (A8) and (A9) hold.

Violating (A8): the rule with $h(\sigma^2, \pi) = 2\kappa_{\text{adv}}\pi\sigma^2(1 + q^2)$ makes the adverse-selection increment depend on q. The skew is unchanged so (A1)–(A6) hold (note that (A4) is unaffected since the perturbation vanishes at $\pi = 0$). At each fixed q, the increment is linear in π with strictly positive coefficient $2\kappa_{\text{adv}}\sigma^2(1 + q^2)$, so (A7) and (A9) hold.

Violating (A9): the rule with $h(\sigma^2, \pi) = \kappa_{\text{adv}}\sigma^2\pi/(1 + \pi)$ makes the adverse-selection increment concave rather than linear in π. Non-additivity follows by direct check: we have $h(\sigma^2, 1) = \kappa_{\text{adv}}\sigma^2/2$ while $2h(\sigma^2, 1/2) = 2\kappa_{\text{adv}}\sigma^2/3 > \kappa_{\text{adv}}\sigma^2/2$. The skew is unchanged so (A1)–(A6) hold, and the increment is strictly increasing in π on $[0, 1]$ (its derivative is equal to$\kappa_{\text{adv}}\sigma^2/(1 + \pi)^2 > 0$) and depends only on (σ^2, π), so (A7) and (A8) hold.

Each counterexample is measurable, so the corresponding axiom carries content that cannot be reduced to a measurability requirement. □

Irredundancy is structurally significant. It establishes that no axiom is redundant, and that the axiom system is tight. Removing any axiom enlarges $\mathcal{Q}$, and adding any further axiom (consistent with the existing system) either is redundant or restricts $\mathcal{Q}$ to a strict subset.

6 Structural Separability and Empirical Identification

This section establishes the first of the two main payoffs of the theory. The three economic parameters and the four state variables play isolated nonoverlapping roles in the quoting rule. The separability is not assumed. It is forced by the axiom system. The consequence (Corollary 6.3) is that each of the three structural primitives is point-identified from a single, distinct moment of the observed rule, with no joint estimation required. This resolves the conflation problem present in the optimization-based literature, where the same parameter often enters multiple components of the quoting rule and identification requires auxiliary modeling commitments.

Theorem 6.1 (Structural Separability)**.** *Under the Full Characterization (Theorem 4.1), the quoting rule satisfies three separability properties.*

(i) Separability of economic primitives. *Each parameter appears in exactly one structural component. The parameter γ_0 appears only in the skew, κ_{inv} appears only in the inventory spread, and κ_{adv} appears only in the adverse-selection spread.*

(ii) Additive separability of the spread. *The spread is $\delta(s) = f_{\text{inv}}(\sigma^2, \lambda) + f_{\text{adv}}(\sigma^2, \pi)$, where it holds that $f_{\text{inv}}(\sigma^2, \lambda) = 2\kappa_{\text{inv}}\sigma^2/\lambda$ depends only on inventory-related variables and $f_{\text{adv}}(\sigma^2, \pi) = 2\kappa_{\text{adv}}\pi\sigma^2$ depends only on information-related variables.*

(iii) Information modularity. *Each state variable other than σ^2 affects exactly one structural component. The inventory q affects only the skew, the intensity λ only the inventory spread, and the informed fraction π only the adverse-selection spread. The variance σ^2 affects all components uniformly.*

Axiom dependency. Theorem 6.1 relies on the full nine-axiom system.

Proof. All three claims follow by direct inspection of (1). □

The axioms enforce a complete modular decomposition. Each parameter corresponds to one economic primitive (preference, cost, information), and each state variable affects exactly one structural component, with no cross-effects.

Remark 6.2 (Contrast with Utility-Based Models)**.** The structural separability is a distinctive feature of the axiomatic approach. In utility-based models such as Avellaneda-Stoikov, the risk-aversion parameter γ enters both the skew (as the inventory-hedging coefficient) and the spread (through a risk-premium term). The two contributions are bundled within the single parameter, and identification requires joint estimation across multiple moments. In Ho-Stoll, the inventory-cost coefficient enters both the skew and the spread similarly. The axiomatic theory characterizes the family of quoting rules in which preference, cost, and information primitives are mathematically decoupled. This decoupling is not assumed. It is a consequence of the axiom system.

The empirical content of Structural Separability is direct. Each parameter is identified from a single, distinct moment of the observable quoting rule.

Corollary 6.3 (Empirical Identification)**.** *Under the Full Characterization, the parameters are identified by*

$$\gamma_0 = -\frac{1}{\sigma^2}\frac{\partial m}{\partial q}, \qquad \kappa_{\text{inv}} = \frac{\lambda\,\delta|_{\pi=0}}{2\sigma^2}, \qquad \kappa_{\text{adv}} = \frac{1}{2\sigma^2}\frac{\partial \delta}{\partial \pi}.$$

The map from observable quoting behavior to the parameter space $\mathbb{R}^3_{++}$ is a bijection, and each parameter is recoverable from a single distinct moment.

Proof. Each parameter is a distinct moment by Theorem 6.1: γ_0 is a partial derivative of the skew, κ_{inv} the spread level at $\pi = 0$, κ_{adv} a partial derivative of the spread. The bijection is Theorem 5.1. □

The identification corollary delivers a sharp empirical strategy. The slope of the mid-quote with respect to inventory identifies γ_0. The spread at zero informed fraction identifies κ_{inv}. The slope of the spread with respect to π identifies κ_{adv}. There are no joint identification problems, no nonlinear estimation procedures, and no distributional assumptions beyond the axioms.

6.1 Application to Empirical Microstructure

Three separate regressions on a sample of the maker's quotes recover the three parameters. The mid-quote on inventory yields γ_0, the spread under low informed flow yields κ_{inv}, and the spread response to informed flow yields κ_{adv}. Each regression isolates one parameter without interference from the others. The identification through partial derivatives also serves as a model test. If the observed quoting rule does not satisfy (1), then at least one axiom is violated, and the form of the violation is informative. A nonlinear skew points to (A2), a σ^4-scaling to (A3), a π-dependent skew to (A6), and so on. The axiom system thus serves as a diagnostic framework for empirical anomalies.

7 The Limit Order Book and the Recovery Theorem

This section establishes the second of the two main payoffs of the theory. The characterization so far concerns the best bid and ask, that is, the maker's quotes for a single unit of inventory change. In practice, a market maker posts a limit order book, a schedule of prices at multiple levels of inventory. Extending the theory to the full order book yields a stronger empirical result. The entire latent inventory cost function κ is observable in the depth profile by integration. In the optimization-based tradition, the cost coefficient is a free parameter to be calibrated under each chosen specification. Here the cost function is an object the theory predicts to be directly observable. Section 8 reads the result structurally as identifying the order book with the graph of the Legendre-Fenchel conjugate pair (κ, κ^*).

This section adopts the *dynamic-schedule reading* of the order book. The n-th level a_n is the maker's ask at the counterfactual post-trade inventory $q + n - 1$, traversed as her inventory varies through successive same-side fills. This is what the axiomatic theory delivers at each level via the marginal-indifference condition (A4) at the corresponding inventory state. Empirically, the dynamic schedule is recovered not from a single snapshot of posted depth, but from the panel of her quotes as her inventory varies through the integer levels.

7.1 The Order Book under General Convex Cost

We relax the linearity assumption C5 to a general convex nondecreasing cost κ with $\kappa(0) = 0$, retaining C1–C4 and C6. The maker now posts a schedule of asks $a_1, a_2, \ldots$ and bids $b_1, b_2, \ldots$, where a_n is her quote at inventory $q + n - 1$ (after $n - 1$ prior bid-fills) and b_n at $q - n + 1$ (after $n - 1$ prior ask-fills). At each level the maker is indifferent between filling and remaining at the pre-trade inventory, the level-by-level analog of (A4).

Remark 7.1 (Schedule Shape)**.** Under C5, the dynamic ask schedule decreases in n (the skew dominates the constant marginal carrying cost), reflecting that as inventory accumulates, the ask falls to attract sellers. Symmetrically, the bid schedule increases in n. This shape is a distinctive prediction and is what the Recovery Theorem exploits.

Theorem 7.2 (Order Book under Convex Cost)**.** *Under (A1) through (A8), together with C1 through C4 and C6 (general convex nondecreasing κ, with linearity C5 dropped), and under the dynamic-schedule reading of the order book described in the preceding remark, the n-th ask is*

$$a_n - \mu = -\gamma_0 \sigma^2 (q + n - 1) + \frac{\sigma^2 \, \kappa'_+(|q + n - 1|)}{\lambda} + \kappa_{\text{adv}} \pi \sigma^2,$$

and the n-th bid is

$$b_n - \mu = -\gamma_0 \sigma^2 (q - n + 1) - \frac{\sigma^2 \, \kappa'_+(|q - n + 1|)}{\lambda} - \kappa_{\text{adv}} \pi \sigma^2,$$

for $n = 1, 2, \ldots$, where κ'_+ denotes the right derivative.

Axiom dependency. Theorem 7.2 relies on the full nine-axiom system.

Proof. The proof tracks each of the three components of the quoting rule at the pre-trade inventory of level n, then assembles.

Level-by-level indifference. At level n, the maker has accumulated $n-1$ prior long positions, bringing her inventory to $q+n-1$, and a_n is the price at which she sells one unit from this state. (A4) applied at this pre-trade inventory equates her per-trade revenue with the marginal carrying cost $\sigma^2\,\kappa'_+(|q+n-1|)/\lambda$. By C2, this marginal cost is well-defined regardless of the sign of $q+n-1$.

