

Evolution of Market Microstructure in the Age of AI

From Double Auctions to Agentic Markets

Irene Aldridge

AbleBlox

irene.aldridge@gmail.com

September 2026

Abstract

Market microstructure studies how trading rules turn orders into prices and allocations. Those rules have been rebuilt repeatedly: for floor traders, electronic limit order books and high-frequency trading, batch auctions and dark pools, blockchains run by automated market makers and block builders, and now AI agents that discover, pay for and compete over resources. This survey traces that evolution through one question: *can a market allocate scarce goods efficiently and fairly without participants revealing everything they know and want?* Each technological shift moved the binding constraint of market design, from trader rationality to speed, to control over transaction ordering, to the dimensionality of what participants can report. Impossibility results persist; only their cost moves. We review evidence on high-frequency trading, exchange and bilateral venues, automated market makers, extractable value and algorithmic collusion, and identify allocation as the missing layer of the agent-commerce protocol stack.

Keywords: market microstructure; mechanism design; high-frequency trading; batch auctions; dark pools; automated market makers; maximal extractable value; AI agents; matching; Nash social welfare.

JEL: D44, D47, D82, G14, G18, G23.

Contents

1	Introduction	3
2	Foundations: What a Trading Mechanism Can and Cannot Do	4
2.1	The double auction and the clearing price	4
2.2	The two-sided impossibility	5
2.3	Why real double auctions work anyway	5
2.4	Information and the market maker	5
3	The Speed Era: Continuous Trading and the Arms Race	6
3.1	Algorithms become the market makers	6
3.2	The latency arms race	6
3.3	Frequent batch auctions	6
3.4	Batch mechanisms in practice	7
3.5	The cost of batching: delayed price discovery	7

4	The Information Era: Lit, Dark and Batch Venues	8
4.1	Three venues as one queueing system	8
4.2	Model and welfare ranking	9
4.3	Regulatory relevance	10
5	Bilateral Markets: Dealers, Search and Requests for Quote	10
5.1	Why bilateral trading persists	10
5.2	Search, bargaining and dealer networks	10
5.3	Electronic requests for quote	11
5.4	Bilateral trading on-chain and between agents	11
6	Designing for What Participants Report: Revelation and Automated Market Makers	12
6.1	The revelation principle and its limits	12
6.2	Prediction markets and the first automated market maker	13
7	On-Chain Microstructure: Blockchains, AMMs and Extractable Value	14
7.1	What institutional adoption requires	14
7.2	Constant-function market makers	15
7.3	Batching returns, with a builder in charge	15
7.4	The research response: restrain the builder, or rescue the AMM	16
7.5	Hiding the book on-chain	17
8	AI Agents as Market Participants	18
8.1	From algorithms to learning agents	18
8.2	The agent commerce stack	18
8.3	Evidence from agentic commerce	19
8.4	What agents can cheaply report	20
9	Allocation Without Rankings: Matching Mechanisms for Agents	20
9.1	Allocation without money	20
9.2	A feature-based model	21
9.3	Spectral projection	21
9.4	Guarantees and diagnostics	22
9.5	Evidence from a simulated agentic-shopping market	24
10	Synthesis: What Each Design Buys and What It Costs	24
11	Conclusion	25




Figure 1: The evolution of market microstructure, from foundational auction theory through automated, high-speed and on-chain markets to agentic marketplaces. Filled markers denote designs whose ideas recur on-chain or in agentic markets later in the survey. The agent protocols of 2024–2025 standardize tools, payment and identity. The allocation layer of agentic marketplaces is the subject of Sections 8–9.

1 Introduction

Every market must answer the same three questions: who trades, at what price, and in what order. Market microstructure theory, from the dealer models of [Glosten and Milgrom \(1985\)](#) and [Kyle \(1985\)](#) to the textbook synthesis of [O’Hara \(1995\)](#), studies how the answers depend on the rules of the trading venue and on what participants know. For most of that history the participants were people, and the rules were written for them. Three successive waves of automation have since changed who, or what, sits on each side of the book.

The first wave was electronic: continuous limit order books replaced the floor, and algorithms replaced the order clerk. The second was high-frequency: when execution moved to microseconds, speed itself became a scarce resource that firms competed for ([Budish et al., 2015](#); [Aquilina et al., 2022](#)). The third, still underway, is *agentic*: software agents built on large language models now search for, negotiate over and pay for goods and services with little or no human involvement ([Allouah et al., 2025](#); [Zhu et al., 2025](#)). A parallel technological shift, the public blockchain, rebuilt trading infrastructure from scratch. It replaced dealers with algorithmic automated market makers (AMMs) and exchanges with block builders who decide which transactions are executed and in what order ([Daian et al., 2020](#); [Xu et al., 2023](#)).

This survey argues that the lens that makes sense of all these changes is *market design*: the explicit engineering of allocation rules subject to incentive, information and computational constraints. Viewed this way, the history of microstructure is a sequence of redesigns. Each was prompted by a technology that made the previous design’s assumptions fail, and each was bounded by the same small set of impossibility theorems. The recurring question is whether scarce goods can be allocated efficiently and fairly without requiring participants to reveal everything they know and want.

The argument in brief. Section 2 recalls the foundational result: no bilateral trading mechanism can be efficient, incentive-compatible, individually rational and budget-balanced at once ([Myerson and Satterthwaite, 1983](#)). Because this impossibility cannot be engineered away, every later design chooses which property to give up and where the loss shows up. We trace that choice through four eras.

1. **Speed** (Section 3). Continuous limit order books reward being first, creating a socially wasteful latency arms race. Frequent batch auctions change what “simultaneous” means and convert

the speed race into price competition (Budish et al., 2015). The price of that cure is timeliness: between clearing times no price can form, so periodic auctions prevent timely price discovery (Du and Zhu, 2017; Lee et al., 2026; European Securities and Markets Authority, 2019).

2. **Information** (Section 4). Once speed is neutralized, the information structure of the book becomes the lever. Recent queueing-theoretic work ranks welfare as dark pools > lit exchanges > batch auctions under moderate arrival rates (Aldridge, 2026b). This refines rather than overturns the batch-auction case. At the opaque extreme sit bilateral dealer markets, where prices come from search and bargaining; electronic requests for quote, on-chain intents and agent-to-agent negotiation are their modern descendants (Section 5).
3. **Ordering control** (Sections 6–7). Blockchains reused Hanson’s automated market maker (Hanson, 2003) and a batch architecture reminiscent of Budish et al. (2015), but gave control of each batch to a block builder. The result, maximal extractable value (MEV), is an institutionalized form of the front-running that batch auctions were designed to remove (Daian et al., 2020; Qin et al., 2022).
4. **Dimensionality** (Sections 8–9). AI agents now have protocols for tools (Anthropic, 2024), payment (Coinbase Developer Platform, 2025) and identity (Ethereum Improvement Proposals, 2025), but none of these decides who gets a scarce resource. Agents can cheaply report a few feature weights but not full rankings, so a new allocation problem appears. Spectral feature matching (Aldridge, 2026a) is one answer: its guarantee is conditional, and its failures can be checked in advance.

Section 10 draws these threads together and Section 11 concludes.

Scope and positioning. Excellent surveys exist on each strand: classical microstructure (O’Hara, 1995, 2015), high-frequency trading (Aldridge, 2013), decentralized exchanges (Xu et al., 2023), and matching theory. Our contribution is to read them as one sequence of design problems with a common structure, and to carry that structure forward into markets populated by AI agents. The survey draws on three companion pieces by the author: an analysis of lit, dark and batch mechanisms (Aldridge, 2026b), a literature review of crypto microstructure prepared for the Federal Reserve Bank of New York (Aldridge, 2025), and a working paper on multi-dimensional matching (Aldridge, 2026a). Where we report their results we say so explicitly.

2 Foundations: What a Trading Mechanism Can and Cannot Do

2.1 The double auction and the clearing price

The canonical trading mechanism is the double auction. Buyers submit bids, sellers submit asks, and the mechanism sets a price at which trade clears. Sorting bids in descending order traces out a demand schedule; sorting asks in ascending order traces out a supply schedule; a uniform clearing price sits where the two cross. Every buyer whose bid is above the price and every seller whose ask is below it trades.

The design question is whether traders will submit their true values. For one-sided auctions Vickrey (1961) gave a celebrated answer: if the winner pays the *second*-highest bid, then the price a bidder pays never depends on her own bid, so shading can only lose profitable trades and overbidding can only win unprofitable ones. Truthful bidding is a weakly dominant strategy in every state of the world.

Table 1: The Myerson–Satterthwaite impossibility forces a choice of which property to give up.

Give up	What remains	Example
Exact efficiency	Budget balance, truthfulness, participation	McAfee (1992) trade reduction
Budget balance	Efficiency, truthfulness, participation	VCG with an external subsidizer
Individual rationality	Efficiency, budget balance, truthfulness	Forced-participation schemes
Truthfulness	Efficiency, budget balance, participation	Posted-price bargaining

2.2 The two-sided impossibility

Two-sided markets are harder. With private values on both sides, each trader’s report affects the price she faces, and incentives to shade appear on both sides of the book. Chatterjee and Samuelson (1983) showed that in the bilateral case strategic bidding leaves a band of efficient trades near the diagonal of the buyer-value/seller-cost plane unconsummated. Myerson and Satterthwaite (1983) proved that this is not a failure of any particular design but a mathematical impossibility.

Result 1 (Myerson and Satterthwaite, 1983). *When a buyer’s and a seller’s privately known valuations are independently distributed with overlapping supports, no bilateral trading mechanism is simultaneously ex-post efficient, Bayesian incentive-compatible, interim individually rational and budget-balanced without an outside subsidy.*

The theorem forces a choice, summarized in Table 1. McAfee (1992) gives the best-known escape: compute the efficient number of trades, then drop the marginal buyer–seller pair if needed so that prices can be set from the excluded bids. The mechanism is dominant-strategy truthful and budget-balanced; the cost is the loss of at most one near-indifferent trade. Vickrey–Clarke–Groves mechanisms keep efficiency and truthfulness but run a deficit that someone must subsidize.

