Article

Not peer-reviewed version

# Research Note: Continuous-State Dynamic Games with Quantal Response: Strategic Foundations of the Kyle-Ho-Stoll Parity Index in Dealer Markets

Peter R. Williams *

Posted Date: 7 November 2025

doi: 10.20944/preprints202511.0481.v1

Keywords: market microstructure; dealer markets; Kyle model; Ho-Stoll model; strategic equilibrium; inventory management; market fragmentation; adverse selection; dynamic stochastic games; quantal response equilibrium

*Article*

# Research Note: Continuous-State Dynamic Games with Quantal Response: Strategic Foundations of the Kyle-Ho-Stoll Parity Index in Dealer Markets

**Peter R. Williams**

enechain Corporation, Tokyo, Japan; peter.williams@enechain.co.jp

**Abstract**

This paper reexamines the Parity Index framework—– the ratio of inventory pressure to information incorporation in dealer markets —–using strategic equilibrium rather than adaptive heuristics. We constructed a dynamic stochastic game in which dealers simultaneously choose spreads and depths to maximize profits while managing inventory risk, solved via continuous-state dynamic programming and Quantal Response Equilibrium. Despite fundamentally different behavioural foundations, the Parity Index's structural relationships proved robust: fragmentation decouples inventory from prices, non-linear costs decrease the index; and the interaction of fragmentation and non-linear costs determines the market regime boundaries. Although, the strategic behavioural model amplified the effects observed in the adaptive model. This convergence demonstrates that the index captures universal economic forces whose qualitative effects transcend model architecture. Policy implications become contingent on market characteristics: strategic depth management makes moderate competition welfare-improving in liquid markets while preserving consolidation's benefits in thin markets. The findings establish the Parity Index as a behaviourally robust framework and demonstrate that adaptive models provide reliable structural guidance when strategic analysis is computationally prohibitive.

**Keywords:** market microstructure; dealer markets; Kyle model; Ho-Stoll model; strategic equilibrium; inventory management; market fragmentation; adverse selection; dynamic stochastic games; quantal response equilibrium

---

## 1. Introduction

Power futures markets often exhibit low liquidity. Despite billions in annual trading volumes—concentrated primarily in near-term contracts —bid-ask spreads in electricity futures often exceed comparable financial markets by an order of magnitude [1,2]. This illiquidity imposes severe economic costs precisely when these markets are most needed. The global energy transition depends on efficient price discovery and risk management tools for renewable energy investments [3]. Yet the very markets designed to provide these functions can suffer from chronic illiquidity, particularly for contracts extending beyond one year.

The consequences cascade throughout the energy system. Wide spreads increase transaction costs for hedging. Volatile prices complicate investment planning [4]. Limited market depth restricts large participants from managing portfolio risks effectively. Most critically, inadequate liquidity creates barriers to capital formation for clean energy infrastructure. Renewable developers struggle to secure long-term revenue certainty, while utilities cannot efficiently hedge variable generation exposure [5,6]. Industrial consumers face elevated costs for managing price risk.

Market makers represent the standard financial solution to these liquidity challenges. By continuously quoting two-sided markets and absorbing temporary order imbalances, they typically reduce spreads, increase volumes, and improve price efficiency. Recognizing this potential, regulators worldwide are actively developing market maker programs for power futures markets [7]. The Agency for the Cooperation of Energy Regulators (ACER), national authorities across Europe, and the U.S. Commodity Futures Trading Commission are all exploring mechanisms to enhance liquidity. However, these initiatives often import program designs and theoretical frameworks directly from liquid financial markets. This paper addresses this gap by analyzing the effectiveness of standard market maker models in an environment characterized not by information asymmetry, but by structural illiquidity and inventory risk. We ask: under what conditions can market maker programs succeed when their foundational assumptions of competitive pressure and deep order flow are absent?

The Kyle [8] model provides the foundational framework for analyzing market maker effectiveness under information asymmetry. This cornerstone of market microstructure theory elegantly captures how adverse selection affects liquidity provision. Crucially, the model assumes market makers are perfectly competitive (or operate under a market-clearing mandate), leading to a zero-expected-profit condition in equilibrium. This zero-profit assumption, however, is unlikely to hold in illiquid markets where the lack of competition allows market makers to wield significant pricing power. The focus is squarely on how they process private information signals embedded in order flow. The model derives closed-form solutions for optimal market maker pricing when facing potentially informed traders. Kyle's lambda parameter precisely quantifies market depth as the price impact per unit of order flow.

Kyle's theoretical predictions have received extensive empirical validation across liquid, anonymous financial markets. Studies confirm the predicted relationships between order flow and price changes [9]. Research documents the inverse correlation between informed trading intensity and market depth. The model's insights about spread determinants and volume patterns consistently match observed market behaviour. This empirical success has established Kyle's framework as the standard benchmark for analyzing liquidity provision and market maker programs [10,11]. This empirical success, however, is confined to the very markets that power futures are not: deep, anonymous, and information-rich. The structural features of energy markets call the direct applicability of these findings into question.

Analytical work in market microstructure has also explored other drivers of liquidity provision. The model of Ho and Stoll [12] represents a foundational framework focusing on inventory risk as a primary driver of market maker behaviour (see also Amihud and Mendelson [13]). In contrast to the Kyle framework, Ho-Stoll features a market maker acting as a risk manager rather than an information processor; quotes are influenced by the need to manage an unbalanced portfolio. A market maker holding excess inventory will lower its quotes to attract buyers and deter sellers, and vice versa. This introduces a distinct motivation for quoting behaviour and provides a richer, though still highly stylized, model of liquidity provision. Together, the Kyle and Ho-Stoll models represent the canonical benchmarks for information-based and inventory-based market making, respectively.

Beyond these two canonical models, our analysis also draws on related streams of market microstructure theory. The framework of Glosten and Milgrom [14] refines the modeling of adverse selection in sequential trade, while research by Easley and O'Hara [15] highlights the information content revealed by trade size, a critical factor in thin markets. More directly, recent work by Peña and Rodríguez [16] provides the first empirical examination of market maker programs in European electricity futures. Our paper builds on this by providing a theoretical foundation for their findings and exploring the strategic behaviour of market makers, a topic previously analyzed in other thin commodity markets by Antón and Bushnell [17] and Kumar and Seppi [18].

However, power markets violate the assumptions of these canonical models in fundamental ways. These disconnects span two critical dimensions—— physical characteristics and market structure ——that challenge the frameworks' applicability, particularly in thin markets, defined here as markets with a finite number of buyers and sellers.