Components at level n. By Lemma 4.3, the mid-quote shift at inventory q' is $-\gamma_0\sigma^2 q'$, so the skew contribution to $a_n-\mu$ is $-\gamma_0\sigma^2(q+n-1)$. This uses only (A1), (A2), (A3), (A5), (A6), so the contribution is independent of the convexity of κ or C5. By the level-by-level (A4) condition (Step above), the inventory half-spread at level n is $\sigma^2\kappa'_+(|q+n-1|)/\lambda$, where κ'_+ denotes the right derivative, finite on $\mathbb{R}_+$ by C6 and C3 (with $\kappa'_+(0)=0$ permitted). Under C5, this reduces to $\kappa_{\text{inv}}\sigma^2/\lambda$ at every level. By (A8) and Lemma 4.6, the adverse-selection half-spread is $\kappa_{\text{adv}}\pi\sigma^2$, constant in n.

Assembly. The n-th ask is the sum of the mid-quote at $q+n-1$ and the half-spread at level n:

$$a_n=\mu-\gamma_0\sigma^2(q+n-1)+\frac{\sigma^2\,\kappa'_+(|q+n-1|)}{\lambda}+\kappa_{\text{adv}}\pi\sigma^2,$$

giving the stated formula. For the n-th bid, the maker considers the symmetric transition with the inventory decreasing to $q-n+1$. By C2, the inventory half-spread $\sigma^2\kappa'_+(|q-n+1|)/\lambda$ depends only on $|q-n+1|$, so the bid-side formula is the symmetric reflection of the ask-side:

$$b_n=\mu-\gamma_0\sigma^2(q-n+1)-\frac{\sigma^2\,\kappa'_+(|q-n+1|)}{\lambda}-\kappa_{\text{adv}}\pi\sigma^2.$$

□

7.2 The Recovery Theorem

The entire cost function κ is recoverable from the observable limit order book by integration across levels; its full shape is encoded in the depth profile, not a free parameter to be calibrated separately.

Theorem 7.3 (Recovery of the Inventory Cost)**.** *Under the conditions of Theorem 7.2, the cost function κ is recoverable from the observable order book. For the maker at inventory $q=0$, the cost function on $[0,Q]$ is*

$$\kappa(Q)=\int_0^Q \kappa'_+(x)\,dx,$$

where the cost gradient is recovered from the order book at continuous depth x via the identity $\kappa'_+(x)=(\lambda/\sigma^2)(a(x)-\mu+\gamma_0\sigma^2 x-\kappa_{\text{adv}}\pi\sigma^2)$, by Theorem 7.2. At integer depths only, the left-endpoint Riemann sum

$$\frac{\lambda}{\sigma^2}\sum_{n=1}^{Q}\left(a_n-\mu+\gamma_0\sigma^2(n-1)-\kappa_{\text{adv}}\pi\sigma^2\right)$$

recovers $\kappa(Q)$ exactly under C5 (where κ'_+ is constant on $\mathbb{R}_+$) and is a lower bound on $\kappa(Q)$ under general convex κ (where κ'_+ is nondecreasing, so the left-endpoint sum understates

the integral). For a maker at arbitrary inventory q, the same construction recovers the cost increment $\kappa(q+Q)-\kappa(q)$ over the interval $[q, q+Q]$ by integrating the level-by-level cost gradients.

Axiom dependency. Theorem 7.3 relies on the full nine-axiom system.

Proof. From Theorem 7.2, the cost gradient at level n is

$$\kappa'_+(|q+n-1|) = \frac{\lambda}{\sigma^2}\left(a_n - \mu + \gamma_0\sigma^2(q+n-1) - \kappa_{\text{adv}}\pi\sigma^2\right).$$

The formula presupposes γ_0 and κ_{adv}, identified by Corollary 6.3 from the best-bid-best-ask quotes, so the recovery is sequentially identifiable.

The smooth integral formula follows by the fundamental theorem of calculus with $\kappa(0)=0$ (C1). Under C5, $\kappa'_+(x)=\kappa_{\text{inv}}$ is constant on $\mathbb{R}_+$, so the Riemann sum equals $Q\kappa_{\text{inv}}=\kappa(Q)$ exactly. Under general convex κ, κ'_+ is nondecreasing, so $\kappa'_+(x)\geq\kappa'_+(n-1)$ for $x\in[n-1,n)$, giving $\int_0^Q \kappa'_+(x)\,dx \geq \sum_{n=1}^{Q}\kappa'_+(n-1)$.

At arbitrary q, the same construction recovers $\kappa(q+Q)-\kappa(q)$ by integrating the level-by-level gradients. □

The Recovery Theorem has a striking empirical implication. The maker's inventory cost function is not a hidden primitive that must be calibrated separately. Its full shape is read off the order book depth profile by integration of the observable quotes. In contrast to Ho-Stoll or Avellaneda-Stoikov, where the cost coefficient is a free parameter fit against data, the axiomatic theory predicts that the cost function itself is observable, not merely its scalar summary. The information content of the order book is therefore greater than that of the best-bid-best-ask alone: the depth profile encodes the entire cost geometry, not merely its value at the touch.

Remark 7.4 (Empirical Identification Strategy)**.** A complete identification strategy combines Corollary 6.3 with Theorem 7.3. From the best quotes, γ_0 is identified from the slope of m in q and κ_{adv} from the slope of δ in π. The function κ is then recovered in full from the order book by integrating the cost-gradient quotes across depth levels. The three identifications are distinct, use different moments, and impose no cross-restrictions.

8 The Legendre-Fenchel Duality of Market Making

The spread-cost relationship under general convex cost (Theorem 7.2) has a deep structural reformulation. Market making is, mathematically, an instance of Legendre-Fenchel convex conjugate duality applied to the inventory cost function. We refer to Rockafellar (1970) and Hiriart-Urruty and Lemaréchal (2001) for the standard treatment used in this section.

8.1 The Legendre-Fenchel Conjugate

Let $\kappa\colon \mathbb{R}_+\to\mathbb{R}_+$ be a convex nondecreasing function with $\kappa(0)=0$. By symmetry of κ in q (C2), restricting to the positive half-line loses no information. Its Legendre-Fenchel conjugate is

$$\kappa^*(p) = \sup_{q\geq 0}\{pq-\kappa(q)\}.$$

The conjugate κ^* is itself convex on $\mathbb{R}_+$ and nondecreasing whenever κ is. Geometrically, $\kappa^*(p)$ is the maximum vertical gap between the line $q \mapsto pq$ and the graph of κ, and equivalently $-\kappa^*(p)$ is the intercept of the supporting line of κ with slope p. At points where κ is differentiable, the supremum is attained at the unique q^* with $\kappa'(q^*) = p$. This is the heart of conjugate duality. The forward map $q \mapsto p = \kappa'(q)$ and the inverse map $p \mapsto q = (\kappa^*)'(p)$ are inverses of each other, and they satisfy the Legendre-Fenchel identity

$$\kappa(q) + \kappa^*(p) = pq \quad \text{whenever } p = \kappa'(q).$$

The convex conjugate has a long history, from Legendre's eighteenth-century work to Fenchel's twentieth-century generalization to nondifferentiable convex functions. It underlies the Lagrangian-Hamiltonian duality of classical mechanics, the energy-free-energy duality of thermodynamics, the expenditure-indirect-utility duality of consumer theory, and the primal-dual structure of convex optimization.

Our cost specifications correspond to three special cases. For *linear cost* $\kappa(q) = \kappa_{\text{inv}} q$ (canonical under C5), the conjugate is the indicator of $(-\infty, \kappa_{\text{inv}}]$,

$$\kappa^*(p) = \begin{cases} 0 & \text{if } p \le \kappa_{\text{inv}}, \\ +\infty & \text{if } p > \kappa_{\text{inv}}, \end{cases}$$

since $(p - \kappa_{\text{inv}})q$ has supremum 0 at $q = 0$ when $p \le \kappa_{\text{inv}}$ and is unbounded above otherwise. For *quadratic cost* $\kappa(q) = \kappa_{\text{inv}} q^2/2$, the first-order condition $p = \kappa_{\text{inv}} q$ gives $\kappa^*(p) = p^2/(2\kappa_{\text{inv}})$, also quadratic with reciprocal coefficient. The quadratic family is self-dual up to coefficient reciprocation, with κ_{inv} passing to $1/\kappa_{\text{inv}}$ under the conjugate. At $\kappa_{\text{inv}} = 1$, the cost and its conjugate coincide identically. For *power cost* $\kappa(q) = q^p/p$ with $p > 1$, the first-order condition $\tilde{p} = q^{p-1}$ gives $\kappa^*(\tilde{p}) = \tilde{p}^{p^*}/p^*$, where $p^* = p/(p-1)$ is the conjugate exponent. The duality swaps p and p^*.

8.2 The Duality Theorem

We now show how the quoting rule, under (A4) and a smooth convex inventory cost, encodes both members of the Legendre-Fenchel conjugate pair (κ, κ^*). The maker's indifference condition relates the inventory level q to the half-spread, which under appropriate rescaling plays the role of the marginal price conjugate to q. The bid and ask quotes jointly trace the subdifferential graph of κ.

Theorem 8.1 (Legendre-Fenchel Duality)**.** *Under (A4) with smooth strictly convex inventory cost κ, the quoting rule encodes the Legendre-Fenchel conjugate pair (κ, κ^*) as follows.*

1. *The forward map $q \mapsto \tilde{p} = \kappa'(q)$ is the derivative of κ. By (A4), the rescaled half-spread is $\tilde{p} = (\lambda/\sigma^2)(\delta(q)/2 - \kappa_{\text{adv}}\pi\sigma^2)$.*
2. *The inverse map $\tilde{p} \mapsto q = (\kappa^*)'(\tilde{p})$ is the derivative of the conjugate κ^*.*
3. *The Legendre-Fenchel identity holds at the indifference point, $\kappa(q) + \kappa^*(\tilde{p}) = q\tilde{p}$.*

Axiom dependency. Theorem 8.1 relies on (A4) together with the smooth strictly convex cost structure, and is therefore a Layer I result.