2.3 Why real double auctions work anyway

If the theory is so pessimistic, why do continuous double auctions work well in practice? Laboratory evidence answered early. Smith (1962) found that continuous double auctions reach near-full efficiency almost immediately, even with few traders and no knowledge of others’ values. Gode and Sunder (1993) went further: “zero-intelligence” traders who bid randomly, subject only to a budget constraint that forbids loss-making trades, achieve allocative efficiency close to that of human traders. Unconstrained random traders do not.

The lesson has shaped market design ever since: *market structure, not trader sophistication, does most of the work*. It also foreshadows the agentic era. If the rules of a well-designed market can extract efficiency from nearly random traders, the rules matter even more when the traders are software whose objectives and strategies may be opaque (Section 8).

2.4 Information and the market maker

The classical microstructure models add the other essential ingredient: asymmetric information. In Glosten and Milgrom (1985) a competitive dealer sets bid and ask quotes knowing that some counterparties are better informed; the spread compensates for adverse selection. In Kyle (1985)

Table 2: The latency race as a symmetric prisoner’s dilemma (Budish et al., 2015).

	Rival stays slow	Rival invests in speed
You stay slow	Both keep profits	You are picked off
You invest in speed	You pick off the rival	Both pay the cost; rents competed away

an informed trader camouflages orders among noise traders and a market maker sets prices from aggregate order flow. Both frameworks identify the *liquidity provider* as the agent who bears adverse selection and the *information structure* of the market as the determinant of prices. The same two elements reappear, in algorithmic form, in every design that follows: the high-frequency market maker, the dark pool, the AMM’s liquidity pool and the agent marketplace.

3 The Speed Era: Continuous Trading and the Arms Race

3.1 Algorithms become the market makers

Electronic continuous limit order books process each order the instant it arrives, under price-time priority. By the late 2000s, market making on most venues had been taken over by high-frequency trading (HFT) firms that post limit orders and earn the spread (Menkveld, 2013). The early evidence on liquidity was largely favorable. Algorithmic trading narrowed spreads and made quotes more informative (Hendershott et al., 2011). HFTs trade in the direction of permanent price changes and against transitory pricing errors (Brogaard et al., 2014). O’Hara (2015) argued that high-frequency markets require rethinking microstructure theory itself, since the trade is no longer the natural unit of analysis. The May 2010 Flash Crash showed the fragility side of the same ecosystem: HFTs did not trigger the crash but amplified its dynamics as liquidity evaporated (Kirilenko et al., 2017).

3.2 The latency arms race

Budish et al. (2015) identified a structural flaw in continuous trading itself. When a public signal moves the value of an asset, every liquidity provider’s quote becomes stale at once, and the first trader to reach the exchange can “snipe” it. Because correlations between closely related instruments that hold at human timescales break down at microsecond horizons, such opportunities recur mechanically, all day, every day. The resulting race has the structure of a prisoner’s dilemma (Table 2): each firm must invest in speed to avoid being picked off, and in equilibrium all firms pay the fixed cost while the arbitrage rents are competed back to the spread paid by investors.

Aquilina et al. (2022) measured these races directly using message-level data from the London Stock Exchange, including failed orders that are invisible in ordinary trade data. They find that races are frequent, last microseconds, and account for a substantial share of trading volume. The implied “latency arbitrage tax” on liquidity is economically meaningful in aggregate, which confirms that the arms race is a first-order feature of continuous markets rather than a theoretical curiosity.

3.3 Frequent batch auctions

The proposed remedy discretizes time. Under frequent batch auctions, orders arriving within a short interval are treated as simultaneous and cleared together in a uniform-price sealed-bid double auction (Budish et al., 2014, 2015). Within a batch, price-time priority is replaced by pro-rata allocation at the clearing price, so arriving a microsecond earlier confers no advantage. Discreteness




Figure 2: Continuous trading versus frequent batch auctions. Batching removes the value of arriving microseconds earlier within an interval.

converts competition on speed into competition on price. Because liquidity providers no longer need to price in the risk of being sniped, the model predicts narrower spreads.

3.4 Batch mechanisms in practice

Periodic auctions are not only a theoretical proposal. In Europe, periodic auction books run alongside continuous lit trading on several venues (for example, the Cboe Europe periodic auctions book and Deutsche Börse’s Xetra auction mechanisms). Their growth after the MiFID II dark-trading caps prompted a formal review by the European Securities and Markets Authority ([European Securities and Markets Authority, 2018, 2019](#)), whose findings on price formation we discuss below. In the United States, IEX adopted a different, non-batch response to the same race: a uniform 350-microsecond delay on inbound orders. The batch interval is itself a design parameter with costs on both sides. Longer intervals allow prices to go stale, while shorter ones leave residual sniping risk, so total cost is U-shaped in the interval length.

3.5 The cost of batching: delayed price discovery

Batching buys protection from sniping by giving up timeliness. In a continuous book every order that crosses the spread is a trade, and every trade is a price that others can observe and act on. In a periodic auction, no trade and hence no transaction price can form between clearing times. Information that arrives just after an auction can reach prices only at the next one, while quotes and resting interest are, by design, stale. *Periodic auctions therefore prevent timely price discovery*: the question is only how large the delay is and who bears it. Theory, a natural experiment in exchange redesign and regulatory evidence all point to the same cost.

Theory. [Madhavan \(1992\)](#) compares continuous and periodic call markets. Pooling orders in a call market can sustain trading under information asymmetry severe enough to shut down a continuous market, but prices are revealed only at the discrete clearing times. [Du and Zhu \(2017\)](#) solve for the trading frequency that maximizes allocative efficiency in a model of sequential double auctions. When news arrives at scheduled times, the optimal frequency matches the frequency of information arrival. When news arrives stochastically, as most market-moving information does, the optimal frequency *far exceeds* the information arrival rate, because infrequent auctions leave traders unable to respond promptly to new information. Section 4 adds a welfare channel: in [Aldridge \(2026b\)](#), batching forces even the most urgent traders to wait $T/2$ on average, and this is one of the two reasons batch auctions rank below lit exchanges.

Table 3: Three market architectures as variants of one queueing system (Aldridge, 2026b).

Architecture	Order book	Service discipline	Market-order cutoff depends on book?
Lit exchange	Publicly observable	First-come-first-served	Yes: $C^*(V, B)$
Dark pool	Hidden until trade reported	First-come-first-served	No: constant $\bar{C}^*$
Batch auction	Accumulates over interval T	Service in random order	No: $\tau^*(V, C)$

Evidence from a market redesign. The Taiwan Stock Exchange offers a rare natural experiment. Twu and Wang (2018) study the shortening of Taiwan’s call-auction interval from 25 to 5 seconds between 2010 and 2014, and find mixed effects on pricing-efficiency measures but better overall market quality at the shorter interval. Lee et al. (2026) exploit Taiwan’s 2020 switch from frequent batch auctions to continuous trading. Liquidity and price efficiency improved significantly for mid- and small-capitalization stocks, which also earned positive abnormal returns, whereas the effect on large-capitalization efficiency was marginal. Faster traders captured more of the gains, and individual investors in smaller stocks lost more. Batching thus slows price discovery most where information is scarcest and matters most: in less liquid stocks.

Regulatory evidence. ESMA’s review of European periodic auctions found that many were not genuinely price-forming: 99.9% of orders submitted to them were pegged to reference prices set elsewhere, and about 80% crossed at the midpoint of the continuous market’s spread (European Securities and Markets Authority, 2019). ESMA proposed treating mid-point-only systems as non-price-forming. Batch venues that import their prices from continuous books do not perform price discovery; they consume it.

The design implication is a genuine trade-off, not a verdict against batching. Shorter intervals reduce staleness but restore part of the speed race, and batching may remain valuable at the open and close or in thin instruments. As a general replacement for continuous trading, however, periodic auctions exchange one cost, the latency arms race, for another, delayed price discovery.

The broader lesson of the speed era is that market design can neutralize an arms race without banning speed or restricting who may trade. It does so by *redefining what counts as simultaneous*. The same lever, control over the sequencing of orders, reappears with very different consequences on blockchains (Section 7).

4 The Information Era: Lit, Dark and Batch Venues

4.1 Three venues as one queueing system

Eliminating the speed race does not by itself establish that batch auctions maximize welfare. Modern equity markets offer at least three architectures: lit exchanges, dark pools and batch auctions. Aldridge (2026b) models all three as the same queueing system with one element changed each time: the information traders have about the book and the discipline by which orders are served (Table 3).

Two foundational queueing results point in opposite directions. Leshno (2022) shows that serving agents in random order can outperform first-come-first-served because it smooths the mismatch caused by fluctuations in waiting times. That argument favors batch auctions over lit books. Che




Figure 3: Simulated aggregate welfare at baseline calibration ($n = 20,000$ traders, $\lambda = 5$, $\Delta = 2$, $T = 1$). Participation rates are 46.1%, 43.6% and 39.9%. The dark pool exceeds the lit exchange by 17% and the lit exchange exceeds batch by 52%. Source: Aldridge (2026b), Table 3.

and Tercieux (forthcoming) show that an uninformed first-come-first-served queue, which reveals nothing beyond recommendations, is optimal because hiding queue position removes the incentive to exit. That argument favors dark pools. Neither model was built for financial markets, which have an escape valve absent from classical queues: an impatient trader can bypass the queue entirely by paying the spread with a market order.