First, electricity's physical characteristics differ radically from financial assets [19,20]. Unlike storable commodities, electricity cannot be easily inventoried, which profoundly shapes hedging needs and prevents market makers from managing positions through a physical spot market. The grid itself imposes structural constraints: transmission bottlenecks fragment the market into location-specific prices, while generation faces hard capacity and ramping limits that restrict supply elasticity [21]. The trading structure also departs from financial archetypes; while futures trade continuously, their pricing is inextricably linked to discrete, high-volume daily auctions in the underlying physical market, creating a distinct informational heartbeat.

Second, the behavioural and informational assumptions of canonical models fail. The assumption of anonymity is frequently violated; participants are registered entities with known generation or load obligations, allowing for reputation-based strategies that are absent from anonymous markets. Instead of a single informed trader facing a sea of uninformed noise traders, the market consists of strategic, oligopolistic generators with known market power [22,23] and load-serving entities with highly predictable hedging needs. Crucially, information often arrives through public channels, such as weather forecasts or grid condition reports, rather than the private signals central to Kyle's framework. This public information structure and non-anonymous, strategic environment fundamentally alters the nature of liquidity provision.

The systematic violation of these assumptions renders traditional equilibrium analysis inadequate for power markets. In thin markets, participants recognize their impact on prices and adapt their behaviour accordingly, moving away from the price-taking behaviour assumed in competitive equilibrium models. Analytical models achieve tractability only by eliminating these defining features—strategic interaction, non-storability, and imperfect competition. Relaxing these assumptions simultaneously destroys closed-form solutions.

Agent-based models (ABMs) offer a promising methodological approach to address these analytical challenges [24]. ABMs allow heterogeneous agents to embody distinct operational characteristics, information sets, and realistic behavioural rules—including adaptive heuristics and strategic behaviours that are difficult to model analytically. While this departs from the strict assumption of idealized payoff maximization used in equilibrium analysis, it allows us to model the complex interactions and emergent market dynamics characteristic of real-world power trading [25]. This flexibility is essential for exploring environments where traditional assumptions are violated and equilibrium analysis becomes intractable [26].

Our companion paper developed an agent-based diagnostic that nests the information-based pricing of Kyle [8] and the inventory-based quoting of Ho and Stoll [12] in a single measurable statistic: the *Impact–Inventory Parity* $\Psi := -\beta_{\mathrm{Ho}}/\beta_{\mathrm{Kyle}}$. Under competitive, liquid conditions with linear inventory costs, $\Psi = 1$ (information pressure and inventory pressure balance). That paper showed how two forces systematically depress $\Psi$: (i) limited competition in order-flow capture (effective exclusivity) and (ii) convex inventory costs/constraints. It further identifies an adaptive-behaviour break— a covariation remainder —where impact becomes state-dependent on inventory, breaking static parity even when structure is unchanged. Across broad simulation sweeps calibrated to thin power futures, the modal regime was inventory-dominated rather than information-dominated, with parity in information and inventory quoting a minorityc case ($< 8\%$ of market conditions). Policy experiments indicate that in thin markets, raising effective flow capture (e.g., via exclusivity or quote-obligation design) stabilises $\Psi$ more than simply adding additional market makers. The paper closes by calling for two extensions: empirical estimation of $\Psi$ in real futures data and a relaxation of the price-taking assumption to allow strategic best-response behaviour among dealers.

Thin electricity futures differ from deep, anonymous financial markets along exactly the margins that make strategic interaction first-order. Market makers and large hedgers are few, identifiable, capital-constrained, and repeatedly interact; quotes shape not only transitory execution but also rivals' future inventory and risk capacity. Order arrival is a function of relative quotes and displayed depth, so each dealer internalises a non-trivial demand system rather than facing exogenous noisy flow. Public information (weather, load, grid constraints) compresses asymmetry while inventory and funding constraints dominate; in that environment, the zero-expected-profit discipline of competitive equilibria is weak, and best responses to rivals (and to program rules) become the primitive. A game-theoretic lens captures behaviours that competitive/adaptive benchmarks abstract away: strategic widening or shading of spreads, inventory parking, intertemporal trade-offs under limits, and tacit coordination risks. For market-design questions— exclusivity, minimum depth, quote obligations —these strategic responses are the mechanism through which rules move spreads, depth, volatility, and ultimately $\Psi$.

The following section details our best-response strategic model. Preliminary results are given in Section 3, and the implications of these results discussed in Section 4.

## 2. Methods

The companion paper's competitive/adaptive pricing benchmark abstracts from strategic interaction: dealers take order flow and fundamentals as exogenous, responding myopically to inventory costs and information signals. While this provides a clean baseline for isolating the Kyle and Ho-Stoll mechanisms, it omits a first-order feature of thin power futures markets: agents are few, identifiable, and repeatedly interacting, with quotes that shape rivals' execution probabilities and future states. In such environments, Nash equilibrium—not price-taking—is the appropriate solution concept.

This section develops a game-theoretic benchmark that retains the same market primitives (order-flow process, information structure, inventory costs) but replaces adaptive price-taking with strategic best-response behaviour among finite dealers. We define the market as a dynamic stochastic game, characterize its equilibrium, and describe our computational approach. The resulting simulated paths allow a clean comparison: how do the reduced-form diagnostics $\{\beta_{\text{Kyle}}, \beta_{\text{Ho}}, \Psi\}$ differ when dealers optimize strategically rather than competitively?

### *2.1. Game Definition*

Players, states, and timing.

A finite set of dealers $\mathcal{M} = \{1, \dots, M\}$ compete to intermediate order flow over an infinite discrete-time horizon $t = 0, 1, 2, \dots$. Each period $t$ is characterized by a public state $\theta_t$ summarizing order-arrival intensity, information quality, and the unobservable fundamental value $V_t$, and by private states $(I_{1t}, \dots, I_{Mt})$ giving each dealer's inventory. The aggregate state is

$$x_t \;=\; (I_{1t}, \dots, I_{Mt}, \theta_t) \;\in\; \mathcal{X} \;:=\; \mathbb{R}^M \times \Theta. \tag{1}$$

The state evolves according to a controlled Markov process: inventories adjust via executions, fundamentals follow the same random walk as in the companion paper, and $\theta_t$ tracks persistent market conditions.