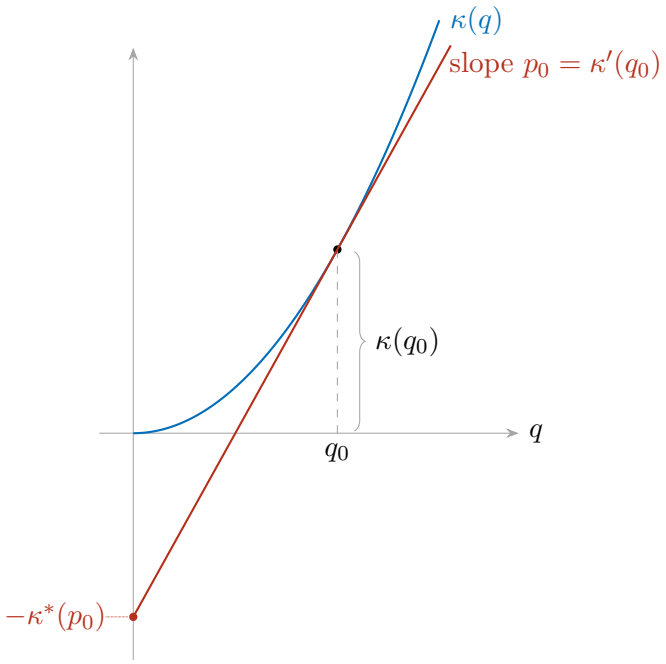


Figure 2: Legendre-Fenchel geometry for the quadratic cost $\kappa(q) = q^2/2$. The tangent to κ at inventory level q_0 has slope $p_0 = \kappa'(q_0)$, which is the marginal price, and vertical intercept $-\kappa^*(p_0)$. Sweeping q_0 across all inventory levels traces out both κ and its conjugate κ^* simultaneously: the slopes recover the forward map $q_0 \mapsto p_0$, and the intercepts recover the dual values $p_0 \mapsto \kappa^*(p_0)$. The quadratic is self-dual: $\kappa^*(p) = p^2/2$.

Proof. Part 1: (A4) rearranges to $\sigma^2\kappa'(q)/\lambda = \delta(q)/2 - \kappa_{\text{adv}}\pi\sigma^2$, equivalently $\kappa'(q) = \tilde{p}$ with the indicated rescaling.

Part 2: by strict convexity, $\kappa' \colon \mathbb{R}_+ \to \kappa'(\mathbb{R}_+)$ is a continuous strictly increasing bijection, hence a homeomorphism, and its inverse coincides with $(\kappa^*)'$ (Rockafellar (1970), Thm. 26.5). A convex but not strictly convex κ would admit flat segments where κ' is constant and not invertible, so strict convexity is needed for Part 2. We have that $q = (\kappa^*)'(\tilde{p}) \iff \tilde{p} = \kappa'(q)$ holds on the relevant domain.

Part 3: the Young-Fenchel inequality $\kappa(q) + \kappa^*(\tilde{p}) \geq q\tilde{p}$ holds for all $(q, \tilde{p})$ and is an equality at the indifference point, where $\tilde{p} = \kappa'(q)$. The equality follows from the definition of κ^* as a supremum, which is attained at q when $\tilde{p} = \kappa'(q)$. Figure 2 illustrates the geometry. □

8.3 The Market Making Interpretation of the Special Cases

The conjugate examples in Section 8.1 are now reread as statements about market making. Under *linear cost* (C5), the maker has constant marginal cost. The conjugate κ^* is the indicator of $(-\infty, \kappa_{\text{inv}}]$, meaning the maker is willing to fill at any half-spread up to $\kappa_{\text{inv}}\sigma^2/\lambda$, with the indifference condition fixing it exactly. Under *quadratic cost* $\kappa(q) = \kappa_{\text{inv}}q^2/2$, the order book at level n posts a half-spread proportional to $|q + n - 1|$ (Theorem 7.2), giving a

depth profile linear in n with slope $2\kappa_{\text{inv}}\sigma^2/\lambda$. The dual picture has inventory proportional to half-spread with inverse coefficient, the two readings equivalent under Legendre-Fenchel. Under *power cost* $\kappa(q) = q^p/p$ for $p > 1$, the depth profile is concave in n for $p < 2$ (sublinear cost growth) and convex for $p > 2$ (superlinear growth). The dual has inverse curvature, with cost and dual exponents connected by $1/p + 1/p^* = 1$.

Remark 8.2 (Recovery and Duality)**.** Theorem 8.1 unifies the Recovery Theorem with the duality structure. The spread-cost relation $\tilde{p} = \kappa'(q)$ is the first-order expression of Legendre-Fenchel duality, and the half-spread reveals the derivative of κ. The Recovery Theorem is the integral expression of the same duality, by which κ is recovered through integrating κ' across the depth profile. Both κ and κ^* are recoverable from observable market data.

Market making, viewed through the duality structure, belongs to convex analysis. The maker's inventory cost and her marginal-price function are conjugate variables under the Legendre-Fenchel transform, and the limit order book reveals both simultaneously: the bid and ask quotes trace out the derivative κ' level by level, while the residuals against the tangent lines encode κ^*. The two functions are recoverable from the same observable object, with no auxiliary calibration required. The reformulation identifies market making as a particular instance of a general structure recurring across physics (thermodynamic duality), economics (consumer theory), and optimization theory (dual programs), placing the bid-ask spread alongside the Hamiltonian, the expenditure function, and the dual objective as instances of the same convex-conjugate machinery.

9 Phase Transitions and the Three Liquidity Regimes

Combining the spread formula (2) with a disagreement bound between maker and counterparties yields a sharp boundary in parameter space separating a functioning regime from a frozen regime, with the functioning region partitioning further by which spread channel dominates.

9.1 The Disagreement Bound and the Phase Boundary

Consider a counterparty with belief μ_C. The counterparty buys at the ask a only if $\mu_C > a$, which requires $\mu_C - \mu > \delta/2$, and symmetrically at the bid. Recall from Section 2 that $\delta_{\max} := 2\max_C |\mu_C - \mu|$ is the maximum admissible spread. The market generates positive trade volume if and only if $\delta < \delta_{\max}$.

Proposition 9.1 (Liquidity Phase Boundary)**.** *Given $\delta_{\max} > 0$, structural parameters $\kappa_{\text{inv}}, \kappa_{\text{adv}} > 0$, volatility $\sigma^2 > 0$, and $\pi \in [0, 1]$, the market is functioning if and only if*

$$\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi < \frac{\delta_{\max}}{2\sigma^2}. \tag{3}$$

Axiom dependency. Proposition 9.1 relies on the core (Layer I) together with the structural choice (A3).

Proof. By (2), $\delta = 2\sigma^2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$. The viability condition $\delta < \delta_{\max}$ rearranges to the stated inequality. □

9.2 Critical Thresholds

The phase condition admits three natural extreme cases. The critical liquidity at fixed π is

$$\lambda^*(\pi) = \frac{\kappa_{\text{inv}}}{\delta_{\max}/(2\sigma^2) - \kappa_{\text{adv}}\pi},$$

with the market functioning for $\lambda > \lambda^*$. More informed flow raises the threshold. The critical informed fraction at $\lambda \to \infty$ is $\pi^* = \delta_{\max}/(2\kappa_{\text{adv}}\sigma^2)$, and for $\pi > \pi^*$ the market is frozen regardless of liquidity, since the adverse-selection premium alone exceeds $\delta_{\max}$. If it holds that $\delta_{\max} > 2\kappa_{\text{adv}}\sigma^2$, then $\pi^* > 1$ and adverse selection alone cannot freeze the market regardless of π. The frozen regime then requires sufficiently low λ. The critical volatility at fixed (λ, π) is

$$\sigma^2_{\max}(\lambda, \pi) = \frac{\delta_{\max}}{2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)}.$$

A spike in uncertainty can push the market across the phase boundary even with constant (λ, π).

9.3 The Three Liquidity Regimes

Within the functioning region, the spread admits a substructure governed by which of its two components dominates.

Theorem 9.2 (Three Liquidity Regimes)**.** *Define the crossover liquidity $\lambda_\times = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$. Assume the ordering condition*

$$\kappa_{\text{adv}}\pi < \frac{\delta_{\max}}{4\sigma^2},$$

equivalently $\lambda^*(\pi) < \lambda_\times$. The parameter space then partitions into three liquidity regimes.

Frozen regime *($\lambda < \lambda^*(\pi)$). The equilibrium spread exceeds $\delta_{\max}$, so no trade occurs.*

Inventory-dominated regime *($\lambda^*(\pi) < \lambda < \lambda_\times$). The inventory spread $2\kappa_{\text{inv}}\sigma^2/\lambda$ exceeds the adverse-selection spread $2\kappa_{\text{adv}}\pi\sigma^2$. The market is functioning, but inventory cost is the binding consideration in the maker's quoting.*

Adverse-selection-dominated regime *($\lambda > \lambda_\times$). The adverse-selection spread exceeds the inventory spread. As $\lambda \to \infty$, the spread converges to $2\kappa_{\text{adv}}\pi\sigma^2$, the pure adverse-selection spread.*

Axiom dependency. Theorem 9.2 relies on the core (Layer I) together with the structural choice (A3).