4.2 Model and welfare ranking

In Aldridge (2026b) each trader i with valuation V_i , side $s_i \in \{-1, +1\}$ and waiting cost C_i trades off execution price, waiting and a fixed cost K :

$$U_i(p, w_i) = s_i(V_i - p) - C_i w_i - K. \quad (1)$$

The specification is identical across mechanisms; only the information structure and service discipline differ. Only on the lit exchange do traders condition their strategy on the visible book state B . The paper’s main result ranks welfare.

Result 2 (Aldridge, 2026b, Theorem 2). *For arrival rates $\lambda \in [\lambda_{\min}, \lambda_{\max}]$, bounded adverse selection $\Delta < \Delta_{\max}$ and a dark-pool reporting delay $\delta < T/2$,*

$$W_{\text{DARK}} > W_{\text{LIT}} > W_{\text{BATCH}}.$$

Two distinct inefficiencies drive the ranking, and execution prices cancel from both because they are transfers between buyer and seller. *Dark beats lit* because observing the book induces costly timing games in which traders delay or rush to improve queue position; opacity removes the incentive, while urgent traders can still execute instantly with market orders. *Lit beats batch* because batching forces even the most urgent traders to wait $T/2$ on average, and pro-rata rationing pushes marginal traders out of the market altogether.

Figure 3 reports the calibrated magnitudes. Over a 20×20 grid of arrival rates and valuation dispersions, the full ranking holds almost everywhere. It fails only in very thin markets ($\lambda \lesssim 2.5$), where just $W_{\text{DARK}} > W_{\text{LIT}}$ survives, and the failure region shrinks as dispersion grows. The model also identifies the real cost of opacity. It is not that uninformed traders receive worse prices, which is only a transfer. It is that marginal traders anticipate poor pricing and exit, and the surplus from their trades is lost. This *participation channel* connects the result to the long-standing debate on whether dark pools harm price discovery (Zhu, 2014).

4.3 Regulatory relevance

Regulators on both sides of the Atlantic have recently revisited these architectures. In December 2022 the U.S. Securities and Exchange Commission proposed four equity-market-structure rules, including an Order Competition Rule that would have required certain segmented retail orders to be exposed to auctions before internalization, and a Regulation Best Execution (U.S. Securities and Exchange Commission, 2022). Both of those proposals were withdrawn in June 2025 (U.S. Securities and Exchange Commission, 2025). In April 2026 ESMA launched a call for evidence on the structure of European equity markets covering lit continuous trading, closing and frequent batch auctions, systematic internalisers, dark trading and post-trade transparency (European Securities and Markets Authority, 2026). The welfare decomposition above has a direct implication for such reviews: shortening batch intervals helps batch auctions, but transparency mandates do not, because opacity was never the batch auction’s problem. Mandating pre-trade disclosure in dark venues would remove the very feature that generates their welfare advantage. This refines rather than overturns Section 3. Batch auctions still eliminate the speed race, but on this welfare measure they are not the ceiling.

5 Bilateral Markets: Dealers, Search and Requests for Quote

5.1 Why bilateral trading persists

Exchanges are the exception, not the rule. Most corporate and municipal bonds, interest-rate and credit swaps, foreign exchange and many large crypto trades change hands *bilaterally*: a customer negotiates with a dealer, one counterparty at a time, over the counter (OTC). Bessembinder et al. (2020) survey the microstructure of these fixed-income markets. Bilateral trade is also the literal setting of the Myerson–Satterthwaite theorem (Section 2): two parties with private valuations bargain over a single trade, so some efficient trades must fail. The canonical bargaining model is the alternating-offers game of Rubinstein (1982), in which patience determines how the surplus is split. Bilateral markets sit at the opaque extreme of the venue spectrum in Section 4: there is no public book, quotes are private, and each trader sees only the prices she is offered.

5.2 Search, bargaining and dealer networks

Duffie et al. (2005) model OTC markets as search markets. Investors must find a counterparty before they can trade, and prices are set by bargaining. Bid–ask spreads therefore reflect search frictions and the relative bargaining power of investors and dealers, not only inventory and adverse selection. Spreads narrow as investors find counterparties more easily or can meet more dealers (see also Duffie, 2012). Zhu (2012) shows how opacity shapes this search. A seller who contacts dealers sequentially reveals information through repeated contact, and dealers who infer that the seller has been turned down elsewhere quote worse prices. Real OTC markets are organized as networks. Li and Schürhoff (2019) document a core–periphery structure of dealers in which trading costs and speed of execution depend on a dealer’s position in the network.

Empirically, bargaining power in opaque bilateral markets favors the dealer, especially against small customers. In municipal bonds, Green et al. (2007) find that dealers earn substantial markups that decline sharply with trade size, so retail investors pay the most. Post-trade transparency changes the balance. When the TRACE system began publishing corporate bond trade prices, execution costs fell for the newly transparent bonds (Bessembinder et al., 2006; Edwards et al., 2007). In bilateral markets, publishing prices after the fact performs part of the price-discovery work that a public order book performs before the trade.

5.3 Electronic requests for quote

Electronic trading has entered bilateral markets through the request for quote (RFQ). A customer sends a trade request to several dealers at once, and the best response wins. [Hendershott and Madhavan \(2015\)](#) show that these one-sided electronic auctions are a viable source of liquidity even in inactively traded bonds. They describe the RFQ as a natural compromise between bilateral search and the continuous double auction of an exchange. Using platform and regulatory data, [O’Hara and Zhou \(2021\)](#) find that electronic RFQ trading, though still limited in scope, has substantial effects on transaction costs and execution quality across electronic, voice and inter-dealer trading. Regulation has pushed in the same direction: the Dodd–Frank Act required many swaps to trade on swap execution facilities. Yet [Riggs et al. \(2020\)](#) find that customers of index credit default swaps typically contact only a few dealers. The benefit of wider competition is limited by a winner’s curse, since dealers who win a widely shopped request fear they have mispriced, and by customer–dealer relationships.

The RFQ is itself a periodic, sealed-bid auction triggered by one trader. It inherits both sides of Section 3.5. It concentrates competition at a moment chosen by the customer, but it produces no continuous public price, and shopping a request widely leaks information about the customer’s intentions.

5.4 Bilateral trading on-chain and between agents

Bilateral mechanisms are returning in decentralized finance. Instead of trading against an AMM pool, a user can sign an *intent*, a statement of what she wants to trade and at what worst price, and let competing “solvers” or market makers bid to fill it. These order-flow auctions are on-chain RFQs. [Bachu et al. \(2024\)](#) measure the price improvement they deliver relative to AMM execution on two major platforms, finding average gains of about 4–5 basis points that come mainly from better access to liquidity for larger trades. [Canidio and Henneke \(2024\)](#) design fair combinatorial auctions for batches of such intents. Filling an intent privately also keeps it out of the public mempool, which is a defense against the extractable value of Section 7.

Agentic markets are bilateral by default. When a user’s agent negotiates directly with a seller’s agent, the interaction is bilateral bargaining under private information, and the Myerson–Satterthwaite logic applies: some efficient trades will fail, and the split of surplus depends on bargaining power. [Zhu et al. \(2025\)](#) find that such unmediated agent-to-agent negotiations can produce systematically imbalanced outcomes. The x402 protocol of Section 8 implements the simplest bilateral mechanism of all, a posted price: the server names a price and the client pays or walks away. Posted prices are simple and give the buyer no reason to misreport, but they lose every trade in which the buyer’s value lies below the price and the seller’s cost below the buyer’s value.

Table 4 summarizes. Across all these settings, the evolution of bilateral markets repeats the pattern of the survey: technology adds competition (RFQ, intents), transparency (TRACE) or automation (agents) to one-to-one trading, but the Myerson–Satterthwaite bound means that some centralized mechanism is needed to recover the efficient trades that bilateral bargaining leaves on the table.

Table 4: Bilateral trading mechanisms, from voice dealers to agents.

Mechanism	How a price is formed	Main friction
Voice OTC dealer market	Sequential search and bilateral bargaining (Duffie et al., 2005)	Search costs; dealer bargaining power; opacity (Green et al., 2007)
Electronic RFQ	One-sided sealed-bid auction among invited dealers (Hendershott and Madhavan, 2015)	Few dealers contacted; winner’s curse; information leakage (Riggs et al., 2020)
Post-trade transparency (TRACE)	Bilateral quotes, publicly reported afterwards	Price discovery after the fact (Bessembinder et al., 2006; Edwards et al., 2007)
On-chain intents / order-flow auctions	Solvers compete to fill a signed intent (Bachu et al., 2024)	Concentration among solvers; fairness across intents (Canidio and Henneke, 2024)
Agent-to-agent negotiation	Bargaining between software agents	Imbalanced outcomes without a mediator (Zhu et al., 2025)
Posted price (x402)	Seller names a price; buyer accepts or leaves	Lost trades below the posted price

Table 5: Escape routes from the Gibbard–Satterthwaite theorem.

Restrict	Example
Outcomes to two	Majority rule is strategy-proof between two options
The preference domain	Single-peaked preferences make the median voter rule strategy-proof; low-dimensional feature preferences (Section 9)
Allow money	Vickrey–Clarke–Groves mechanisms, the route taken by auctions

6 Designing for What Participants Report: Revelation and Automated Market Makers

6.1 The revelation principle and its limits

Designing a mechanism appears to require searching an unbounded space of games: ascending auctions, sealed bids, multi-round negotiations. The revelation principle collapses the search. Any equilibrium outcome of any mechanism can be reproduced by a direct mechanism in which agents truthfully report their types and the mechanism plays their equilibrium strategies for them (Myerson, 1979, 1981). The simplification has a cost that becomes central in the agentic era. “Truthful” means truthful about *what the mechanism asks for*, and the principle does not guarantee that the truthful direct mechanism is simple, fast or computable.