At the start of each period, after observing $x_t$, dealers simultaneously post quotes. For parsimony and analytical clarity, we adopt symmetric quoting around a common reference mid-price $m_t$. Note that the mid-price $m_t$ is a public statistic (e.g., the previous period's transaction price or a volume-weighted average). Allowing asymmetric quotes (separate bid and ask decision variables per dealer) is straightforward but doubles the dimensionality of the action space without changing the economic content. Our focus is on strategic interaction in spreads and depth, not quote asymmetry. Each dealer $i$ chooses a quote vector

$$q_i \;=\; (w_i, D_i) \;\in\; \mathcal{Q} \;\subseteq\; \mathbb{R}^2_+, \tag{2}$$

where $w_i \geq 0$ is the half-spread (quoted distance from $m_t$ on both bid and ask sides) and $D_i \geq 0$ is displayed depth (quantity available at the quoted prices). The action space $\mathcal{Q}$ is compact: $w_i \in [0, \bar{w}]$ and $D_i \in [0, \bar{D}]$, reflecting regulatory or technological bounds on quote aggressiveness and size.

Order arrival and execution shares.

Aggregate order flow arrives according to a Poisson process with state-dependent intensity $\Lambda(\theta_t)$, capturing time-varying market activity. Conditional on an arrival, the executing dealer is determined by a *smooth demand system* over relative quotes. We adopt a random-utility microfoundation: each arriving trader $n$ receives a vector of idiosyncratic match values $\{\epsilon_{ni}\}_{i \in \mathcal{M}}$ for trading with each dealer, drawn i.i.d. from a Type-I extreme value distribution. The trader's total utility from executing with dealer $i$ is

$$U_{ni} \;=\; -\gamma(\theta_t)\,w_i \;+\; \xi(\theta_t)\,D_i \;+\; \epsilon_{ni}, \tag{3}$$

where $\gamma(\theta_t) > 0$ captures the trader's spread sensitivity (willingness to pay for price improvement) and $\xi(\theta_t) > 0$ reflects the value of depth (protection against price impact or partial fills). The trader chooses the dealer $i^* = \arg\max_{i\in\mathcal{M}} U_{ni}$. Aggregating over the distribution of $\epsilon_{ni}$, the probability that an arriving order is routed to dealer $i$ is given by the logit share formula [27]:

$$s_i(q_i, q_{-i}; \theta_t) \;=\; \frac{\exp\{-\gamma(\theta_t)w_i + \xi(\theta_t)D_i\}}{\sum_{j\in\mathcal{M}} \exp\{-\gamma(\theta_t)w_j + \xi(\theta_t)D_j\}}, \qquad \sum_j s_j = 1, \quad s_i \in (0,1). \tag{4}$$

This specification satisfies the standard smoothness and monotonicity properties required for existence and uniqueness of pure-strategy equilibria [28]. Crucially, $s_i$ depends on *relative* quotes: narrowing one's spread or increasing depth raises execution share at rivals' expense, the essence of strategic substitutability in this market.

Period payoff.

Dealer $i$'s expected profit in period $t$, given action profile $(q_i, q_{-i})$ and state $x_t$, is

$$\pi_i(q_i, q_{-i}; x_t) \;=\; \underbrace{\Lambda(\theta_t)\, s_i(q_i, q_{-i}; \theta_t)}_{\text{expected executions}} \times \underbrace{\big[w_i\,(1 - \mathcal{A}(\theta_t))\big]}_{\text{net margin per unit}} - \underbrace{c\big(I_{it} + \Delta I_i(q_i, q_{-i}; \theta_t)\big)}_{\text{inventory holding cost}}. \tag{5}$$

We now justify each component.

**Gross margin.** When dealer $i$ executes a unit at half-spread $w_i$, the gross revenue is $w_i$ (capturing the bid-ask spread earned on a round-trip, or the deviation from fair value on a one-sided trade).

**Adverse selection.** Not all spread is profit: informed traders systematically trade against the dealer when they possess superior information. Following Glosten and Milgrom [14], we model this via a *reduced-form adverse-selection parameter* $\mathcal{A}(\theta_t) \in [0,1)$, which represents the fraction of gross margin lost to informed order flow. Specifically, $\mathcal{A}(\theta_t)$ is the probability-weighted information content of order flow:

$$\mathcal{A}(\theta_t) \;=\; \rho(\theta_t) \cdot \frac{\operatorname{Var}(\tilde{v} \mid \text{informed arrival})}{\operatorname{Var}(\tilde{v})}, \tag{6}$$

where $\rho(\theta_t)$ is the proportion of informed traders in the arrival population and the ratio captures the informativeness of their signals relative to prior uncertainty. In equilibrium with competitive information revelation, $\mathcal{A}(\theta_t)$ is pinned down by the Kyle lambda (price impact per unit flow) and volatility [8]. We calibrate $\mathcal{A}(\theta_t)$ from the companion paper's estimated $\beta_{\text{Kyle}}$ to ensure consistency across benchmarks. The net margin per unit executed is thus $w_i(1 - \mathcal{A}(\theta_t))$, capturing the trade-off between spread revenue and information leakage.

**Inventory cost.** Dealers face a convex cost $c(I)$ of holding inventory, reflecting capital constraints, risk aversion, and regulatory costs. We adopt the quadratic baseline

$$c(I) \;=\; \frac{\kappa}{2}\, I^2, \tag{7}$$

with $\kappa > 0$ governing the severity of inventory aversion. This specification is standard in market microstructure [12,29] and ensures concavity of the profit function. Extensions with hard position limits or piecewise-linear costs are straightforward.

**Inventory transition.** Expected executed quantity for dealer $i$ is

$$Q_i(q_i, q_{-i}; \theta_t) \;=\; \Lambda(\theta_t)\, s_i(q_i, q_{-i}; \theta_t)\, \bar{q}(\theta_t), \tag{8}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 November 2025 doi:10.20944/preprints202511.0481.v1



where $\bar{q}(\theta_t)$ is the mean order size (drawn from the same distribution as in the companion paper). The inventory change is

$$\Delta I_i(q_i, q_{-i}; \theta_t) \;=\; \sigma_i \, Q_i(q_i, q_{-i}; \theta_t), \tag{9}$$

where $\sigma_i \in \{-1, +1\}$ is the trade sign (buy or sell from the dealer's perspective), drawn from the signed-flow process used in the companion paper. This ensures that the information and impact environment is identical across benchmarks; only the dealers' decision rule changes.

### 2.2. Equilibrium Concept and Approximation

The full dynamic game.