Proof. The frozen-functioning boundary is given by Proposition 9.1. The inventory-versus-adverse-selection boundary is $\kappa_{\text{inv}}/\lambda = \kappa_{\text{adv}}\pi$, which solves to $\lambda = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi) = \lambda_\times$. The ordering condition $\kappa_{\text{adv}}\pi < \delta_{\max}/(4\sigma^2)$ is equivalent to $\lambda^*(\pi) < \lambda_\times$, since

$$\begin{aligned}
\lambda^*(\pi) < \lambda_\times &\iff \frac{\kappa_{\text{inv}}}{\delta_{\max}/(2\sigma^2) - \kappa_{\text{adv}}\pi} < \frac{\kappa_{\text{inv}}}{\kappa_{\text{adv}}\pi} \\
&\iff \kappa_{\text{adv}}\pi < \frac{\delta_{\max}}{2\sigma^2} - \kappa_{\text{adv}}\pi \\
&\iff \kappa_{\text{adv}}\pi < \frac{\delta_{\max}}{4\sigma^2}.
\end{aligned}$$

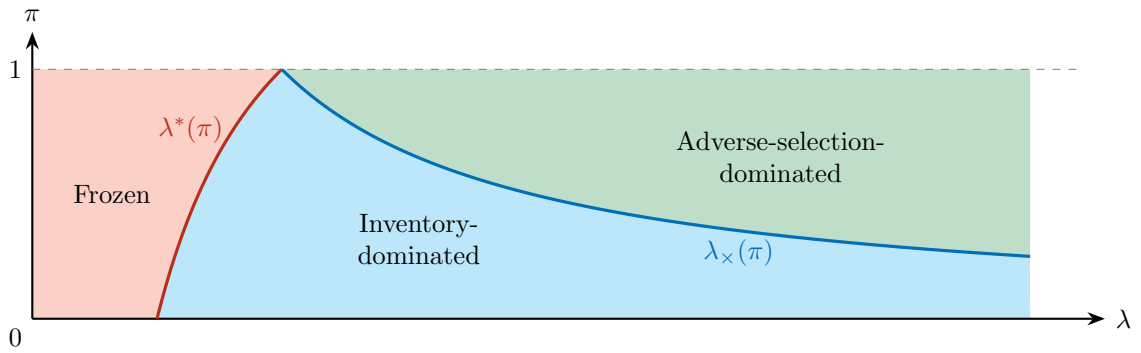


Figure 3: Phase diagram in (λ, π) space. The frozen region $\{\lambda < \lambda^*(\pi)\}$ (salmon) lies behind the phase boundary $\lambda^*(\pi) = \kappa_{\text{inv}}/(\delta_{\max}/(2\sigma^2) - \kappa_{\text{adv}}\pi)$. The functioning region partitions into inventory-dominated $\{\lambda^*(\pi) < \lambda < \lambda_\times(\pi)\}$ (blue) and adverse-selection-dominated $\{\lambda > \lambda_\times(\pi)\}$ (green), with crossover $\lambda_\times(\pi) = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$. Illustrative values $\kappa_{\text{inv}} = \kappa_{\text{adv}} = \sigma^2 = 1$, $\delta_{\max} = 4$. Under failure of the ordering condition (Remark 9.3), the inventory-dominated region is empty.

Under this ordering condition, the two thresholds partition $\mathbb{R}_{++}$ into the three intervals stated. The asymptotic identity $\lim_{\lambda\to\infty} \delta = 2\kappa_{\text{adv}}\pi\sigma^2$ follows by direct substitution into (2). □

Remark 9.3 (The Intermediate Adverse-Selection Regime)**.** When the ordering condition $\kappa_{\text{adv}}\pi < \delta_{\max}/(4\sigma^2)$ fails but the phase condition still holds, the crossover liquidity $\lambda_\times$ lies in the frozen region. At $\lambda = \lambda_\times$ the spread equals $4\kappa_{\text{adv}}\pi\sigma^2 \geq \delta_{\max}$. The inventory-dominated interval is empty, and the decomposition degenerates to two regimes, namely frozen $(\lambda < \lambda^*(\pi))$ and adverse-dominated $(\lambda > \lambda^*(\pi))$.

Edge case. The crossover liquidity $\lambda_\times$ is defined only for $\pi > 0$. At $\pi = 0$, the adverse-selection spread vanishes and the decomposition collapses to a single functioning regime, with the phase condition reducing to $\kappa_{\text{inv}}/\lambda < \delta_{\max}/(2\sigma^2)$.

Empirical markets can be classified by regime, and the boundaries are sharp. Increases in σ^2, π, or κ_{inv}, or decreases in λ, all push the market toward the frozen boundary. Figure 3 shows the partition.

Remark 9.4 (Market Freezes as Phase-Boundary Crossings)**.** The phase boundary admits a structural reading consistent with empirically observed market freezes. During the 2008 financial crisis and several subsequent stress episodes, trading ceased in a number of markets despite the apparent presence of willing makers, a pattern difficult to rationalize within standard optimization models since each maker individually has positive expected profit from quoting. Within the axiomatic framework, such freezes correspond to crossings of the phase boundary (3): a spike in volatility, adverse selection, or inventory cost (or a collapse in arrival intensity) can push the implied spread beyond $\delta_{\max}$, eliminating the willing-counterparty population. The maker's individual profitability is not the issue. The aggregate market clears only when the spread is below the maximum disagreement, and that condition can fail even when makers are profitable in expectation.

9.4 Information-Minimizing Liquidity

The crossover liquidity $\lambda_\times$ admits a second characterization in information-theoretic terms, under explicit auxiliary modeling. Adopt the following modeling premises, used solely for the per-trade information content: (M1) the maker's prior over v is $\mathcal{N}(\mu, \sigma^2)$; (M2) a counterparty trade at the ask a (respectively bid b) is interpreted as evidence that the counterparty's reservation value exceeds a (respectively falls below b), inducing a posterior update on v; (M3) the Kullback-Leibler divergence between posterior and prior is proportional to $(a-\mu)^2/\sigma^2$ to leading order in $(a-\mu)/\sigma$, with constant $\kappa_{\text{info}} > 0$ depending on the update model but not on (λ, σ^2).

Under (M1)–(M3), one trade conveys $\kappa_{\text{info}}(a-\mu)^2/\sigma^2$ to leading order. At zero inventory and zero skew, $a - \mu = \delta/2$, giving information rate

$$R_{\text{info}}(\lambda) = \kappa_{\text{info}}\, \lambda\sigma^2 \Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi\Big)^2.$$

Theorem 9.5 (Minimum Information Rate at the Crossover)**.** *Under (M1)–(M3), the information rate $R_{\text{info}}(\lambda)$ attains a unique stationary point at $\lambda = \lambda_\times = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$. At $\lambda_\times$ the two spread components are equal, the stationary point is a minimum, and R_{info} is U-shaped on $\mathbb{R}_{++}$. The location of the minimum is independent of κ_{info}.*

Axiom dependency. The proof follows from the axiomatic spread formula together with (M1)–(M3). The coincidence with $\lambda_\times$ is a feature of the combined setup, not a free-standing consequence of the axioms.

Proof. Write $R_{\text{info}}(\lambda) = \kappa_{\text{info}}\sigma^2 f(\lambda)$ with $f(\lambda) = \lambda(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)^2$. Substituting $u = \sqrt{\lambda}$ gives $f = (\kappa_{\text{inv}}/u + \kappa_{\text{adv}}\pi u)^2$. The bracket is minimized by AM-GM at $u^* = \sqrt{\kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)}$ and diverges as $u \to 0^+$ and $u \to \infty$. In λ coordinates, the minimum is located at $\lambda_\times = (u^*)^2 = \kappa_{\text{inv}}/(\kappa_{\text{adv}}\pi)$. □

The crossover liquidity, identified by structural regime classification in Theorem 9.2, is also the least informative liquidity under (M1)–(M3). A different information model would generically place the minimum elsewhere. The leading-order approximation in (M3) breaks down where the spread approaches $\delta_{\max}$, so empirical implications should be read as conditional on (M1)–(M3).

10 Comparative Statics

The closed-form characterization makes the comparative statics immediate. Each follows by direct differentiation, and the interpretive content is in the structural reading of signs and magnitudes.

10.1 First-Order Derivatives, Cross-Partials, and Elasticities

The closed-form characterization makes all partial derivatives of m and δ immediate. Table 2 collects them together with the nonzero cross-partials and the leading elasticities.

The mid-quote tracks belief one-for-one and slopes downward in inventory at rate γ_0, autonomous from the trading environment conditional on (μ, q); the volatility-interacted

Quantity	$\partial/\partial q$	$\partial/\partial\mu$	$\partial/\partial\sigma^2$	$\partial/\partial\lambda$	$\partial/\partial\pi$
m	$-\gamma_0\sigma^2$	1	$-\gamma_0 q$	0	0
δ	0	0	$2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$	$-2\kappa_{\text{inv}}\sigma^2/\lambda^2$	$2\kappa_{\text{adv}}\sigma^2$

Cross-partial	Value	Elasticity	Value
$\partial^2 m/\partial q\,\partial\sigma^2$	$-\gamma_0$	$\epsilon_{\delta,\lambda}$	$-(\kappa_{\text{inv}}/\lambda)/(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$
$\partial^2\delta/\partial\sigma^2\,\partial\lambda$	$-2\kappa_{\text{inv}}/\lambda^2$	$\epsilon_{\delta,\pi}$	$(\kappa_{\text{adv}}\pi)/(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$
$\partial^2\delta/\partial\sigma^2\,\partial\pi$	$2\kappa_{\text{adv}}$	$\epsilon_{\delta,\sigma^2}$	1
$\partial^2\delta/\partial\lambda\,\partial\pi$	0	$\epsilon_{m-\mu,q}$	1

Table 2: Comparative statics of the characterized quoting rule under linear inventory cost (C5). Relaxing C5 to general convex κ gives $\partial\delta/\partial q = 2\sigma^2\kappa''(|q|)/\lambda \geq 0$.

cross-partials identify γ_0 and κ_{adv} from separate regressions, matching Corollary 6.3; the vanishing $\partial^2\delta/\partial\lambda\partial\pi$ is a direct consequence of additive separability and yields a testable cross-restriction; the liquidity elasticity $\epsilon_{\delta,\lambda}$ equals -1 in the inventory-dominated regime, approaches 0 in the adverse-selection-dominated regime, and crosses $-1/2$ at $\lambda_\times$, localizing the market within the three-regime structure of Section 9; and the unit elasticities $\epsilon_{\delta,\sigma^2}$ and $\epsilon_{m-\mu,q}$, direct consequences of (A3) and (A2), signal axiom violations when they depart from unity.