Gibbard (1973) and Satterthwaite (1975) then showed that with three or more outcomes and unrestricted preferences, the only strategy-proof social choice rules are dictatorial. Escaping the theorem requires restricting one of its conditions (Table 5). Each escape route corresponds to a family of practical mechanisms, and the one relevant for AI agents, restricting the preference domain, is developed in Section 9.

Table 6: Trading mechanisms of major prediction markets.

Venue	Mechanism	Source
Iowa Electronic Markets	Continuous double auction; market and limit orders	Berg et al. (2008)
Intrade	Exchange with trader-posted bids and asks	Rothschild and Sethi (2016)
Kalshi	CFTC-designated contract market; order book of resting bids and asks	U.S. Commodity Futures Trading Commission (2020) ; Kalshi (2026)
Polymarket	Initially an AMM; now a hybrid CLOB with off-chain matching and on-chain settlement	Tsang and Yang (2026)
Hanson-style markets	Logarithmic market scoring rule (automated market maker)	Hanson (2003, 2007)

6.2 Prediction markets and the first automated market maker

Prediction markets restrict the domain to a single uncertain event, and good mechanisms become possible. A contract’s price is a forecast: the level at which no one wants to trade further. [Wolfers and Zitzewitz \(2004\)](#) show that such prices forecast outcomes about as well as expert models.

Prediction markets run on limit order books. In practice, the prediction markets that carry real volume are organized as limit order books, the double auction of Section 2 in continuous form. The Iowa Electronic Markets operate a continuous double auction with market and limit orders ([Berg et al., 2008](#)). Intrade was an exchange on which traders posted their own bids and asks, and [Rothschild and Sethi \(2016\)](#) use its transaction-level data to study trading strategies and market microstructure. The two largest venues today follow the same design (Table 6). Kalshi, designated by the CFTC as a contract market in November 2020 ([U.S. Commodity Futures Trading Commission, 2020](#)), displays resting bids and asks for each event contract ([Kalshi, 2026](#)). Polymarket began with an automated market maker and moved to a hybrid central limit order book, in which an operator matches signed orders off-chain and trades settle on-chain ([Tsang and Yang, 2026](#)).

The order book’s weakness is thin liquidity. A thinly traded contract has few resting orders, so incoming traders often find no counterparty.

[Hanson \(2003, 2007\)](#) addressed thin liquidity by replacing the matched counterparty with an algorithm. Under the logarithmic market scoring rule (LMSR), every trader transacts against a market maker whose prices derive from a smooth convex cost function over the vector q of outstanding shares in each of n outcomes:

$$C(q) = b \log \left(\sum_{i=1}^n e^{q_i/b} \right), \quad p_i(q) = \frac{\partial C}{\partial q_i} = \frac{e^{q_i/b}}{\sum_j e^{q_j/b}}. \quad (2)$$

Buying Δq costs $C(q + \Delta q) - C(q)$. Prices are always quoted and sum to one, and the market maker’s worst-case loss is bounded by $b \log n$ regardless of volume. The liquidity parameter b is a designed subsidy: larger b buys deeper liquidity at a larger, but known, maximum cost.

The LMSR is the conceptual ancestor of the constant-function market makers that dominate decentralized finance. A cost function replaces the counterparty; liquidity is always available, 24

Table 7: Indicative maximum throughput, transactions per second. Figures as compiled in Aldridge (2025) from Visa, NYSE, *Forbes* (2009) and public blockchain sources. Capacity definitions differ across venues, so comparisons are order-of-magnitude only.

System	Max. transactions per second
NASDAQ	250,000
NYSE	86,000
Visa	65,000
Ethereum (base layer)	15–30 (up to $\sim 1,000$ with Layer-2 rollups)
Bitcoin	3–7

hours a day, without human supervision; and the designer’s problem becomes choosing the shape of the curve.

The two market types have migrated in opposite directions. Decentralized spot exchanges moved from order books to AMMs (Section 7), while prediction markets, once volume arrived, moved from AMMs to order books, as Polymarket’s redesign illustrates. The pattern is consistent with the division of labor suggested by the LMSR’s bounded subsidy. An algorithmic market maker is worth its cost when counterparties are scarce, but once enough traders arrive, a limit order book lets them quote their own prices and removes the need for a subsidized counterparty.

7 On-Chain Microstructure: Blockchains, AMMs and Extractable Value

7.1 What institutional adoption requires

A blockchain is a replicated database in which each block carries a cryptographic summary of its predecessor, so altering history requires recomputing the entire chain (Haber and Stornetta, 1991; Nakamoto, 2008). Its promise for finance is 24/7 automated settlement through smart contracts (Cong and He, 2019). Aldridge (2025) distills institutional requirements into six market-quality bars: (i) fast and certain settlement; (ii) institutional-grade throughput; (iii) predictable transaction (“gas”) costs; (iv) minimal extractable value, fraud and attack surface; (v) efficient matching of buyers and sellers; and (vi) privacy of business terms. Bars (iv) and (v) are market-design problems and occupy most of this section.

Public blockchains fall short on the first two bars (Table 7). Throughput is four to five orders of magnitude below traditional venues. Layer-2 “rollups” batch transactions off-chain and post results back. Optimistic rollups assume validity unless challenged with a fraud proof during a withdrawal window of about seven days. Zero-knowledge rollups prove validity cryptographically at higher computational cost. Settlement finality varies from probabilistic and never absolute on Bitcoin, to economic finality after roughly a quarter of an hour on proof-of-stake Ethereum, to deterministic sub-second finality on newer chains designed for payments such as Circle’s Arc, which also offers opt-in shielded balances (Circle Internet Group, 2025). Traditional securities settle on a T+1 cycle but with legal certainty. Designers also face the “scalability trilemma” between throughput, security and decentralization. Buterin’s heuristic for decentralization is whether each half of a system, cut in two, could continue operating independently (Buterin, 2017). One practical compromise is quasi-decentralization: geographically distributed nodes run by a small set of vetted operators.




Figure 4: A constant-product market maker. The pool moves from P_1 to P_2 along the curve; the quoted price at each point is the slope of the tangent.

7.2 Constant-function market makers

Decentralized exchanges (DEXs) largely abandoned the order book. Instead, liquidity providers (LPs) deposit reserves (x, y) of two assets into a pool governed by a trading function, and every trade moves the pool along a curve (Angeris and Chitra, 2020; Xu et al., 2023):

$$f(x, y) = k, \quad \text{Uniswap: } x \cdot y = k, \quad p = -\frac{dy}{dx} = \frac{y}{x}. \quad (3)$$

The marginal price is the slope of the curve. Buying the risky asset depletes its reserve and raises its price along a pre-specified path (Figure 4). A convex curve approaches both axes asymptotically, so price rises without bound as a reserve becomes scarce, a useful property for an unmanned market maker. Lehar and Parlour (2025) study the equilibrium of Uniswap liquidity provision and its relation to order-book markets.

The economics of the LP resemble those of the classical dealer. Revenue is a fee per trade. Costs are inventory risk and adverse selection: arbitrageurs trade against stale pool prices whenever the external market moves. Milionis et al. (2022) formalize this cost as *loss-versus-rebalancing* (LVR), the shortfall of an LP position relative to a continuously rebalanced portfolio; the related “divergence” or “impermanent” loss measures the LP’s shortfall against passive holding. Capponi and Jia (2021) show that AMM liquidity provision can be unprofitable in volatile, thin markets. Table 8 contrasts AMMs with limit order books.


7.3 Batching returns, with a builder in charge

Ethereum’s move to proof of stake in 2022 organized transaction processing into a batch architecture. Pending transactions sit in a public mempool; every 12 seconds, a specialized *block builder* selects transactions and chooses their order; a validator proposes the block with its stake at risk; every trade then fills against an AMM pool (Figure 5). The architecture is reminiscent of frequent batch auctions, but with two differences that invert their purpose. First, there is no uniform clearing price: trades within a block execute sequentially along the AMM curve, so their *order* determines their price. Second, the party that sets the order sees the order flow first and may insert its own trades.

The consequence is *maximal extractable value* (MEV): the profit available to whoever controls inclusion and ordering. Daian et al. (2020) documented priority-gas-auction bots front-running DEX trades and warned of consensus instability. Qin et al. (2022) quantified the extracted value across sandwich attacks, arbitrage and liquidations. The canonical sandwich attack (Figure 5, bottom)

Table 8: Decentralized exchanges with AMMs versus traditional limit order books. Adapted from Aldridge (2025).

Property	DEX with AMM	Traditional exchange
Pricing	Constant-function curve	Central limit order book
Matching	One-to-pool: traders never meet each other	One-to-one: buyer meets seller
Hours	24×7	Mostly business hours
Quotes vary with	Pool inventory and liquidity	Dealer discretion, uncertainty
Price path	Pre-specified by the curve	Depends on market-maker participation
Settlement	On-chain, atomic with the trade	T+1 through clearing



A sandwich attack inside one block (illustrative prices): buy low, sell high around the user, risk-free.

Figure 5: Top: Ethereum’s post-merge transaction pipeline. Bottom: how control of intra-block ordering yields maximal extractable value.

buys ahead of a pending user order, lets the order execute at a worse price, and sells immediately after. Adams et al. (2025) decompose the realized cost of Uniswap swaps into benign slippage and adversarial “reordering slippage”. They find that gas dominates costs for small swaps, while price impact and slippage dominate for large ones. Because the curve’s price path is deterministic, a trade and its reversal need not commute, which is the formal basis of reordering attacks (Bartoletti et al., 2022).

Aldridge (2025) summarizes the irony. Periodic batching was designed to eliminate the need for continuous market makers and the rents of those who move first. On Ethereum it hosts AMMs that must fill every trade and hands first-mover rents to the builder, reproducing the front-running complaints once made about HFT. At the same time, crypto still needs AMMs, because they are what make asynchronous 24/7 trading possible when no counterparty is present.