The natural equilibrium concept for this setting is a *Markov-perfect equilibrium* (MPE) in stationary strategies $\{\sigma_i : \mathcal{X} \to \mathcal{Q}\}_{i \in \mathcal{M}}$. Each dealer $i$ solves

$$V_i(x) \;=\; \max_{q_i \in \mathcal{Q}} \mathbb{E}\Big[\pi_i(q_i, \sigma_{-i}(x); x) \;+\; \beta \, V_i(x') \;\Big|\; x, q_i, \sigma_{-i}(x)\Big], \tag{10}$$

where $\beta \in (0,1)$ is the discount factor, $x' = (I'_1, \ldots, I'_M, \theta')$ is the next-period state, and the expectation integrates over execution realizations and exogenous shocks to $\theta'$. An MPE is a fixed point: $\sigma_i \in \arg\max$ for each $i$, simultaneously.

Computing the full MPE is intractable for three reasons: (i) the state space $\mathcal{X} = \mathbb{R}^M \times \Theta$ is high-dimensional and continuous; (ii) the continuation value $V_i(x')$ depends on rivals' future strategies $\sigma_{-i}$, requiring a nested fixed point; and (iii) the state is stochastic with a non-trivial transition kernel. Standard approaches (value function iteration, policy iteration) face the curse of dimensionality and do not scale beyond $M = 2$ or 3 dealers.

Stage-game approximation with shadow inventory costs.

We adopt a tractable approximation that preserves the strategic structure while eliminating the dynamic programming burden. The key insight is that, in stationary equilibrium with fast mean-reversion of inventories, the continuation value is *approximately separable* in own inventory:

$$V_i(x) \;\approx\; V_i^{\text{base}}(\theta) \;-\; \lambda \, I_i^2, \tag{11}$$

where $V_i^{\text{base}}(\theta)$ captures the value of future quoting opportunities (common across inventory levels) and $-\lambda I_i^2$ is a quadratic penalty for deviating from zero inventory. This approximation is valid under two conditions:

(i) **Fast mean-reversion:** Inventories fluctuate around zero with high frequency, so the marginal value of inventory $\partial V_i / \partial I_i$ is approximately linear in $I_i$ (hence $V_i$ quadratic).

(ii) **Competitive continuation:** Rivals' strategies $\sigma_{-i}$ are sufficiently aggressive (spreads remain tight) that future execution opportunities are not drastically altered by current inventory positions, so cross-effects $\partial V_i / \partial I_j$ are second-order.

Both conditions are satisfied in our calibrated environment: mean reversion times are $O(10)$ periods, and spreads remain within a few ticks of competitive levels even in the strategic equilibrium (verified ex-post; see Results).

Under (11), the marginal continuation value of a unit inventory change is

$$\beta \left[V_i(I_{it} + \Delta I_i) - V_i(I_{it})\right] \;\approx\; -\beta\lambda \left[(I_{it} + \Delta I_i)^2 - I_{it}^2\right] \;=\; -\beta\lambda \left[2 I_{it} \Delta I_i + (\Delta I_i)^2\right]. \tag{12}$$

Substituting into (10) and dropping the state-independent constant $V_i^{\text{base}}(\theta)$, we obtain an *augmented stage game* with effective period payoff

$$\tilde{\pi}_i(q_i, q_{-i}; x_t) \;=\; \pi_i(q_i, q_{-i}; x_t) \;-\; \beta\lambda \left[2 I_{it} \Delta I_i(q_i, q_{-i}; \theta_t) \;+\; \Delta I_i(q_i, q_{-i}; \theta_t)^2\right]. \tag{13}$$

The first term, $-\beta\lambda \cdot 2I_{it}\Delta I_i$, is a *linear shadow cost* penalizing inventory accumulation proportional to current position; the second term, $-\beta\lambda \cdot (\Delta I_i)^2$, is a *variance penalty* discouraging large executions. Together, they internalize the dealer's dynamic inventory-management problem in a static objective.

Calibration of $\lambda$.

The shadow-cost parameter $\lambda$ must be chosen to match the continuation value curvature in the true MPE. We calibrate $\lambda$ by equating the marginal cost of inventory in (13) to that implied by the inventory cost function $c(I)$:

$$\beta\lambda \cdot 2I_{it} \;=\; c'(I_{it}) \;\implies\; \lambda \;=\; \frac{c'(I_{it})}{2\beta I_{it}}. \tag{14}$$

For the quadratic cost $c(I) = (\kappa/2)I^2$ in (7), this yields $\lambda = \kappa/(2\beta)$. In practice, we set $\lambda = \kappa$ (corresponding to $\beta \approx 0.5$, a typical value for intraday horizons), which we verify ex-post by checking that mean squared inventories in simulation are consistent with the assumed curvature.

Nash equilibrium of the stage game.

With the augmented payoff (13), we seek a Nash equilibrium of the period-$t$ stage game: a profile $q^*(x_t) = (q_1^*(x_t), \ldots, q_M^*(x_t))$ such that

$$q_i^*(x_t) \;\in\; \arg\max_{q_i \in \mathcal{Q}} \; \tilde{\pi}_i(q_i, q_{-i}^*(x_t); x_t), \qquad \forall\, i \in \mathcal{M}. \tag{15}$$

**Existence and uniqueness.** We verify that the standard sufficient conditions for existence and uniqueness of pure-strategy Nash equilibria [28] are satisfied:

(a) **Compactness:** $\mathcal{Q} = [0, \bar{w}] \times [0, \bar{D}]$ is compact and convex.

(b) **Continuity:** $\tilde{\pi}_i(q_i, q_{-i}; x_t)$ is continuous in $(q_i, q_{-i})$ (logit shares are $C^\infty$, inventory costs are $C^2$).

(c) **Concavity:** For fixed $q_{-i}$, $\tilde{\pi}_i$ is strictly concave in $q_i$. To see this, note that the logit share $s_i(q_i, q_{-i})$ is log-concave in $q_i$ (standard result for exponential families), revenue $w_i s_i$ is log-concave when spreads are not too wide [30], and the inventory cost $c(\cdot)$ is strictly convex. The sum inherits concavity.

(d) **Diagonal strict concavity:** The Jacobian matrix of the aggregate payoff $\sum_i \tilde{\pi}_i$ is negative definite (Rosen's sufficient condition for uniqueness), which we verify numerically for our parameterization.

By Rosen's theorem, a pure-strategy Nash equilibrium $q^*(x_t)$ exists and is generically unique for each state $x_t$.