10.2 Comparison with Existing Models

Ho-Stoll and Avellaneda-Stoikov each predict $\delta = A\sigma^2 + B$ for some structural constants. Glosten-Milgrom predicts $\delta \propto \pi\sigma$ with σ-scaling (the $\beta = 1/2$ neighbor of our family via (A3s)). Our characterization predicts $\delta = 2\sigma^2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$, separating the variance coefficient into an inventory part and an adverse-selection part. In Ho-Stoll the inventory cost enters both spread and skew, and in Avellaneda-Stoikov the risk-aversion parameter enters both components. Neither admits parameter separation from a single regression. The axiomatic decomposition resolves this conflation. The parameter γ_0 is identified from $\partial m/\partial q$, κ_{inv} from the spread at $\pi = 0$, and κ_{adv} from $\partial\delta/\partial\pi$.

10.3 An Empirical Implementation Roadmap

The identification corollary gives the formal mapping from observable quoting behavior to the three parameters. We describe a step-by-step implementation, assuming quote-level data for one or more makers together with inventory and state variables.

Data requirements. Time-stamped observations of the maker's bid and ask, her inventory q, the public belief μ (typically the consolidated mid-quote or a Bayesian estimate), the price variance σ^2 (exponentially-weighted realized variance), the trade arrival intensity λ, and an informed-fraction proxy π (PIN or VPIN).

Step 1, identification of γ_0. Run $m_t - \mu_t = -\gamma_0 \sigma_t^2 q_t + \epsilon_t$, and the slope on $\sigma_t^2 q_t$ identifies γ_0. Include controls for λ, π. Under the axioms these should have zero coefficients (by (A5), (A6)), and a nonzero coefficient signals a violation.

Step 2, identification of κ_{inv}. Restrict to low informed flow (π below the 25th percentile) and run $\delta_t = 2\kappa_{\text{inv}} \sigma_t^2 / \lambda_t + \epsilon_t$. The slope on σ_t^2/λ_t identifies κ_{inv}.

Step 3, identification of κ_{adv}. Run $\delta_t = 2\kappa_{\text{inv}} \sigma_t^2/\lambda_t + 2\kappa_{\text{adv}} \sigma_t^2 \pi_t + \epsilon_t$. The slope on $\sigma_t^2 \pi_t$ identifies κ_{adv}, and the two slopes are jointly identified.

Step 4, recovery of the cost function. If the order book is observable beyond the inside, apply Theorem 7.3. At each level n, the cost gradient $\kappa'(q+n-1)$ is the residual of the n-th ask after subtracting skew and adverse-selection components. Integrate across n to recover κ.

Step 5, diagnostic tests. The axiom system carries testable cross-restrictions as direct consequences: unit elasticity of δ in σ^2 from (A3), no effect of λ, π on the mid-quote from (A5), (A6), and linearity of the spread in π from (A8). A failure signals violation of the corresponding axiom.

Pitfalls. Three issues deserve explicit treatment. *Attenuation bias in κ_{adv}.* The informed fraction π is unobservable and must be proxied (PIN, VPIN, imbalance estimators), each noisy, so classical measurement error attenuates the OLS estimate of κ_{adv} toward zero by the factor $\sigma_\pi^2/(\sigma_\pi^2 + \sigma_e^2)$. The defensible remedy is instrumental variables using lagged proxies or order-flow imbalance shocks. *Instrument quality for κ_{inv}.* Identification requires variation in λ exogenous to (q, σ^2, π). In observational data, λ correlates with σ^2 and π, so candidate instruments include exchange-system events, trading-calendar effects, and structural shifts in maker-rebate regimes. *Boundary-regime nonlinearity.* The linearity assumptions underlying the identification regressions are most questionable near the phase boundary, where the spread approaches $\delta_{\max}$. A practical resolution is to estimate parameters on interior data and test the predicted boundary location out of sample.

10.4 Numerical Illustration on Simulated Quote Data

We illustrate the identification procedure on synthetic data drawn from the characterized rule. The illustration is not a substitute for an empirical study but a transparent demonstration that the three regressions of Section 10.3 recover the true primitives under the maintained axioms.

Data-generating process. We fix true parameters $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) = (0.5, 0.2, 0.1)$ and draw $T = 10{,}000$ independent states. Inventory q is uniform on $[-10, 10]$, belief μ is Gaussian with mean 100 and variance 1, belief variance σ^2 is uniform on $[0.05, 1.0]$, arrival intensity λ is shifted log-normal on $[0.5, \infty)$, and informed fraction π is uniform on $[0, 0.5]$. The mid-quote and spread are generated from (1) plus small Gaussian noise (standard deviation 0.002 on the mid, 0.005 on the spread) representing measurement and discretization error.

Identification results. Table 3 reports estimates from the three regressions of Section 10.3. Step 1 regresses $m - \mu$ on $-\sigma^2 q$ through the origin to recover γ_0. Step 2 restricts to π below the 25th percentile (so $\kappa_{\text{adv}}\pi$ is small) and regresses δ on $2\sigma^2/\lambda$ to recover κ_{inv}. Step 3 runs the joint regression $\delta = 2\kappa_{\text{inv}}\sigma^2/\lambda + 2\kappa_{\text{adv}}\sigma^2\pi + \epsilon$ on the full sample. Each parameter is recovered to within Monte Carlo error of its true value.

Step	Parameter	True	Estimate (SE)	R^2
1	γ_0	0.500	0.5000 (0.0000)	1.000
2	κ_{inv}	0.200	0.2081 (0.0002)	0.995
3	κ_{inv}	0.200	0.1999 (0.0001)	0.998
	κ_{adv}	0.100	0.1004 (0.0003)	

Table 3: Identification estimates from $T = 10{,}000$ simulated states. Standard errors in parentheses. Step 1 uses regression through the origin; Step 2 uses the π below the 25th-percentile subsample ($n = 2{,}500$); Step 3 uses the full sample.

Cross-restrictions. The axioms imply that λ and π should not enter the mid-quote regression (by (A5), (A6)). Augmenting the Step 1 regression with both yields coefficients 2.8×10^{-5} on λ and -1.2×10^{-4} on π, both within sampling error of zero. A nonzero coefficient on either would falsify the corresponding axiom.

Recovery of the cost function. We illustrate Theorem 7.3 under a quadratic cost specification, $\kappa(q) = \frac{1}{2}\kappa_{\text{inv}} q^2$ with $\kappa_{\text{inv}} = 0.2$, generating a ten-level order book at the fixed state $(\mu, \sigma^2, \lambda, \pi, q) = (100, 0.5, 1.0, 0.2, 0)$. For each level n, we recover the cost gradient via

$$\widehat{\kappa'}_+(n-1) = \frac{\lambda}{\sigma^2}\big(a_n - \mu + \gamma_0\sigma^2(n-1) - \kappa_{\text{adv}}\pi\sigma^2\big).$$

The recovered gradients track the true linear gradient $\kappa'_+(x) = 0.2x$ to within observation noise. Summing across levels gives the left-endpoint Riemann estimate $\widehat{\kappa}(10) = 9.00$, against the true value $\kappa(10) = 10.00$. The Riemann sum underestimates by exactly the amount predicted in Theorem 7.3: for convex κ, the left-endpoint sum is a lower bound on the integral. Under the linear specification C5, the underestimate disappears.

11 The Skew-Spread Ratio and Critical Inventory

The full characterization yields a dimensionless quantity, the skew-spread ratio, classifying the maker's positioning into structural regimes.

Theorem 11.1 (Skew-Spread Ratio)**.** *Under the Full Characterization, the ratio of skew magnitude $\gamma_0\sigma^2|q|$ to half-spread $\sigma^2(\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi)$ is*

$$R(q, \lambda, \pi) = \frac{\gamma_0|q|}{\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi},$$

independent of σ^2.

Axiom dependency. Theorem 11.1 relies on the full nine-axiom system.

Proof. Direct from the closed form, with σ^2 canceling. □

The ratio R is a dimensionless summary statistic of the maker's positioning. The critical inventory q^* is defined by $R(q^*, \lambda, \pi) = 1$, namely

$$q^* = \frac{\kappa_{\text{inv}}/\lambda + \kappa_{\text{adv}}\pi}{\gamma_0}.$$

For $|q| < q^*$, the half-spread exceeds the skew, and the maker operates in the *spread-dominated regime*, with her quoting conservative in inventory positioning and revenue coming primarily from the bid-ask spread. For $|q| > q^*$, the skew exceeds the half-spread, and the maker enters the *skew-dominated regime*, with the skew adjustment becoming the principal mechanism of inventory management.

Remark 11.2 (Regime Classification by q^*)**.** The critical inventory q^* is dependent on $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}})$ and (λ, π) but not on σ^2. At $\pi = 0$, $q^* = \kappa_{\text{inv}}/(\gamma_0\lambda)$ scales as $1/\lambda$: more liquid markets cross the regime earlier at smaller inventory. At $\lambda \to \infty$, $q^* = \kappa_{\text{adv}}\pi/\gamma_0$ is determined by adverse selection alone. At $q = 0$, $R = 0$ and quoting is purely spread-driven. The band $|q| \le q^*$ has an operational reading as the maker's working inventory range under business-as-usual quoting, with excursions beyond pushing the rule into the skew-dominated regime.