7.4 The research response: restrain the builder, or rescue the AMM

The literature separates two problems: *who orders the batch*, and *who prices the trade*. Table 9 organizes representative proposals along this divide.

Two strands connect directly to earlier sections. First, several of the most effective MEV remedies are batch auctions with uniform pricing (Canidio and Fritsch, 2023; Chan et al., 2025): the on-chain literature has rediscovered Budish et al. (2015), now as a defense against the builder

Table 9: Representative responses to MEV and AMM unprofitability.

Proposal	Studies
<i>Restrain the builder: ordering and extraction</i>	
Formal theory of reordering attacks on AMMs	Bartoletti et al. (2022)
Batch all trades and execute at the post-batch marginal price (function-maximizing AMM), eliminating LVR and sandwiches	Canidio and Fritsch (2023)
Process each block as a batch under a potential function, making the AMM arbitrage-resilient	Chan et al. (2025)
Dynamic default slippage limits to make sandwiches unprofitable	Chemaya et al. (2023)
Verifiable ordering that minimizes intra-block price volatility	McLaughlin et al. (2024)
<i>Rescue the AMM: pricing and liquidity provision</i>	
Curves that maximize the number of trades intermediated	Goyal et al. (2023)
Optimal fees for constant-function market makers	Fritsch (2021)
Mechanism-design framework in which traders report their price expectations	Milionis et al. (2024)
Predictable loss and optimal liquidity provision via stochastic control	Cartea et al. (2024)
Missed optimization opportunities in lending, flash swaps and liquidations	Yaish et al. (2023)
Parallel execution through incentive-compatible sharded pools	Chen et al. (2024)
Programmable “hooks” that customize pools on the fly (Uniswap v4)	Bachu et al. (2025)

rather than the HFT. Second, [Milionis et al. \(2024\)](#) bring the revelation principle of Section 6 into AMM design. Each trader reports a price $\hat{p}$ drawn from a distribution known to the market maker and receives $x(\hat{p}) = g(p_0) - g(\hat{p})$ units of the risky asset for a payment $y(\hat{p})$. The designer chooses the demand curve g^* that maximizes expected profit $\mathbb{E}[y(\hat{p}) - p_0 x(\hat{p})]$, subject to an update rule $\pi(p_0, \hat{p})$ that lies between p_0 and $\hat{p}$, is non-decreasing in $\hat{p}$, and leaves the price unchanged when $\hat{p} = p_0$. The result is an order-book-like AMM in which each trader calls her own price, with truthful reporting designed in rather than hoped for.

7.5 Hiding the book on-chain

Section 4 suggests a third direction: rather than restraining whoever sees the order flow, stop publishing it. Table 10 contrasts three designs. In a “digital dark pool” layer-1 chain proposed by [Aldridge \(2025\)](#), traders submit prices *and* quantities that are not displayed on-chain, no validator orders transactions, and only execution prices are posted. The design follows the logic of [Milionis](#)

Table 10: Three designs compared by what is visible before the trade and who controls ordering.

Design	Visible before trade	Who orders trades	How trades match
Lit exchange	Full order book	Price-time priority	Buyer to seller
Ethereum with AMM	Pending orders in the mempool	Block builder, for a fee	Trader to pool
Digital dark pool L1	Nothing; only output prices are posted	A matching rule; no validators	Traders’ own prices and quantities, matched directly

et al. (2024) in that traders call their own prices, but it omits periodic batching, in line with the $W_{\text{DARK}} > W_{\text{LIT}} > W_{\text{BATCH}}$ ranking. Payments-oriented chains with opt-in privacy (Circle Internet Group, 2025) point in a similar direction for settlement. These designs are early, and their equilibrium properties on a permissionless network remain open questions.¹

8 AI Agents as Market Participants

8.1 From algorithms to learning agents

Algorithmic trading automated the *execution* of human decisions. The current wave automates the decisions themselves, first with reinforcement learning and now with agents built on large language models (LLMs). The leading concern this raises for microstructure is tacit collusion. Calvano et al. (2020) showed that Q-learning pricing algorithms in repeated oligopoly learn supra-competitive prices sustained by punishment of deviations, without communication. Dou et al. (2025) bring the question to financial markets. Informed speculators trading with reinforcement learning autonomously sustain collusive, supra-competitive profits by under-reacting to information, without agreement, communication or intent, which harms price informativeness and liquidity. They identify two mechanisms. The first is collusion through punishment threats, which arises only when price efficiency and information asymmetry are not too high. The second is collusion through homogenized learning biases (“artificial stupidity”), which persists even in efficient markets. For LLM agents, Agrawal et al. (2025) find in simulated continuous double auctions that direct communication between seller agents increases collusive tendencies, that propensities to collude differ across models, and that oversight and urgency instructions change behavior.

These results sharpen the lesson of Gode and Sunder (1993). Market structure can extract efficiency from traders who are individually naive, but that argument assumes their errors are independent. Learning agents trained on similar data, or running the same underlying model, may share their errors, and a mechanism that is robust to random traders need not be robust to correlated ones.

8.2 The agent commerce stack

For agents to transact rather than merely chat, three pieces of infrastructure are needed: a way to reach external tools and data, a way to pay without a human in the loop, and a way to be identified and held accountable. Each now has an open standard (Figure 6).

¹Disclosure: the author is affiliated with AbleBlox, which develops the digital dark pool design described here.




Figure 6: The emerging protocol stack for agent commerce. Connectivity, payment and identity are standardized; allocation of scarce resources is not.

- **Tools: the Model Context Protocol.** Released by Anthropic in November 2024, MCP defines how an agent’s client discovers and calls the tools, resources and prompts exposed by any MCP server wrapping a database, file system or API (Anthropic, 2024). It turns an M -agents-by- N -tools integration problem into an $M + N$ standardization problem. In December 2025 MCP was donated to the Agentic AI Foundation, a directed fund under the Linux Foundation co-founded by Anthropic, Block and OpenAI (Linux Foundation, 2025).
- **Payment: x402.** Launched by Coinbase in May 2025, x402 revives the long-dormant HTTP 402 **Payment Required** status code (Coinbase Developer Platform, 2025). A server answers a request for a paid resource with a 402 response listing payment terms; the client signs a stablecoin payment payload, retries the request with the payload in a header, and a facilitator verifies and settles the payment on-chain before the resource is returned. No account or API key is required.
- **Identity: ERC-8004.** The “Trustless Agents” draft standard, proposed in August 2025, defines three on-chain registries (Ethereum Improvement Proposals, 2025). An Identity Registry issues each agent an ERC-721 token that resolves to a registration file. A Reputation Registry records feedback from clients. A Validation Registry lets independent validators record checks of an agent’s work.

Two observations follow for microstructure. First, the stack settles on the infrastructure of Section 7. x402 pays in stablecoins and ERC-8004 identities live on-chain, so agent commerce inherits whatever finality, cost predictability and extractable value those chains deliver. Second, the stack is *connective* rather than *allocative*. Table 11 makes the gap explicit.

8.3 Evidence from agentic commerce

Early evidence suggests that leaving allocation to unmediated agent interaction is costly. Allouah et al. (2025) evaluate AI shopping agents choosing among products and document systematic biases and strong dependence on the underlying model, with implications for sellers and platforms. Zhu et al. (2025) model agent-to-agent negotiation and transactions in consumer markets and find that outcomes without a mediating allocation mechanism can be systematically imbalanced across agents. This is exactly the gap that classical market design was built to fill: the allocation question behind

Table 11: Which question each agent protocol answers.

Question	Answered by
Can the agent reach the resource?	MCP (Anthropic, 2024)
Can the agent pay for it?	x402 (Coinbase Developer Platform, 2025)
Is the agent who it claims to be, and has it behaved well?	ERC-8004 (Ethereum Improvement Proposals, 2025)
Who gets the resource when many agents want the same one?	Not yet answered

double auctions, batch auctions and top trading cycles reappears at the top of the agent stack, now for a non-human population.

8.4 What agents can cheaply report

The revelation principle says a designer may restrict attention to direct mechanisms, but a direct mechanism is useful only if agents can report what it asks for. Classical matching mechanisms ask for complete rankings or cardinal utilities over every object. For an agent acting on a user’s brief, a ranking over a catalog of hundreds of products is expensive to elicit and to reason about. A short vector of feature weights (price sensitivity, battery life, sound quality, and so on) is cheap. The allocation problem for agents is therefore a problem of *low-dimensional reports*: how to aggregate many feature-weight vectors into one fair and efficient assignment of scarce objects.

9 Allocation Without Rankings: Matching Mechanisms for Agents

9.1 Allocation without money

Some goods cannot, legally or ethically, be allocated by price. [Shapley and Scarf \(1974\)](#) posed the housing market, in which each agent owns one house and may prefer another’s. Gale’s Top Trading Cycles (TTC) algorithm has each agent point to the owner of her most preferred remaining house, trades along every resulting cycle, removes those agents, and repeats. TTC is Pareto efficient, individually rational and strategy-proof, and [Ma \(1994\)](#) proved it is the *unique* mechanism with all three properties. The same cycle logic underlies school choice ([Abdulkadiroğlu and Sönmez, 2003](#)) and kidney exchange, where incompatible patient–donor pairs swap donors along cycles, and chains started by altruistic donors clear the pool further ([Roth et al., 2004, 2005](#)). Deferred acceptance ([Gale and Shapley, 1962](#)) plays the corresponding role for two-sided markets. TTC shows that removing money does not require giving up efficiency or truthful incentives.