### 2.3. Computation of the Equilibrium

We compute the equilibrium $q^*(x_t)$ using two complementary methods, both of which are standard in applied game theory and industrial organization:

Method 1: Damped best-response iteration.

Starting from an initial guess $q^{(0)}$, we iterate

$$q_i^{(k+1)} \;=\; (1-\eta)\, q_i^{(k)} \;+\; \eta\, \mathrm{BR}_i(q_{-i}^{(k)}; x_t), \tag{16}$$

where $\mathrm{BR}_i(q_{-i}; x_t) \in \arg\max_{q_i} \tilde{\pi}_i(q_i, q_{-i}; x_t)$ is $i$'s best response to $q_{-i}$, and $\eta \in (0,1]$ is a damping parameter. For concave games, this iteration converges to the unique equilibrium when $\eta$ is sufficiently small [31]. We set $\eta = 0.3$ and terminate when $\|q^{(k+1)} - q^{(k)}\| < 10^{-4}$, which typically occurs in 20–50 iterations.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 November 2025 doi:10.20944/preprints202511.0481.v1



Each best response $\mathrm{BR}_i$ is computed by grid search over $\mathcal{Q}$ (discretized into $50 \times 50$ grid points) or, for finer accuracy, by gradient ascent on $\tilde{\pi}_i$ using automatic differentiation. The logit share function (4) and inventory cost (7) are smooth, so standard optimization routines (L-BFGS-B) converge rapidly.

Method 2: Logit quantal-response equilibrium (QRE).

As a robustness check and to avoid potential cycling in best-response iteration, we also compute equilibria via the logit QRE [32]. A QRE with rationality parameter $\mu > 0$ is a fixed point $\tilde{q}(\mu)$ where each dealer's action is drawn from a smoothed best-response distribution:

$$\mathbb{P}\{q_i \mid q_{-i}, x_t\} \propto \exp\{\mu \, \tilde{\pi}_i(q_i, q_{-i}; x_t)\}. \tag{17}$$

As $\mu \to \infty$, the QRE converges to the Nash equilibrium (agents become perfectly rational). For finite $\mu$, the QRE provides a regularized solution that avoids knife-edge discontinuities and facilitates smooth convergence.

We compute the QRE fixed point by iterating the QRE mapping (17) until convergence, then increase $\mu$ gradually (continuation method) and re-solve, tracing the solution path toward the Nash limit. In practice, $\mu = 100$ is sufficient to approximate Nash equilibria to within simulation noise. This method produces nearly identical results to damped best-response (correlation $> 0.99$) but is more robust to initialization and avoids local cycling.

### *2.4. Computational Protocol and Alignment with the Companion Paper*

This section explains how we implemented the strategic market game defined above using the same computational infrastructure, market environments, and econometric measurement protocols established in [33]. The sole modification was the replacement of adaptive behavioural heuristics with Nash equilibrium computation at each time step. All market regime parameters were kept the same, and the full range explored in the companion paper. This design isolated the impact of strategic optimization on market outcomes while holding all other structural and environmental factors constant.

At each period $t$, rather than having the market maker follow the companion paper's adaptive pricing heuristics, we computed the Nash equilibrium $q^*(x_t)$ using damped best-response iteration as specified in the previous Section. The market maker observes the current state $x_t = (I_{1t}, \ldots, I_{Mt}, \theta_t)$, solves the optimization problem defined in Equation 15 to determine optimal quotes $(w_i^*, D_i^*)$, and posts these equilibrium quotes. Order flow then realizes according to the logit share function in Equation 4, inventories update, and the process continues. This substitution represents the minimal change necessary to transition from adaptive heuristics to game-theoretic best response.

Third, we applied the identical econometric measurement apparatus developed in the companion paper. After each simulation run completed, we estimated the Kyle coefficient through the regression $\Delta m_t = \alpha_K + \beta_{\text{Kyle}} Q_t + \varepsilon_t$, estimate the Ho-Stoll coefficient through the cointegrating regression $m_t - V = \alpha_H + \beta_{\text{Ho}} I_{t-1} + u_t$ with Augmented Dickey-Fuller testing, calculate the Parity Index $\Psi = -\beta_{\text{Ho}}/\beta_{\text{Kyle}}$, and aggregate results across $N_{\text{sim}} = 1000$ independent replications using Rubin's Rules. This ensures that any differences in measured outcomes reflect genuine behavioural distinctions rather than measurement inconsistencies. Fourth, we conducted systematic comparative analysis. For each configuration tested in the companion paper, we computed the behavioural deltas: $\Delta\beta_{\text{Kyle}} = \beta_{\text{Kyle}}^{\text{strategic}} - \beta_{\text{Kyle}}^{\text{adaptive}}$, $\Delta\beta_{\text{Ho}} = \beta_{\text{Ho}}^{\text{strategic}} - \beta_{\text{Ho}}^{\text{adaptive}}$, and $\Delta\Psi = \Psi^{\text{strategic}} - \Psi^{\text{adaptive}}$. These differences quantify how strategic optimization modifies the market outcomes observed under adaptive heuristics. Positive values of $\Delta\Psi$ indicate that strategic behaviour amplifies inventory dominance relative to the adaptive benchmark, while negative values indicate strategic behaviour attenuates this tendency.

## 3. Results

The overall conclusion from these rerun simulations was that strategic optimization amplified the market microstructure patterns documented in the companion paper. Where the adaptive behavioural

 doi:10.20944/preprints202511.0481.v1



model revealed systematic relationships between market fragmentation, inventory costs, and the Parity Index $\Psi$, the strategic equilibrium model intensified these effects through endogenous dealer responses. Dealers actively adjusted spreads $w_i$ and depths $D_i$ in response to inventory pressure and competitive conditions, creating feedback mechanisms that magnified the original adaptive patterns. Across all eight key findings from the companion paper, strategic behaviour either strengthened the observed relationships or revealed starker regime boundaries, confirming that the adaptive heuristics captured genuine economic forces rather than behavioural artifacts.

### *3.1. Competition Effect: Enhanced Fragmentation Impact*

The companion paper established that the Parity Index tracks the inverse of dealer concentration: $\Psi \approx 1/\phi$ when inventory costs are linear $\kappa = 0$. Strategic behaviour amplified this relationship substantially. In the strategic equilibrium, each dealer's capture share $s_i(q;\theta)$ becomes endogenously state-dependent through sensitivity to spreads $\gamma(\theta)$ and depth valuations $\xi(\theta)$. When these sensitivities are high, equilibrium shares concentrate near $1/M$ in symmetric configurations, strengthening the decoupling between individual dealer inventories and aggregate information flow.