12 The Bertrand Equivalence

The axiomatic theory derives the quoting rule from indifference axioms on a single maker. A natural question is whether the resulting formula depends essentially on this monopolistic-indifference foundation, and we show that it does not. The quoting rule (1) arises also as the unique symmetric Nash equilibrium of a Bertrand-competitive market with free entry of homogeneous makers, in the framework of Bertrand (1883) treated in standard form in Tirole (1988) and Vives (1999).

12.1 The Competitive Market

Throughout Section 12 the Bertrand setting is fixed by six conditions, namely $N \ge 2$ homogeneous makers with cost parameters $(\kappa_{\text{inv}}, \kappa_{\text{adv}}, \gamma_0)$ and instantaneous inventory disutility $U(q) = \frac{1}{2}\gamma_0\sigma^2q^2$, continuous prices and quantities, no capacity constraints, uniform random tie-breaking, free entry at the stage-game zero-profit level, and a per-state stage-game equilibrium concept. Dynamic evolution across repeated stages lies outside the scope of this section. The quadratic inventory disutility is a modeling choice within the Bertrand framework, not a consequence of the axiom system, and it matches the standard inventory-management form (Avellaneda-Stoikov, Cartea-Jaimungal). Each maker simultaneously posts $(b_i, a_i) \in \mathcal{B}$, and trades execute at the best quote.

12.2 The Bertrand Equivalence Theorem

The standard Bertrand argument predicts that competition drives prices to marginal cost, here the per-trade carrying cost plus the expected adverse-selection loss. The equilibrium prices coincide with the axiomatic quoting rule in the continuous-quantity limit, with an $O(\Delta)$ discretization artifact in the inventory-skew adjustment that vanishes as $\Delta \to 0$.

Theorem 12.1 (Bertrand Equivalence)**.** *Under the setting of Section 12, the unique symmetric Nash equilibrium of the Bertrand-competitive market is the quoting rule of Theorem 4.1,*

$$b = \mu - \gamma_0\sigma^2 q - \sigma^2\Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi\Big), \quad a = \mu - \gamma_0\sigma^2 q + \sigma^2\Big(\frac{\kappa_{\text{inv}}}{\lambda} + \kappa_{\text{adv}}\pi\Big).$$

The spread structure is derived from zero-profit competition independently of the disutility specification. The skew $-\gamma_0\sigma^2 q$ depends on the chosen quadratic form, with a different inventory disutility yielding a different skew.

Axiom dependency. The Bertrand derivation itself does not invoke the axiom system, and the equivalence to the closed-form rule of Theorem 4.1 requires the full nine-axiom system on the axiomatic side.

Proof. We prove the theorem in three steps, namely identification, existence, and uniqueness among symmetric equilibria.

Step 1, identification. At a symmetric equilibrium, each maker posts the same quote (b, a) and faces a fraction $1/N$ of total flow. The per-trade economic margin at the ask decomposes into three components: the per-trade carrying cost $\kappa_{\text{inv}}\sigma^2/\lambda$, the expected adverse-selection loss $\kappa_{\text{adv}}\pi\sigma^2$, and the marginal change in inventory disutility.

With instantaneous disutility $U(q) = \frac{1}{2}\gamma_0\sigma^2 q^2$, a trade at the ask reduces inventory from q to $q - \Delta$, giving $U(q-\Delta) - U(q) = -\gamma_0\sigma^2\Delta q + \frac{1}{2}\gamma_0\sigma^2\Delta^2$ and per-unit reduction $\gamma_0\sigma^2 q + O(\Delta)$. Taking $\Delta \to 0$, the residual $\frac{1}{2}\gamma_0\sigma^2\Delta$ vanishes and the per-unit reduction is $\gamma_0\sigma^2 q$, positive when $q > 0$ (the trade dumps unwanted inventory) and negative when $q < 0$.

Bertrand competition drives the per-trade margin to zero. A maker undercutting by $\epsilon > 0$ captures all the buy-side flow but forgoes ϵ per trade. The undercut is profitable if and only if the per-trade margin at a is strictly positive, so at equilibrium the margin must vanish:

$$\underbrace{(a-\mu)}_{\text{revenue net of mid}} = \underbrace{\frac{\kappa_{\text{inv}}\sigma^2}{\lambda}}_{\text{carrying}} + \underbrace{\kappa_{\text{adv}}\pi\sigma^2}_{\text{adverse selection}} - \underbrace{\gamma_0\sigma^2 q}_{\text{inventory-disutility reduction}}.$$

Solving and applying the symmetric reasoning on the bid side:

$$a = \mu - \gamma_0\sigma^2 q + \frac{\kappa_{\text{inv}}\sigma^2}{\lambda} + \kappa_{\text{adv}}\pi\sigma^2, \qquad b = \mu - \gamma_0\sigma^2 q - \frac{\kappa_{\text{inv}}\sigma^2}{\lambda} - \kappa_{\text{adv}}\pi\sigma^2,$$

matching (1). The risk-premium term $\gamma_0\sigma^2 q$ emerges from the quadratic disutility together with the zero-margin condition. A different functional form for U would yield a different skew, while the spread structure is determined by zero-profit competition alone.

Step 2, existence. At the stated quotes each maker earns zero expected profit at the trading margin. Undercutting raises the fill share from $1/N$ to one but reduces the per-trade

margin to negative, since the equilibrium half-spread equals the per-trade marginal cost. Overshooting reduces the fill share to zero. The stated quotes constitute a Nash equilibrium.

Step 3, uniqueness among symmetric equilibria. Suppose (b', a') is a symmetric equilibrium with $a' \neq a$, and let

$$\mathrm{MC}(q) = \kappa_{\mathrm{inv}}\sigma^2/\lambda + \kappa_{\mathrm{adv}}\pi\sigma^2 - \gamma_0\sigma^2 q.$$

If $a' > a$, posting $a'' = a' - \epsilon$ for small $\epsilon > 0$ captures the full buy-side flow (rather than the $1/N$ share at a') with strictly positive per-trade margin $a' - a - \epsilon > 0$. For small enough ϵ, the full-flow profit exceeds the $1/N$-share profit at a' since $N \geq 2$, contradicting equilibrium. Continuous prices ensure that no minimum tick obstructs the choice of ϵ. If $a' < a$, the per-trade margin is strictly negative, so withdrawing earns zero expected profit, also a profitable deviation. Hence, $a' = a$, and symmetrically $b' = b$.

The argument extends to free entry as $N \to \infty$. The equilibrium condition is per-trade rather than flow-based, so it is unaffected by N. In the free-entry limit, each maker's $1/N$ share tends to zero but the equilibrium quotes are unchanged. □

The Bertrand Equivalence has a structural reading. The same closed-form quoting rule arises from two structurally distinct economic models, a monopolistic maker subject to indifference axioms and many competitive makers subject to free entry. Axioms (A4), (A8), (A9) each encode a single-maker break-even condition that carries the structural content of competitive zero-profit pricing. The indifference axioms are competitive logic written in single-maker language.

12.3 The Profit Ceiling

With the spread driven to marginal cost by competitive entry, the per-trade economic profit collapses to zero, and the spread becomes a pure cost-recovery instrument.

Corollary 12.2 (Profit Ceiling)**.** *Under the Full Characterization, the following hold.*

(i) The maker's expected per-trade economic profit is zero, by the two break-even axioms (A4) and (A8).

(ii) The deterministic adverse-selection component of the spread, $2\kappa_{\mathrm{adv}}\pi\sigma^2$, is bounded above by $2\kappa_{\mathrm{adv}}\sigma^2$, attained at the maximum informed-trader fraction $\pi = 1$.

The realized per-trade profit on a single fill against an informed trader depends on the unobserved $v - \mu$ and is not bounded by any function of the structural parameters. Only its expectation is.

Axiom dependency. Corollary 12.2 relies on (A4) and (A8), the two break-even axioms of the indispensable core.

Proof. Statement (i) is the conjunction of two break-even conditions. Axiom (A4) equates the inventory half-spread to marginal carrying cost per trade, and (A8) equates the adverse-selection half-spread to the expected information rent per informed trade. Statement (ii) is immediate from the form of $2\kappa_{\mathrm{adv}}\pi\sigma^2$ and the bound $\pi \in [0, 1]$. □

In expectation, the maker operates as a cost-covering service provider rather than a profit extractor. The bid-ask spread is not a margin retained as economic rent. It compensates the maker exactly for the marginal cost of carrying inventory and the expected information rent extracted by informed counterparties. Empirical work inherits a clean reading. Estimates of dealer profitability based on quoted spreads recover cost components rather than margin, and persistent positive economic rent must be attributed to barriers to entry or to costs outside the structural specification.

13 Structural Invariants of the Axiomatic Family

The axioms in $\mathcal{A}$ identify a three-parameter family of quoting rules. Different choices of structural primitives (cost function κ, variance scaling exponent, equivariance generator) would yield different families. We close the structural development with a meta-theorem identifying what is preserved across all such choices.

The class of *admissible structural primitives* consists of triples (κ, h, β). Here $\kappa\colon \mathbb{R} \to \mathbb{R}_+$ is convex nondecreasing in $|q|$ with $\kappa(0) = 0$ and symmetric in q (the cost-structure conditions C1, C2, C3, C6), $h\colon \mathcal{D}_h \to \mathbb{R}$ is a strictly increasing differentiable map from an admissible price domain (the equivariance generator), and $\beta \in (0, 1]$ is the variance-scaling exponent of (A3) or (A3s). The canonical choice is $(\kappa(q) = \kappa_{\text{inv}}|q|, h = \text{id}_{\mathbb{R}}, \beta = 1)$.