With cardinal utilities the picture changes. [Hylland and Zeckhauser \(1979\)](#) allocate probability shares through a pseudo-market with budgets and prices, obtaining ex-ante efficiency but not strategy-proofness. [Zhou \(1990\)](#) proved that no mechanism for cardinal one-sided matching is simultaneously ex-ante efficient, strategy-proof and symmetric. Subsequent work seeks approximate versions of all three, usually measured by Nash social welfare (NSW), the product of agents’ gains, which balances efficiency and fairness ([Nash, 1950](#); [Caragiannis et al., 2019](#)). Table 12 summarizes representative mechanisms. All of the cardinal ones require agents to report complete utilities over objects.

Table 12: Cardinal matching mechanisms and what they require agents to report. Adapted from Aldridge (2026a).

Mechanism	Agents port	re-	Method	Welfare guarantee	Incentives
Hylland and Zeckhauser (1979)	Utilities over objects	over	Pseudo-market with budgets and prices	Ex-ante Pareto optimal	Strategic bidding possible
Chawla et al. (2010)	Multi-parameter valuations		Sequential posted prices, one dimension at a time	$O(1)$ approximation	Truthful in expectation
Devanur et al. (2015)	Utilities over objects	over	Random assignment within a configuration LP	2-approximation to NSW	Envy-free in expectation
Abebe et al. (2020)	Utilities over objects	over	Random sampling and configuration LPs	$O(1)$ approximation to NSW	Truthful in expectation
Spectral feature matching (Aldridge, 2026a)	Feature weights		One SVD projection, then a sort	Exact projected-NSW optimum; conditional true-NSW bound	Stable to noise; not strategy-proof

9.2 A feature-based model

Consider I agents and J objects with capacities M_j summing to I . Objects are described by feature vectors $f_j \in \mathbb{R}^X$ and agents by preference weights $u_i \in \mathbb{R}^X$, with X small. Utility is the inner product

$$U_{ij} = u_i \cdot f_j = \sum_{x=1}^X u_{ix} f_{jx}, \quad (4)$$

which is additive across features and so rules out complementarities. An agent reports w_i , which may differ from u_i . Fairness is measured relative to uniform random assignment. Agent i 's disagreement point is $o_i = \frac{1}{J} \sum_j U_{ij}$, and

$$\text{NSW}(P \mid U) = \prod_{i=1}^I (\mathbb{E}[U_i \mid P] - o_i), \quad (5)$$

defined to be zero if any agent falls below her disagreement point. An NSW maximizer is Pareto efficient, leaves no agent worse off than random assignment, and is invariant to positive rescaling of any agent's utilities. For reporting, Aldridge (2026a) fixes a clipped version, $\text{NSW}_\epsilon = \prod_i \max(\text{gain}_i, \epsilon)$ with $\epsilon = 0.01$, which is always disclosed alongside the rate of individual-rationality (IR) violations. This matters because clipping makes NSW_ϵ optimistic precisely when agents are being harmed.

9.3 Spectral projection

Collect object features in $F \in \mathbb{R}^{J \times X}$ with singular value decomposition $F = U \Sigma V^\top$. By the Eckart–Young theorem (Eckart and Young, 1936), the leading right singular vector v_1 is the direction of maximum variance of the object features, and the rank-one approximation $\sigma_1 u_1 v_1^\top$ minimizes reconstruction error. The mechanism projects both sides onto v_1 and sorts.

Algorithm 1 Spectral feature matching (Aldridge, 2026a)

Require: feature matrix $F \in \mathbb{R}^{J \times X}$; reports $w_1, \dots, w_I \in \mathbb{R}^X$; capacities M_j

- 1: Compute v_1 , the leading right singular vector of F (oriented so that its entries are predominantly positive)
 - 2: Gauge-fix reports: $\hat{w}_i \leftarrow w_i / \|w_i\|$ $\triangleright$ removes dependence on report scale
 - 3: Project: $\tilde{f}_j \leftarrow f_j \cdot v_1$ for all j ; $\hat{a}_i \leftarrow \hat{w}_i \cdot v_1$ for all i
 - 4: **Round diagnostic:** $n_+ \leftarrow \#\{i : \hat{a}_i > 0\}$; $C_+ \leftarrow \sum_{j: \tilde{f}_j > \bar{f}} M_j$ where $\bar{f}$ is the mean projected score
 - 5: **if** $n_+ \neq C_+$ **then return** FAIL $\triangleright$ no deterministic allocation can be individually rational for all
 - 6: **end if**
 - 7: Within each sign class, sort agents by $\hat{a}_i$ and objects by $\tilde{f}_j$ and match in order, respecting capacities
 - 8: **Margin diagnostic:** report μ^* , the minimum projected gain, against 4Δ with $\Delta = \sqrt{X-1} \sigma_2$
-

The procedure runs in $O(\min(J^2 X, JX^2) + I \log I + J \log J)$ time, dominated by the SVD and two sorts. It is deterministic and symmetric by construction: identical reports receive identical treatment.

The gauge. A naive version that projects raw reports has a scale bug. Because $\hat{a}_i$ scales with $\|w_i\|$, rescaling one agent’s report without changing her preferences can flip the entire allocation and move NSW from positive to zero. Normalizing reports to unit length before projecting (a “gauge” choice) removes the dependence and restores the scale invariance that NSW itself enjoys.

A worked example. Take three products with features (performance, design) $f_1 = (7.65, 1.82)$, $f_2 = (5.45, 3.62)$, $f_3 = (3.42, 1.93)$, and three agents reporting $w_1 = (8, 3)$, $w_2 = (6, 7)$, $w_3 = (5, 4)$ (Aldridge, 2026a, Example 8.1). The singular values are $\sigma_1 = 10.80$ and $\sigma_2 = 1.83$, and $v_1 = (0.92, 0.38)$: products differ mainly in performance (Figure 7). Projected product scores are 7.76, 6.42 and 3.90. Raw projected agent scores are 8.54, 8.23 and 6.15, so plain sorting assigns $A_1 \rightarrow P_1$, $A_2 \rightarrow P_2$, $A_3 \rightarrow P_3$. After the gauge fix the scores become 0.999, 0.893 and 0.961, and agents 2 and 3 swap places, which illustrates how much report scale matters.

The example also illustrates the mechanism’s central negative result. All three agents rank the products identically ($P_1 \succ P_2 \succ P_3$), so every allocation must hand the weakest product to someone. Table 13 enumerates all six assignments: none leaves every agent above her disagreement point, and every assignment has NSW equal to zero. The round diagnostic detects this in advance. All three gauge-fixed scores are positive ($n_+ = 3$), but only two products lie above the mean projected score ($C_+ = 2$), so the algorithm returns FAIL.

9.4 Guarantees and diagnostics

Aldridge (2026a) establishes the following properties for the gauge-fixed mechanism. We state them informally and refer to the paper for precise conditions.

1. **Exact projected optimum.** Within the projected one-dimensional space, sorted matching by sign class computes the exact NSW-maximizing allocation. The result is an equality, not an approximation.
2. **Unconditional utilitarian guarantee.** Total true utility is bounded relative to the utilitarian optimum without any individual-rationality claim.




Figure 7: Worked example: products (squares) and agent reports (circles) projected onto the leading singular vector v_1 of the product-feature matrix. Values recomputed from the example in Aldridge (2026a).

Table 13: All deterministic assignments in the worked example. Gains are measured against each agent’s random-assignment disagreement point $o = (51.42, 50.24, 37.36)$. No assignment is individually rational for all three agents. The plain sorted assignment (first row) maximizes total utility.

Assignment (A_1, A_2, A_3)	Gain A_1	Gain A_2	Gain A_3	Total utility
P_1, P_2, P_3	15.24	7.80	−12.54	149.52
P_1, P_3, P_2	15.24	−16.21	4.37	142.42
P_2, P_1, P_3	3.04	8.40	−12.54	137.92
P_2, P_3, P_1	3.04	−16.21	8.17	134.02
P_3, P_1, P_2	−18.27	8.40	4.37	133.52
P_3, P_2, P_1	−18.27	7.80	8.17	136.72

- Conditional NSW guarantee.** A multiplicative bound on true NSW holds when the feasibility condition $n_+ = C_+$ is met and the minimum projected gain clears a margin proportional to the projection error, $\mu^* \geq 4\Delta$ with $\Delta = \sqrt{X-1} \sigma_2$.
- Necessity.** When $n_+ \neq C_+$, *no* deterministic mechanism can guarantee positive NSW. Agents who agree too strongly about which objects are best leave no room for an allocation that beats random assignment for everyone. This is the same structural obstacle, a shared ranking, that drives the impossibility results of Section 2.
- Incentives.** The mechanism is stable to exogenous reporting noise but is *not* strategy-proof, and the paper gives an explicit profitable misreport.

Two diagnostics make the conditional guarantee usable. At the *market* level, the share of variance on the leading direction, $\rho_1 = \sigma_1^2 / \sum_\ell \sigma_\ell^2$, indicates whether “generically better” is a meaningful axis; the paper uses $\rho_1 \geq 0.5$ as a deployment threshold. At the *round* level, the feasibility test $n_+ = C_+$ and the margin test $\mu^* \geq 4\Delta$ are computed from the algorithm’s own output before any true

utilities are observed. Contrast Myerson–Satterthwaite: there the efficiency loss is unconditional and unmeasured. Here it is conditional, and the failure regime is identified exactly and can be checked in advance.

9.5 Evidence from a simulated agentic-shopping market

Aldridge (2026a) tests the mechanism on the application motivating this section: LLM-style shopping agents, each translating a user’s brief into feature weights over six headphone attributes, compete for a limited drop of ten products, one unit each.