However, market-level prices reflect the best quotes across all dealers rather than any single dealer's position. Because the strategic model prohibits inventory-based quote tilting (bid-ask asymmetry), the link between dealer inventory $I_i$ and the market mid-price deviation $m_t - V$ weakens substantially. This drives down the measured Ho-Stoll coefficient $|\beta_{\text{Ho}}|$ more sharply than in the adaptive benchmark, while the Kyle coefficient $\beta_{\text{Kyle}}$ at the market level remains relatively stable. Consequently, $\Psi$ values cluster closer to unity but exhibit steeper declines with fragmentation: the $1/\phi$ slope is more pronounced, indicating that strategic competition amplifies inventory dominance when markets fragment.

Decomposing by dealer-level versus market-aggregated measures confirms this mechanism. Dealer-level Parity Indices $\Psi_i$ (computed from individual inventory versus market mid-price) fall more sharply with market concentration $H = \sum s_i^2$ than in the adaptive model. The strategic equilibrium thus reveals that the competition effect operates even more powerfully when dealers optimize their quote exposure, making fragmentation's structural push toward inventory dominance more severe than adaptive heuristics suggested.

### *3.2. Kurtosis Effect: Strategic Dampening Through Depth Control*

The companion paper documented that non-linear inventory costs $\kappa > 0$ suppress $\Psi$ below unity, with suppression particularly strong in liquid markets, creating visible stratification across market thickness regimes. Strategic behaviour attenuated this pattern. Dealers in the strategic equilibrium can throttle execution risk by adjusting depth $D_i$ or widening spread $w_i$ whenever inventories grow large, providing an active release valve that the adaptive model lacked. This endogenous depth modulation directly reduces inventory variance $\sigma^2(I_i)$ and higher moments.

Because the kurtosis-driven suppression of $\Psi$ scales with inventory variance and tail behaviour, the strategic model exhibits weaker suppression than the adaptive benchmark, particularly in liquid multi-dealer markets where each dealer's flow share is small. The $\Psi$-$\kappa$ curves are flatter, and the cross-market stratification visible in the companion paper's Figure 1 diminishes substantially. Decomposing inventory variance with and without strategic depth adjustment reveals that the depth margin absorbs much of the inventory pressure that would otherwise manifest as price impact. Consequently, the region of parameter space where $\Psi < 1$ due to kurtosis effects shrinks relative to the adaptive model, indicating that strategic depth management partially insulates the Parity Index from non-linear inventory costs.

### *3.3. Joint Effects: Shifted Parity Boundary*

The companion paper identified a curved $\Psi = 1$ boundary in $(\phi, \kappa)$ space, where achieving parity in fragmented markets required substantially higher inventory cost non-linearity. Strategic behaviour shifted this boundary leftward and downward. Because fragmentation lowers each dealer's inventory

 doi:10.20944/preprints202511.0481.v1



variance $\sigma^2(I_i)$ in the strategic equilibrium—dealers with small market shares face smaller individual inventory fluctuations—the threshold $\kappa$ needed to offset fragmentation's effects decreases. When spread sensitivity $\gamma(\theta)$ and depth sensitivity $\xi(\theta)$ are high, the strategic shares become more elastic, amplifying this variance-reduction mechanism.

The curvature of the boundary persists, reflecting the fundamental tension between decoupling (from fragmentation) and variance control (from non-linear costs), but the region where $\Psi < 1$ substantially contracts. Overlaying the adaptive model's $\Psi = 1$ contour on the strategic equilibrium heatmap reveals a systematic leftward-downward shift, with the greatest displacement occurring in regions where dealer shares are most responsive to quote competitiveness. This confirms that strategic optimization amplifies the joint interaction between market structure and inventory management, making the transition to inventory dominance more sensitive to competitive conditions than the adaptive model predicted.

### *3.4. Adaptive Parity Breakdown: Mitigated Collapse*

The companion paper's most dramatic finding was that adaptive spread adjustment $\lambda_{\text{adaptive}} > 0$ drove severe Parity Index collapses, with liquidity amplifying the breakdown: $\Psi \ll 1$ in thick markets, with values approaching 0.05 in the most extreme cases. Strategic behaviour substantially mitigated this collapse. In the strategic equilibrium, state-dependent inventory management migrates from the spread adjustment parameter $\lambda_t$ to the dealers' choice variables $(w_i, D_i)$, mediated through shadow inventory costs in the stage payoff function. This architectural difference eliminates the covariation remainder mechanism that caused the adaptive model's breakdowns.

When dealers face inventory pressure, they widen spreads $w_i$ and reduce depths $D_i$ rather than amplifying price impact. Rival dealers then capture the marginal flow, spreading execution across the market. The Kyle coefficient $\beta_{\text{Kyle}}$ at the market level remains stable or declines modestly, while the Ho-Stoll coefficient $|\beta_{\text{Ho}}|$ falls because individual dealer inventories exert less influence on market prices. Consequently, the dramatic collapses observed in the adaptive model—where $\Psi$ approached zero in liquid monopolistic configurations—do not occur in the strategic equilibrium. The lowest observed values of $\Psi$ remain above 0.3, and the amplification of collapse with liquidity virtually disappears.

Decomposing the change $\Delta\Psi$ into contributions from $\Delta\beta_{\text{Kyle}}$ versus $\Delta\beta_{\text{Ho}}$ confirms this mechanism. The adaptive model's collapses were driven by large increases in $\beta_{\text{Kyle}}$ through the covariation remainder; the strategic model exhibits small changes in $\beta_{\text{Kyle}}$ but larger declines in $|\beta_{\text{Ho}}|$. Analysis of dealer turnover in "best quote" status during extreme inventory episodes reveals extensive flow-switching in the strategic equilibrium, contrasting sharply with the impact amplification that characterized the adaptive breakdown.

### *3.5. Stability and Econometric Performance*

The companion paper documented higher econometric failure rates in thin, fragmented, or low-$\kappa$ configurations, though successful runs exhibited very high $R^2$ values for both Kyle and Ho-Stoll regressions. Strategic behaviour improved stability substantially. The concave game structure and Quantal Response Equilibrium continuation ensure a unique stationary equilibrium for each state, producing better-behaved share and price processes than the adaptive model's heuristic dynamics. Econometric failure rates declined across all parameter regions, with the sharpest improvements occurring in the previously fragile corners (thin markets with many dealers and low inventory costs).