Theorem 13.1 (Invariant Features)**.** *Under (A1)–(A9), or the (A3s) variant with $\beta \in (0, 1]$, any admissible structural primitive (κ, h, β) yields a quoting rule satisfying the following four invariant features.*

1. Inventory half-spread determined by κ'. *The inventory half-spread at level q equals $\sigma^2\kappa'_+(|q|)/\lambda$, a continuous monotone function of q determined by the cost derivative. Under linear cost (C5), this reduces to the constant $\kappa_{\text{inv}}\sigma^2/\lambda$.*

2. Additive spread decomposition. *The spread satisfies $\delta = \delta_{\text{inv}} + \delta_{\text{adv}}$, with δ_{inv} depending on $(q, \sigma^2, \lambda, \kappa)$ and δ_{adv} depending on (σ^2, π, β).*

3. Skew decoupling from π. *The inventory hedging does not depend on the informed-trader fraction.*

4. Variance homogeneity of each component. *Each component of the quoting rule is homogeneous of some degree in σ^2. The degree is one in the canonical case $\beta = 1$, and mixed degrees one and β under (A3s).*

Axiom dependency. Theorem 13.1 relies on (A1)–(A9) and is meta-stated across all admissible structural-primitive choices (κ, h, β).

Proof. Property 1. Relaxing C5 to general convex nondecreasing κ and applying (A4) at each inventory level (via the order-book argument of Theorem 7.2) gives the level-q inventory half-spread $\sigma^2\kappa'_+(|q|)/\lambda$. Under linear cost (C5), $\kappa'_+ \equiv \kappa_{\text{inv}}$ is constant on $\mathbb{R}_+$, and the inventory half-spread reduces to $\kappa_{\text{inv}}\sigma^2/\lambda$.

Property 2. The two distinct contributions to the half-spread are identified separately, namely the inventory contribution by (A4), and the adverse-selection contribution by (A8)

through the increment $h(\sigma^2, \pi) = \delta(q, \mu, \sigma^2, \lambda, \pi) - \delta(q, \mu, \sigma^2, \lambda, 0)$. The two contributions, derived separately in Lemmas 4.5 and 4.6, combine additively. A different cost function changes δ_{inv} but leaves δ_{adv} untouched. A different scaling exponent (A3s with $\beta \neq 1$) changes the variance-homogeneity degree of δ_{adv} but leaves the additivity untouched.

Property 3. The proof of Lemma 4.3 uses (A1), (A2), (A3), (A5), and (A6), none of which constrains π. (A6) in particular makes the skew coefficient γ depend only on σ^2 (and through (A3), on σ^2 via $\gamma_0\sigma^2$), not on π. The skew is determined by inventory and volatility alone, with the trading environment (λ, π) entering only through the spread.

Property 4. Under (A3) (with $\beta = 1$), each component is homogeneous of degree one in σ^2. Under (A3s) with $\beta \in (0, 1]$, the skew and inventory components have degree 1 and the adverse-selection component has degree β; under (A3), all three coincide at degree 1. □

Remark 13.2 (On Multiplicative Price Separability)**.** A natural further candidate for an invariant feature is *multiplicative price separability*. Under any admissible equivariance generator h, the μ-dependence of the quoting rule should factor through h. In the paper's additive setting ($h = \text{id}_{\mathbb{R}}$), this follows from Theorem 4.9. Under the logarithmic alternative $h(\mu) = \log \mu$ on $\mathbb{R}_{++}$, additivity in h becomes multiplicativity in μ. The mid-quote takes the form $m(s) = \mu \cdot \exp(\psi(\cdot))$, and the ratios $a(s)/\mu$ and $\mu/b(s)$ depend only on the non-μ arguments. We do not state this as an invariant feature, since doing so requires a precise specification of the h-coordinate version of (A2), (A4), (A8), outside our scope. The equivariance content of (A2), (A4), (A8) is naturally restated in h-coordinates, and the h-equivariance follows there by the same argument. A precise statement remains for future work.

The Invariant Features Theorem isolates the substantive content of axiomatic market making from matters of structural convention: any quoting rule within the axiomatic family must exhibit inventory half-spread determined by κ', additive spread decomposition, skew decoupling from π, and per-component variance homogeneity, regardless of which admissible structural primitive (κ, h, β) is chosen. Violations of any of these four features place the model outside the axiomatic family entirely. The convention-dependent features (skew form under (C5, linearity) versus general convexity, variance-scaling exponent σ^2 versus σ, additive versus multiplicative price structure) reflect choices an applied theorist makes for a specific market and carry no invariant status; the four invariant features, by contrast, are not imposed separately but are predicted by the axioms themselves.

13.1 Existing Models as Structural Neighbors

None of the standard models is a strict special case of our axiomatic family. Each violates at least one axiom or environmental assumption. What Theorem 13.1 provides is the broader axiomatic universe within which the standard models sit, each at a structurally identifiable position. We summarize the structural commitments of the four canonical models alongside our own in Table 4, then treat each in turn.

Ho-Stoll. The Ho-Stoll model (Ho and Stoll, 1981) violates (C4, autonomy) through its time-to-horizon dependence, and converges in the stationary limit to a point in our family with $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}}) = (\alpha, \alpha/2, 0)$.

Structural commitment	Ours	Ho-Stoll	A-S	C-J	G-M
Closed-form spread	yes	yes	yes	yes	yes
Inventory channel	yes	yes	yes	yes	no
Adverse-selection channel	yes	no	no	no	yes
Steady-state λ	yes	no	no	no	yes
σ^2-scaling of spread	yes	yes	partial	yes	no (σ)
Linear inventory cost (C5)	yes	n/a	no (quad.)	no (quad.)	n/a
Inventory aversion linear in q	yes	yes	yes	yes	n/a
Skew autonomous of λ, π	yes	no	no	no	n/a
Cost function observable	yes	no	no	no	n/a
Number of structural parameters	3	2	3	4	1

Table 4: Structural commitments of our axiomatic family compared with the canonical models of Ho-Stoll, Avellaneda-Stoikov (A-S), Cartea-Jaimungal (C-J), and Glosten-Milgrom (G-M). A 'yes' indicates the commitment is satisfied; 'no' that it is violated or replaced by an incompatible one; 'partial' that the property holds for the spread's leading term but fails for a residual term; 'n/a' that the model lacks the channel that would make the question meaningful. Each model violates at least one commitment of our family, placing it in a structural neighborhood rather than inside it. The number of identified structural parameters refers to free coefficients of the equilibrium quoting rule, not to all primitives of the underlying optimization problem.

Avellaneda-Stoikov. The Avellaneda-Stoikov model (Avellaneda and Stoikov, 2008) extends Ho-Stoll with exponential utility and Poisson arrivals, and its logarithmic term $(2/\gamma)\ln(1+\gamma/k)$ violates (A3); the large-k limit approximates a point in our family.

Cartea-Jaimungal. The Cartea-Jaimungal framework (Cartea et al., 2015) adds quadratic inventory penalties, violating (C5, linearity), and retains the Ho-Stoll time-to-horizon dependence; the axiomatic foundation for the Avellaneda-Stoikov and Cartea-Jaimungal objectives on the preference side is developed in the companion paper (Feys, 2026), connected to the present paper by the Legendre-Fenchel duality of Section 8.

Glosten-Milgrom. The Glosten-Milgrom model (Glosten and Milgrom, 1985) derives a spread proportional to $\pi\sigma$ with no inventory channel, violating (A3) and lying in the neighboring (A3s) family with $\beta = 1/2$.[1]

These models are unified not by membership in a common axiomatic family but by their place within the broader axiomatic universe. Four features are shared across all such families, namely skew determined by κ', additive spread decomposition, skew decoupled from π, and per-component homogeneity in σ^2. Each existing model is the unique closed-form solution of some axiom system within the universe, and the four invariants are forced at the meta-level rather than imposed model by model.

[1]The Kyle model (Kyle, 1985) sits further from our family, since it combines σ-scaling with an endogenous informed-flow intensity and continuous-time strategic dynamics; it occupies a distinct axiomatic neighborhood within Theorem 13.1.

14 Relation to the Axiomatic Tradition

Characterizing an object by a short list of axioms, and reading its structure off the axioms, is a standard move in economic theory. The present application adds the step that matters most for empirical work: the forced functional form is not merely characterized but point-identified from observable moments. The axioms thereby deliver an identification strategy, not only a representation.

14.1 Robustness and Falsifiability

The nine-axiom system implies Theorem 4.1 as a joint conclusion, but the identified structure does not rest on the nine restrictions equally. For each substantive conclusion we ask which minimal subset suffices, the discipline of identifying the core that Arrow (1951) made standard. The three-layer partition consists of an indispensable core (A1), (A4), (A7), (A8), (A9); a structural-choice axiom (A3); and three modularity extensions (A2), (A5), (A6). Figure 4 summarizes the partition.

Layer I *indispensable core*	**Conclusions forced** *with C1–C6*
(A1) Inventory aversion (A4) Competitive break-even (inventory) (A7) Spread monotonicity in π (A8) Competitive break-even (information) (A9) Non-interaction among informed traders	• Additive spread decomposition • Bertrand equivalence (on the spread) • Legendre-Fenchel duality • Qualitative phase boundary
Layer II *structural choice*	**Conclusions added** *with Layer I*
(A3) Variance-proportional risk pricing	• Full quantitative spread formula • Sharp phase boundary • Three-regime decomposition
Layer III *modularity extensions*	**Conclusions added** *with Layers I, II*
(A2) Single inventory-risk attitude (A5) Inventory-pricing autonomy (flow) (A6) Inventory-pricing autonomy (information)	• Closed-form skew $m(s) - \mu = -\gamma_0 \sigma^2 q$ • Canonical bijection with $\mathbb{R}^3_{++}$ • Structural Separability • Recovery Theorem

Figure 4: Layered partition of the axiom system. The nine axioms split into three layers, ordered from the indispensable core (Layer I, darkest) to the most easily deniable modularity extensions (Layer III, lightest). Each layer contributes the structural conclusions listed in its right-hand column, given that the axioms of all inner layers are in force.