A single drop. The full 24-product catalog looks comfortably low-rank ($\rho_1 = 0.912$, effective rank 2.24, with v_1 loading positively on all six features). The ten-product drop still passes the market-level threshold ($\rho_1 = 0.871$), but its projection-error bound $\Delta = 35.2$ is large relative to individual agents’ gains. This is something the market-level ρ_1 cannot reveal. With only mainstream shoppers (Scenario A), all ten gauge-fixed scores are positive but only six products lie above the mean ($n_+ = 10 \neq C_+ = 6$), and the mechanism correctly returns FAIL. With four contrarian shoppers substituted (Scenario B), feasibility holds ($n_+ = C_+ = 6$) and total utility reaches 99.1% of the utilitarian optimum (1,467.6 versus 1,480.8). The margin diagnostic nevertheless fails by three orders of magnitude ($\mu^* = 0.26$ against $4\Delta = 140.8$). It correctly predicts that one agent ends below random assignment and that the mechanism’s NSW_ϵ reaches only about 0.27% of the true optimum found by brute force over all $10! \approx 3.6$ million assignments. A blind search over 200,000 alternative reports finds a misreport that improves the harmed agent’s true gain by 33.6. Truthful reporting is not a best response.

One hundred random markets. Across 100 independently generated markets, the exact feasibility test passes in 22 trials. This is a knife-edge integer condition, neither routine nor rare. Pooled per-agent gains shift clearly rightward relative to random allocation (Figure 8): the mean gain is 19.44 versus -0.02 under random assignment (zero by construction), and the share of agents below their disagreement point falls from 50.3% to 15.6%. On mean $\log \text{NSW}_\epsilon$ (Figure 9) the mechanism beats random assignment by about 11 log-points and is statistically indistinguishable from serial dictatorship. Serial dictatorship uses true, unprojected utilities and requires one best response per agent, but carries no formal guarantee.

Deployment rules. The evidence supports four rules for agentic commerce. First, diagnose every round, not just the market, because a small, scarce drop can be far worse conditioned than the catalog it comes from. Second, run both round checks before allocating. Third, fall back to exact search when they fail; for a ten-unit drop, brute force took 16.5 seconds, affordable for small, high-stakes rounds. Fourth, do not assume that agents report truthfully: production systems need an incentive layer or explicit disclosure to users.

10 Synthesis: What Each Design Buys and What It Costs

Table 14 lines up the designs surveyed. The throughline is that every gain in speed, scale or simplicity is paid for somewhere else in the mechanism.

Five themes recur across the eras.




Figure 8: Per-agent gains over the disagreement point, pooled over 100 random markets; the mechanism (blue) versus random allocation (orange). Source: Aldridge (2026a), Figure 3.




Figure 9: Mean $\log \text{NSW}_\epsilon$ ($\epsilon = 0.01$) with 95% confidence intervals over 100 markets. Mean IR violations per 10 agents: 5.03, 1.48, 1.56 and 1.28. Source: Aldridge (2026a), Table 6.

1. **Impossibility relocates rather than disappears.** The theorems of Myerson and Satterthwaite, of Gibbard and Satterthwaite, and of Zhou recur in new clothing. The designer’s real choice is where the loss falls and whether it is visible.
2. **The binding constraint moves with technology.** For human floor traders the constraint was rationality and information, and Code and Sunder (1993) showed that structure compensates for the first. For HFT it was speed. On blockchains it is control over ordering. For AI agents it is the dimensionality of what can be reported.
3. **Whoever controls the queue captures the rent.** The first arrival in a continuous book, the builder of an Ethereum block and the agent that negotiates without a mediator all profit from sequencing. Batch auctions, opaque books and sort-based matching are different ways of neutralizing sequencing power.
4. **Information structure is a design lever, not a given.** Hiding the book raises welfare in equity markets (Aldridge, 2026b), and publishing the mempool creates MEV on-chain. The same lever should be available to designers of agent marketplaces.
5. **Diagnostics complement guarantees.** When exact guarantees are impossible, the next best thing is a mechanism that says, before the outcome, whether its guarantee holds. The round-level diagnostics of Section 9 are one instance; pre-trade measures of MEV exposure or of participation risk in dark venues would be others.

11 Conclusion

Market microstructure has always been a discipline of trade-offs imposed by impossibility theorems and relaxed by clever design. The age of AI does not change that logic, but it changes its inputs. Participants are faster, more numerous and more correlated. They report less than full preferences. They transact over protocols that standardize connectivity, payment and identity while leaving allocation unaddressed. The designs surveyed here, from Vickrey’s second-price rule to spectral feature matching, show that markets can be made to work under each new constraint, provided the designer is explicit about which property is being given up and makes the resulting loss visible.

Table 14: Eight designs, one throughline: every gain is paid for elsewhere in the mechanism.

Design	What it buys	What it costs
Double auction	Truthfulness or budget balance	Exact efficiency (Myerson and Satterthwaite, 1983)
Frequent batch auction	No latency arms race	Timely price discovery; waiting costs; often not price-forming in practice
Dark pool	No book-timing games; highest welfare at moderate arrival rates	Participation of uninformed marginal traders
Bilateral / RFQ market	Trading in illiquid, customized instruments; discretion	Search costs, dealer bargaining power, no continuous public price
LMSR / AMM	Liquidity with thin or asynchronous trading	A bounded subsidy; loss-versus-rebalancing for LPs; displaced by order books once volume arrives
On-chain AMM with builder	24×7 atomic settlement	MEV for whoever orders the batch
Top trading cycles	Efficiency and truthfulness without money	Requires full ordinal rankings; no transfers
Feature matching	Speed and cheap elicitation for agents	A conditional, disclosed guarantee; manipulable

For markets populated by AI agents, that means choosing information structures and sequencing rules deliberately, preferring mechanisms that announce in advance when their guarantees fail, and treating allocation as a first-class layer of agent infrastructure.

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D. Hartline. A truthful cardinal mechanism for one-sided matching. arXiv preprint arXiv:1903.07797, 2020.
- Austin Adams, Benjamin Y. Chan, Sarit Markovich, and Xin Wan. Don’t let MEV slip: The costs of swapping on the Uniswap protocol. In *Financial Cryptography and Data Security (FC 2024)*, volume 14744 of *Lecture Notes in Computer Science*, pages 172–191. Springer, 2025. doi: 10.1007/978-3-031-78676-1_10.
- Kushal Agrawal, Verona Teo, Juan J. Vazquez, Sudarsh Kunnavakkam, Vishak Srikanth, and Andy Liu. Evaluating LLM agent collusion in double auctions. arXiv preprint arXiv:2507.01413, 2025.
- Irene Aldridge. *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Wiley, Hoboken, NJ, 2nd edition, 2013.
- Irene Aldridge. Crypto: Literature review, best practices and microstructure. Presentation to the Federal Reserve Bank of New York, September 2025.

- Irene Aldridge. Multi-dimensional matching. Working paper, 2026a.
- Irene Aldridge. Comparing market mechanism efficiencies. arXiv preprint arXiv:2605.31072, 2026b. URL <https://arxiv.org/abs/2605.31072>.
- Amine Allouah, Omar Besbes, Josué D. Figueroa, Yash Kanoria, and Akshit Kumar. What is your AI agent buying? Evaluation, biases, model dependence, and emerging implications for agentic e-commerce. arXiv preprint arXiv:2508.02630, 2025.
- Guillermo Angeris and Tarun Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT '20)*, pages 80–91, 2020.
- Anthropic. Model context protocol specification. Open standard, released November, November 2024. URL <https://modelcontextprotocol.io>.
- Matteo Aquilina, Eric Budish, and Peter O’Neill. Quantifying the high-frequency trading “arms race”. *Quarterly Journal of Economics*, 137(1):493–564, 2022.
- Brad Bachu, Xin Wan, and Ciamac C. Moallemi. Quantifying price improvement in order flow auctions. arXiv preprint arXiv:2405.00537, 2024.
- Brad Bachu, Joel Hasbrouck, Fahad Saleh, and Xin Wan. An overview of Uniswap v4 for researchers. SSRN Working Paper 5152215, 2025. URL <https://ssrn.com/abstract=5152215>.
- Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. A theory of automated market makers in DeFi. *Logical Methods in Computer Science*, 18(4), 2022.
- Joyce Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz. Results from a dozen years of election futures markets research. In Charles R. Plott and Vernon L. Smith, editors, *Handbook of Experimental Economics Results*, volume 1, chapter 80. Elsevier, Amsterdam, 2008.
- Hendrik Bessembinder, William Maxwell, and Kumar Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288, 2006.
- Hendrik Bessembinder, Chester S. Spatt, and Kumar Venkataraman. A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis*, 55(1):1–45, 2020.
- Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306, 2014.
- Eric Budish, Peter Cramton, and John Shim. Implementation details for frequent batch auctions: Slowing down markets to the blink of an eye. *American Economic Review: Papers & Proceedings*, 104(5):418–424, 2014.
- Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621, 2015.
- Vitalik Buterin. The meaning of decentralization. Medium essay, 2017. URL <https://medium.com/@VitalikButerin/the-meaning-of-decentralization-a0c92b76a274>.
- Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–3297, 2020.