For successful runs, $R^2$ values remained similarly high or improved slightly, though the Ho-Stoll regression occasionally exhibited weaker statistical power due to the mechanical reduction in $|\beta_{\text{Ho}}|$ from symmetric quote constraints. The strategic model's superior stability reflects the disciplining force of equilibrium selection: dealers cannot pursue locally destabilizing strategies because rivals' best responses preclude such paths. This confirms that the adaptive model's stability challenges stemmed from the behavioural specification rather than fundamental economic constraints, and that strategic optimization naturally stabilizes market dynamics even in challenging configurations.

### 3.6. Global Sensitivity: Amplified Interaction Effects

The companion paper identified three equally important pillars of market stability—dealer count $M$, noise trader intensity $\eta$, and information precision $\sigma_\varepsilon$—and showed that among stable markets, information precision $\sigma_\varepsilon$ dominated Parity Index variation. Strategic behaviour preserved this three-pillar structure while amplifying interaction effects. The fundamental drivers $M$, $\eta$, and $\sigma_\varepsilon$ remain essential because they determine aggregate order flow variance $\Lambda(\theta)$, which the logit-share mechanism still requires for stable execution.

However, conditional on stability, the strategic equilibrium reveals stronger interactions between information precision $\sigma_\varepsilon$ and the new behavioural parameters: spread sensitivity $\gamma(\theta)$ and depth sensitivity $\xi(\theta)$. Because depth competition actively blunts per-dealer inventory excursions, these sensitivities now mediate how information uncertainty translates into inventory pressure. Sobol decomposition under strategic pricing shows similar first-order contributions from $\sigma_\varepsilon$ but substantially larger total-effect indices for $\gamma$ and $\xi$ through interaction terms. This indicates that strategic optimization amplifies the channels through which market microstructure primitives jointly determine the Parity Index, making the regime outcomes more sensitive to the institutional details of quote competition than the adaptive model suggested.

### 3.7. Regime Prevalence: Shift Toward Balanced Markets

The companion paper documented that in the plausible parameter space, approximately 73% of markets exhibited inventory dominance $\Psi < 0.9$, 8% were balanced $0.9 \leq \Psi \leq 1.1$, and 19% showed adverse selection dominance $\Psi > 1.1$. Strategic behaviour shifted this distribution substantially toward balanced markets. The modal value of $\Psi$ moved closer to unity, driven by the systematic dampening of $|\beta_{\text{Ho}}|$ under symmetric quote constraints combined with stable or declining $\beta_{\text{Kyle}}$ from depth competition.

The inventory-dominated share declined from 73% to approximately 49%, while the balanced regime expanded from 8% to 19%. The adverse-selection-dominated share increased modestly from 19% to 32%, with most of the shift occurring in liquid markets where strategic depth management prevents inventory accumulation from dominating price formation. This redistribution confirms that strategic optimization attenuates the structural bias toward inventory dominance documented in the adaptive model, though inventory effects still dominate the majority of configurations. The binding constraint on regime membership remains the Ho-Stoll coefficient, but its reduced magnitude in the strategic equilibrium fundamentally reshapes the landscape of plausible market regimes.

## 4. Discussion and Conclusions

This paper reexamined the eight principal findings of the companion paper through the lens of strategic equilibrium rather than adaptive heuristics. Despite fundamentally different behavioural foundations—– adaptive dealers following gradient -based learning rules versus strategic dealers solving dynamic games —–the market microstructure patterns proved robust. Strategic optimization amplified most effects documented in the adaptive model: fragmentation's push toward inventory dominance became steeper, kurtosis-driven suppression was partially mitigated by endogenous depth control, and the $\Psi = 1$ boundary shifted in predictable ways. The most dramatic divergence occurred in the adaptive parity breakdown, where strategic behaviour eliminated the severe collapses observed under adaptive spread adjustment, yet even here the directional effect—– that state-dependent pricing alters inventory dominance —–remained consistent across specifications.

This convergence despite architectural differences reveals something fundamental about dealer market microstructure. The Parity Index $\Psi = -\beta_{\text{Ho}}/\beta_{\text{Kyle}}$ captures a structural tension between two universal forces: inventory pressure (the dealer's need to offload positions) and information incorporation (the market's response to order flow signals). These forces operate regardless of whether dealers optimize myopically through heuristics or strategically through equilibrium selection. The adaptive model's gradient adjustments and the strategic model's first-order conditions both respond to the same economic gradients: inventory costs, adverse selection risk, and competitive pressure. What changes across specifications is not the existence or direction of these forces, but rather their relative magnitudes and the channels through which they operate.

The game-theoretic sophistication adds nuance rather than contradiction. Strategic dealers internalize how their quotes affect both current execution and future inventory distributions, leading them to adjust depths and spreads in ways that blunt extreme outcomes. Adaptive dealers respond only to recent experience, occasionally producing runaway dynamics when feedback mechanisms align perversely. Yet in both models, fragmentation decouples individual inventories from market prices, non-linear inventory costs compress the Parity Index, and the interaction between competition and risk management determines regime boundaries. The adaptive heuristics thus captured genuine economic structure rather than behavioural artifacts; strategic optimization refines the quantitative predictions while confirming the qualitative architecture.

This robustness has important methodological implications. The computational cost of solving dynamic stochastic games with continuous state spaces remains prohibitive for many applications, particularly when exploring high-dimensional parameter spaces or conducting policy counterfactuals. Our findings suggest that carefully calibrated adaptive models can provide reliable qualitative guidance and reasonable quantitative approximations, at least for first-order effects. The strategic model remains essential for welfare analysis and for understanding responses to regulation—– contexts where anticipatory behaviour and equilibrium selection matter critically –—but the adaptive framework offers a tractable complement for exploring structural relationships and generating empirical predictions. This division of labour between specifications may prove useful as market microstructure research confronts increasingly complex institutional environments.

The policy implications evolved from the companion paper's conclusion but retained their contingent character. The companion paper's policy prescription favoured consolidation over competition in very thin markets, arguing that fragmentation structurally pushes markets toward inventory dominance. Strategic behaviour complicates this conclusion substantially. In the strategic equilibrium, moderate competition can improve stability and reduce inventory pressure by splitting flow and allowing depth margins to absorb trades. This softens the case for universal consolidation, making the optimal policy contingent on market characteristics and behavioural parameters.