Empirical violations localize to specific axiom layers. A violation of the closed-form skew with the spread formula intact is a Layer III deviation, non-σ^2 scaling signals (A3), and q-dependent adverse selection signals (A8). The axiom hierarchy thereby serves as a diagnostic framework, mapping each observed deviation to the responsible axiom.

14.2 The Characterization Tradition in Economic Theory

Axioms that determine a functional form, with the form's structure then read off the axioms, recur across economic theory. The social-choice tradition originates with Arrow (1951), the bargaining tradition with Nash (1950) and Kalai and Smorodinsky (1975), the mechanism-design tradition with Vickrey (1961), Hurwicz (1973), and Myerson (1981), the risk-measurement tradition with Artzner et al. (1999), extended to convex risk measures in Föllmer and Schied (2002) and Föllmer and Schied (2016), and the decision-theoretic tradition with von Neumann and Morgenstern (1944) and Savage (1954). The risk-measurement tradition is the closest relative of the present work. The axioms yield a functional form with a transparent dual representation, and our Legendre-Fenchel duality (Theorem 8.1) plays the analogous role. What distinguishes the present application is the empirical reach of the implied form. In the classical characterizations the axioms deliver a representation, and the parameters are typically taken as given. Here the same axioms separate the parameters into nonoverlapping channels (Theorem 6.1), so each is point-identified from a distinct observable moment. The axiomatic characterization and the identification strategy are, in this application, the same theorem read in two directions.

14.3 Positioning Relative to Optimization Models

The dominant approach in microstructure theory derives the quoting rule from primitive specifications (utility, distribution, equilibrium concept), and the conclusions are sensitive to those specifications. The axiomatic conclusions here are sensitive only to the axioms, which buys three things of direct use to an applied researcher. First, the characterization does not depend on optimization-level modeling choices, since Theorem 4.1 does not depend on whether utility is exponential or logarithmic, or on whether arrivals are Poisson or renewal. Second, the identified parameters do not depend on the equilibrium concept, so the same three moments recover $(\gamma_0, \kappa_{\text{inv}}, \kappa_{\text{adv}})$ whether the maker is read as a single optimizer or as a competitor under free entry, the two reconciled by the Bertrand equivalence. Third, each axiom is independently testable, so the theory furnishes a falsifiable prediction with diagnostic localization. The method also runs in the descriptive direction. Placement of existing market-making models within the family of Theorem 13.1 characterizes each by the axiom system it satisfies, so a comparison across models becomes a comparison across axiom systems. The standard models, namely Ho-Stoll, Avellaneda-Stoikov, Cartea-Jaimungal, Glosten-Milgrom, and Kyle, appear not as competitors but as points within it.

15 Conclusion

We have presented an axiomatic theory of market making. Axioms are imposed on the maker's quoting rule directly, in place of the utility-function, distributional, and equilibrium-concept commitments of the optimization-based tradition, and the structural primitives are read off observable moments of the rule itself. Two payoffs follow.

First, the three economic primitives of market making, namely inventory risk aversion, marginal carrying cost, and adverse-selection cost, are point-identified from three distinct moments of the observed quoting rule. The Structural Separability Theorem (Theorem 6.1)

establishes the decoupling, and the identification corollary (Corollary 6.3) gives an empirical strategy in three separate regressions, with no joint estimation. What the optimization-based literature conflates within a single parameter or two, the axiomatic theory separates into three orthogonal primitives.

Second, the latent inventory cost function, treated in the optimization-based tradition as a hidden parameter to be fitted under each chosen specification, is recoverable from the limit order book depth profile by integration (Theorem 7.3). The cost function is a directly observable object, not a free parameter of the theory. The deeper structural reading (Theorem 8.1) identifies the spread-cost relationship as one half of a Legendre-Fenchel conjugate pair, placing market making within the same convex-conjugate tradition as thermodynamics and consumer theory.

These two payoffs rest on a closed-form characterization (Theorem 4.1) of the quoting rule, in which the mid-quote is linear in inventory and the spread splits additively into inventory and adverse-selection components. Four further results complete the architecture. A canonical bijection between the axiom system and $\mathbb{R}^3_{++}$ (Theorem 5.1) is established with no axiom redundant (Proposition 5.2). A sharp boundary in parameter space determines market viability (Proposition 9.1), and the functioning region partitions into three regimes (Theorem 9.2). The Bertrand Equivalence Theorem 12.1 recovers the same quoting rule from free-entry competition. A meta-level Invariant Features Theorem (Theorem 13.1) identifies features preserved across all admissible structural-primitive choices.

Three directions are left for future work: a dynamic extension replacing the static state with its stochastic counterpart, a multi-asset extension with cross-asset correlation, and a strategic extension replacing κ_{adv} with an equilibrium object derived from optimizing informed traders.

The two main deliverables, namely the three-moment identification strategy and the recovery of the cost function from the order book, apply directly to empirical microstructure and to the practical work of market making.

Declarations

Conflict of interest. The author declares no conflicts of interest.

Funding. No funding was received for this work.

Data availability. No empirical data were used. The numerical illustration in Section 10.4 uses synthetically generated data; no external dataset was analyzed.

References

J. Aczél. *Lectures on Functional Equations and Their Applications*. Academic Press, 1966.

Yakov Amihud and Haim Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53, 1980.

Frank J. Anscombe and Robert J. Aumann. A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34(1):199–205, 1963.

Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

Robert J. Aumann. What is game theory trying to accomplish? In K. J. Arrow and S. Honkapohja, editors, *Frontiers of Economics*. Basil Blackwell, 1985.

Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.

Kerry Back. Insider Trading in Continuous Time. *The Review of Financial Studies*, 5(3):387–409, 1992.

Joseph Bertrand. Théorie mathématique de la richesse sociale. *Journal des Savants*, 67:499–508, 1883.

Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and High-Frequency Trading*. Cambridge University Press, 2015.

Patrick Cheridito, Freddy Delbaen, and Michael Kupper. Dynamic Monetary Risk Measures for Bounded Discrete-Time Processes. *Electronic Journal of Probability*, 11:57–106, 2006.

Gerard Debreu. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Yale University Press, 1959.

Kai Detlefsen and Giacomo Scandolo. Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9(4):539–561, 2005.

Darrell Duffie, Nicolae Gârleanu, and Lasse Heje Pedersen. Over-the-Counter Markets. *Econometrica*, 73(6):1815–1847, 2005.

David Easley and Maureen O'Hara. Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1):69–90, 1987.

David Easley, Nicholas M. Kiefer, Maureen O'Hara, and Joseph B. Paperman. Liquidity, Information, and Infrequently Traded Stocks. *Journal of Finance*, 51(4):1405–1436, 1996.

Frank M. V. Feys. Avellaneda-Stoikov and Cartea-Jaimungal as One Framework: A Forced Uniqueness Theorem for Inventory Market Making. arXiv:2606.01477, 2026. URL https://arxiv.org/abs/2606.01477.

Hans Föllmer and Alexander Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447, 2002.

Hans Föllmer and Alexander Schied. *Stochastic Finance: An Introduction in Discrete Time*. De Gruyter, 4th edition, 2016.

Thierry Foucault, Marco Pagano, and Ailsa Röell. *Market Liquidity: Theory, Evidence, and Policy*. Oxford University Press, 2013.

Marco Frittelli and Emanuela Rosazza Gianin. Putting order in risk measures. *Journal of Banking & Finance*, 26(7):1473–1486, 2002.

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

Christian Genest and James V. Zidek. Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, 1(1):114–135, 1986.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Dealing with the inventory risk: A solution to the market making problem. *Mathematics and Financial Economics*, 7(4):477–507, 2013.

Joel Hasbrouck. *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press, 2007.

Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of Convex Analysis*. Springer, 2001.

Thomas Ho and Hans R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.

Roger D. Huang and Hans R. Stoll. The Components of the Bid-Ask Spread: A General Approach. *Review of Financial Studies*, 10(4):995–1034, 1997.

Leonid Hurwicz. The Design of Mechanisms for Resource Allocation. *American Economic Review*, 63(2):1–30, 1973.

Ehud Kalai and Meir Smorodinsky. Other Solutions to Nash's Bargaining Problem. *Econometrica*, 43(3):513–518, 1975.

Albert S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, 53(6):1315–1336, 1985.

Ananth Madhavan and Seymour Smidt. An Analysis of Changes in Specialist Inventories and Quotations. *Journal of Finance*, 48(5):1595–1628, 1993.

Eric Maskin. Nash Equilibrium and Welfare Optimality*. *Review of Economic Studies*, 66(1):23–38, 1999.

Eric Maskin. Mechanism Design: How to Implement Social Goals. *American Economic Review*, 98(3):567–576, 2008.

Eric Maskin and John Riley. Optimal Auctions with Risk Averse Buyers. *Econometrica*, 52(6):1473–1518, 1984.

Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

John F. Nash, Jr. The Bargaining Problem. *Econometrica*, 18(2):155–162, 1950.

Maureen O'Hara. *Market Microstructure Theory*. Blackwell, 1995.

R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.

Alvin E. Roth. *Axiomatic Models of Bargaining*. Springer, 1979.

Leonard J. Savage. *The Foundations of Statistics*. Wiley, 1954.

Amartya K. Sen. *Collective Choice and Social Welfare*. Holden-Day, 1970.

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

Jean Tirole. *The Theory of Industrial Organization*. MIT Press, 1988.

William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961.

Xavier Vives. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, 1999.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Pierre-Olivier Weill. The Search Theory of Over-the-Counter Markets. *Annual Review of Economics*, 12:747–773, 2020.