- Andrea Canidio and Robin Fritsch. Arbitrageurs’ profits, LVR, and sandwich attacks: Batch trading as an AMM design response. arXiv preprint arXiv:2307.02074, 2023.
- Andrea Canidio and Felix Henneke. Fair combinatorial auction for blockchain trade intents: Being fair without knowing what is fair. arXiv preprint arXiv:2408.12225, 2024.
- Agostino Capponi and Ruizhe Jia. The adoption of blockchain-based decentralized exchanges. arXiv preprint arXiv:2103.08842, 2021.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.
- Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralized finance and automated market making: Predictable loss and optimal liquidity provision. *SIAM Journal on Financial Mathematics*, 15(3):931–959, 2024.
- T.-H. Hubert Chan, Ke Wu, and Elaine Shi. Mechanism design for automated market makers. In *7th Conference on Advances in Financial Technologies (AFT 2025)*, volume 354 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:22, 2025. doi: 10.4230/LIPIcs.AFT.2025.7.
- Kalyan Chatterjee and William Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–851, 1983.
- Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC ’10)*, pages 311–320, 2010.
- Yeon-Koo Che and Olivier Tercieux. Optimal queue design. *Journal of Political Economy*, forthcoming. doi: 10.1086/741622. Working paper version: arXiv:2307.07746.
- Nir Chemaya, Dingyue Liu, Robert McLaughlin, Nicola Ruaro, Christopher Kruegel, and Giovanni Vigna. The power of default: Measuring the effect of slippage tolerance in decentralized exchanges. SSRN Working Paper 4585229; Financial Cryptography and Data Security (FC 2024), 2023. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4585229.
- Hongyin Chen, Amit Vaisman, and Ittay Eyal. SAMM: Sharded automated market maker. arXiv preprint arXiv:2406.05568, 2024.
- Circle Internet Group. Introducing Arc: An open layer-1 blockchain purpose-built for stablecoin finance. Company announcement, August 12, August 2025. URL <https://www.circle.com/blog/introducing-arc-an-open-layer-1-blockchain-purpose-built-for-stablecoin-finance>.
- Coinbase Developer Platform. Introducing x402: A new standard for internet-native payments. Protocol announcement, May 2025. URL <https://www.coinbase.com/developer-platform/discover/launches/x402>.
- Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754–1797, 2019.
- Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927, 2020.

- Nikhil R. Devanur, Jason D. Hartline, and Qiqi Yan. Envy freedom and prior-free mechanism design. *Journal of Economic Theory*, 156:103–143, 2015.
- Winston Wei Dou, Itay Goldstein, and Yan Ji. AI-powered trading, algorithmic collusion, and price efficiency. NBER Working Paper 34054, National Bureau of Economic Research, 2025.
- Songzi Du and Haoxiang Zhu. What is the optimal trading frequency in financial markets? *Review of Economic Studies*, 84(4):1606–1651, 2017.
- Darrell Duffie. *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets*. Princeton University Press, Princeton, NJ, 2012.
- Darrell Duffie, Nicolae Gârleanu, and Lasse Heje Pedersen. Over-the-counter markets. *Econometrica*, 73(6):1815–1847, 2005.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Amy K. Edwards, Lawrence E. Harris, and Michael S. Piwowar. Corporate bond market transaction costs and transparency. *Journal of Finance*, 62(3):1421–1451, 2007.
- Ethereum Improvement Proposals. ERC-8004: Trustless agents. Draft standard, August 2025. URL <https://eips.ethereum.org/EIPS/eip-8004>.
- European Securities and Markets Authority. Call for evidence: Periodic auctions for equity instruments. ESMA press release ESMA71-99-1054, 2018. URL https://www.esma.europa.eu/sites/default/files/library/esma71-99-1054_esma_launches_call_for_evidence_on_periodic_auctions_for_equity_instruments_0.pdf.
- European Securities and Markets Authority. Final report: Call for evidence on periodic auctions for equity instruments. ESMA70-156-1035, 11 June, June 2019. URL <https://www.esma.europa.eu/document/final-reportcall-evidence-periodic-auctions>.
- European Securities and Markets Authority. Call for evidence on the market structure of European equity markets. ESMA74-1119406008-1578, 30 April, April 2026. URL https://www.esma.europa.eu/sites/default/files/2026-04/ESMA74-1119406008-1578_Call_for_Evidence_on_on_the_market_structure_of_European_equity_markets_0.pdf.
- Robin Fritsch. A note on optimal fees for constant function market makers. In *Proceedings of the 2021 ACM CCS Workshop on Decentralized Finance and Security (DeFi '21)*, pages 9–14, 2021. doi: 10.1145/3464967.3488589.
- David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- 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.
- Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

- Mohak Goyal, Geoffrey Ramseyer, Ashish Goel, and David Mazières. Finding the right curve: Optimal design of constant function market makers. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, pages 783–812, 2023. doi: 10.1145/3580507.3597688.
- Richard C. Green, Burton Hollifield, and Norman Schürhoff. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, 20(2):275–314, 2007.
- Stuart Haber and W. Scott Stornetta. How to time-stamp a digital document. *Journal of Cryptology*, 3(2):99–111, 1991.
- Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1): 107–119, 2003.
- Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.
- Terrence Hendershott and Ananth Madhavan. Click or call? Auction versus search in the over-the-counter market. *Journal of Finance*, 70(1):419–447, 2015. doi: 10.1111/jofi.12164.
- Terrence Hendershott, Charles M. Jones, and Albert J. Menkveld. Does algorithmic trading improve liquidity? *Journal of Finance*, 66(1):1–33, 2011.
- Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.
- Kalshi. The orderbook. Kalshi Help Center, 2026. URL <https://help.kalshi.com/en/articles/13823828-the-orderbook>.
- Andrei Kirilenko, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. The flash crash: High-frequency trading in an electronic market. *Journal of Finance*, 72(3):967–998, 2017.
- Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
- Yi-Tsung Lee, Roberto Riccò, and Kai Wang. Frequent batch auctions vs. continuous trading: Evidence from Taiwan. *Journal of Financial Markets*, page 101082, 2026. SSRN working paper 3733682.
- Alfred Lehar and Christine A. Parlour. Decentralized exchange: The Uniswap automated market maker. *Journal of Finance*, 2025. doi: 10.1111/jofi.13405.
- Jacob D. Leshno. Dynamic matching in overloaded waiting lists. *American Economic Review*, 112(12):3876–3910, 2022.
- Dan Li and Norman Schürhoff. Dealer networks. *Journal of Finance*, 74(1):91–144, 2019.
- Linux Foundation. Linux foundation announces the formation of the agentic AI foundation (AAIF), anchored by new project contributions including Model Context Protocol (MCP), goose and AGENTS.md. Press release, December 9, December 2025. URL <https://www.linuxfoundation.org/press/linux-foundation-announces-the-formation-of-the-agentic-ai-foundation>.
- Jinpeng Ma. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):75–83, 1994.

- Ananth Madhavan. Trading mechanisms in securities markets. *Journal of Finance*, 47(2):607–641, 1992.
- R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2): 434–450, 1992.
- Robert McLaughlin, Nir Chemaya, Dingyue Liu, and Dahlia Malkhi. CLVR ordering of transactions on AMMs. arXiv preprint arXiv:2408.02634, 2024.
- Albert J. Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.
- Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. arXiv preprint arXiv:2208.06046, 2022.
- Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. A Myersonian framework for optimal liquidity provision in automated market makers. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 81:1–81:19, 2024.
- Roger B. Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–73, 1979.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.
- Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. White paper, 2008. URL <https://bitcoin.org/bitcoin.pdf>.
- John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- Maureen O’Hara. *Market Microstructure Theory*. Blackwell, Cambridge, MA, 1995.
- Maureen O’Hara. High frequency market microstructure. *Journal of Financial Economics*, 116(2): 257–270, 2015.
- Maureen O’Hara and Xing (Alex) Zhou. The electronic evolution of corporate bond dealers. *Journal of Financial Economics*, 140(2):368–390, 2021.
- Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214, 2022.
- Lynn Riggs, Esen Onur, David Reiffen, and Haoxiang Zhu. Swap trading after Dodd-Frank: Evidence from index CDS. *Journal of Financial Economics*, 137(3):857–886, 2020.
- Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.
- Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. A kidney exchange clearinghouse in New England. *American Economic Review: Papers & Proceedings*, 95(2):376–380, 2005.
- David M. Rothschild and Rajiv Sethi. Trading strategies and market microstructure: Evidence from a prediction market. *Journal of Prediction Markets*, 10(1):1–29, 2016.

- Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
- Mark A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2): 187–217, 1975.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.
- Vernon L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137, 1962.
- Kwok Ping Tsang and Zichao Yang. The anatomy of a blockchain prediction market: Polymarket in the 2024 U.S. presidential election. arXiv preprint arXiv:2603.03136, 2026.
- Mia Twu and Jianxin Wang. Call auction frequency and market quality: Evidence from the Taiwan Stock Exchange. *Journal of Asian Economics*, 57:53–62, 2018.
- U.S. Commodity Futures Trading Commission. CFTC designates KalshiEX LLC as a contract market. Press Release 8302-20, November 4, November 2020. URL <https://www.cftc.gov/PressRoom/PressReleases/8302-20>.
- U.S. Securities and Exchange Commission. Equity market structure proposals: Order competition rule, regulation best execution, amendments to rule 605, and minimum pricing increments and access fees. Proposed rules released December 14, December 2022.
- U.S. Securities and Exchange Commission. Withdrawal of proposed regulatory actions. Notice withdrawing fourteen proposed rules, including the Order Competition Rule and Regulation Best Execution, June 12, June 2025.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2): 107–126, 2004.
- Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. SoK: Decentralized exchanges (DEX) with automated market maker (AMM) protocols. *ACM Computing Surveys*, 55(11):238:1–238:50, 2023.
- Aviv Yaish, Maya Dotan, Kaihua Qin, Aviv Zohar, and Arthur Gervais. Suboptimality in DeFi. Cryptology ePrint Archive, Paper 2023/892, 2023. URL <https://eprint.iacr.org/2023/892>.
- Lin Zhou. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990.
- Haoxiang Zhu. Finding a good price in opaque over-the-counter markets. *Review of Financial Studies*, 25(4):1255–1285, 2012.
- Haoxiang Zhu. Do dark pools harm price discovery? *Review of Financial Studies*, 27(3):747–789, 2014.

Shenzhe Zhu, Jiao Sun, Yi Nian, Tobin South, Alex Pentland, and Jiaxin Pei. The automated but risky game: Modeling agent-to-agent negotiations and transactions in consumer markets. In *Proceedings of the Natural Legal Language Processing Workshop 2025*. Association for Computational Linguistics, 2025. arXiv:2506.00073.