In very thin markets (small $M$, low $\eta$), consolidation remains beneficial because even with logit shares, each dealer faces lumpy flow that overwhelms endogenous adjustments. However, in moderately thin or liquid markets with strong depth sensitivity ($\xi$ high), competition can lower individual dealer inventory variance $\sigma^2(I_i)$ without destroying profit margins, reducing the need for consolidation. When adverse selection is severe (high $\mathcal{A}$), competition risks a winner's-curse race to the bottom; strategic dealers endogenously widen spreads in response, shifting the policy focus from dealer count to information symmetry. Policy counterfactuals varying $M$ alongside taste parameters $\gamma$ and $\xi$ reveal non-monotonic welfare effects of competition once strategic depth management is operative, indicating that regulatory interventions must account for dealers' equilibrium responses rather than treating market structure as exogenous to behaviour.

Together, the companion paper and this strategic extension demonstrate that the Parity Index framework provides a coherent lens for understanding how market structure, inventory management, and information asymmetry jointly determine dealer market outcomes. The Kyle and Ho-Stoll coefficients—though estimated from reduced-form regressions—encode structural relationships that remain stable across radically different behavioural specifications. The index's empirical tractability, combined with its theoretical grounding in both adaptive and strategic models, makes it a promising tool for empirical market microstructure research. Future work can extend this framework to asymmetric dealer configurations, time-varying information environments, and institutional features such as quote transparency and trade reporting delays, zonal pricing models, and coupled markets, confident that the core structural relationships documented here will continue to organize the empirical patterns. The fundamental tension between inventory pressure and information incorporation transcends model specification, providing a robust foundation for understanding dealer intermediation in both theoretical and empirical contexts.

## References

1. Hagemann, S.; Weber, C. Analyzing the liquidity of the German electricity futures market. *Energy Economics* **2010**, *32*, 1113–1125.
2. Frestad, D.; Anderson, S.C.; Bunn, D. The liquidity of the Nordic power market. *The Energy Journal* **2010**, *31*, 1–22.
3. Wüstenhagen, R.; Menichetti, E. Strategic choices for renewable energy investment: Conceptual framework and opportunities for further research. *Energy Policy* **2012**, *40*, 1–10.
4. Weron, R. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* **2014**, *30*, 1030–1081.
5. Hirth, L. The market value of variable renewables: The effect of solar and wind power variability on their relative price. *Energy Economics* **2013**, *38*, 218–236.
6. Gatzert, N.; Vogl, N. Evaluating investments in renewable energy under policy risks. *Energy Policy* **2016**, *95*, 238–252.
7. Peña, J.I.; Rodríguez, R. Market Makers and Liquidity Premium in Electricity Futures Markets. *The Energy Journal* **2022**, *43*.
8. Kyle, A.S. Continuous auctions and insider trading. *Econometrica* **1985**, *53*, 1315–1335.
9. Hasbrouck, J. Measuring the information content of stock trades. *The Journal of Finance* **1991**, *46*, 179–207.
10. Biais, B.; Glosten, L.; Spatt, C. Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets* **2005**, *8*, 217–264.
11. Madhavan, A. Market microstructure: A survey. *Journal of Financial Markets* **2000**, *3*, 205–258.
12. Ho, T.; Stoll, H.R. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* **1981**, *9*, 47–73.
13. Amihud, Y.; Mendelson, H. Dealership market: Market-making with inventory. *Journal of Financial Economics* **1980**, *8*, 31–53.
14. Glosten, L.R.; Milgrom, P.R. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* **1985**, *14*, 71–100.
15. Easley, D.; O'Hara, M. Price, trade size, and information in securities markets. *Journal of Financial Economics* **1987**, *19*, 69–90.
16. Peña, J.I.; Rodríguez, R. Market makers and liquidity premium in electricity futures markets. *The Energy Journal* **2022**, *43*, 143–166.
17. Antón, A.; Bushnell, J.B. Market power in electricity markets: Beyond concentration measures. *The Energy Journal* **2002**, *23*, 65–88.
18. Kumar, P.; Seppi, D.J. Manipulation and the allocational role of prices. *The Review of Financial Studies* **1992**, *5*, 651–678.
19. Stoft, S. *Power system economics: Designing markets for electricity*; IEEE Press/Wiley-Interscience, 2002.
20. Kirschen, D.S.; Strbac, G. *Fundamentals of power system economics*, 2nd ed.; John Wiley & Sons, 2018.
21. Joskow, P.L. California's electricity crisis. *Oxford Review of Economic Policy* **2001**, *17*, 365–388.
22. Borenstein, S.; Bushnell, J.B.; Wolak, F.A. Measuring market inefficiencies in California's restructured wholesale electricity market. *American Economic Review* **2002**, *92*, 1376–1405.
23. Wolfram, C.D. Measuring market power in the British electricity spot market. *American Economic Review* **1999**, *89*, 805–826.
24. Weidlich, A.; Veit, D. A critical survey of agent-based wholesale electricity market models. *Energy Economics* **2008**, *30*, 1728–1759.
25. Bunn, D.W.; Oliveira, F.S. Agent-based simulation—an application to the new electricity trading arrangements of England and Wales. *IEEE Transactions on Evolutionary Computation* **2001**, *5*, 493–503.
26. Ringler, P.; Keles, D.; Fichtner, W. Agent-based modelling and simulation of smart electricity grids and markets–A literature review. *Renewable and Sustainable Energy Reviews* **2016**, *57*, 205–215.
27. McFadden, D. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*; Zarembka, P., Ed.; Academic Press, 1974; pp. 105–142.
28. Rosen, J.B. Existence and uniqueness of equilibrium points for concave *n*-person games. *Econometrica* **1965**, *33*, 520–534.

doi:10.20944/preprints202511.0481.v1




29. Amihud, Y.; Mendelson, H. Dealership market: Market-making with inventory. *Journal of Financial Economics* **1980**, *8*, 31–53.
30. Caplin, A.; Nalebuff, B. Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica* **1991**, *59*, 25–59.
31. Fudenberg, D.; Levine, D.K. *The Theory of Learning in Games*; MIT Press, 1998.
32. McKelvey, R.D.; Palfrey, T.R. Quantal response equilibria for normal form games. *Games and Economic Behavior* **1995**, *10*, 6–38.
33. Williams, P.R. Market Makers in Thin Power Futures Markets: Testing the Kyle Model's Robustness **2025**. submitted to Energies.