AIxMarkets: Blessing or Curse?

Konstantinos (Kostas) E. Zachariadis

Professor of Financial Economics and Director of Research
Queen Mary University of London (QMUL)

Research Associate
Hellenic Observatory, Financial Markets Group & Systemic Risk Centre
London School of Economics (LSE)

New Challenges in Financial and Energy Markets: Math, Data & AI
31 October 2025
Organised by NTUA & Athens Exchange Group

Why This Matters Now

A Changing Market Reality

- A large share of order flow on major venues is algorithmic (share varies by venue and instrument).
- AI systems increasingly make allocation, pricing, and compliance decisions.
- Yet our **economic and market-design theories** still assume human agents.

Framing Question

What happens when markets are run not just *with* AI tools, but *by* AI agents?

Roadmap

1. The Blessing: AI as a Tool
2. The Curse: AI as Agent

The Two Faces of AI in Economics

An Economist's View of Modern AI

We have moved beyond simple automation to **end-to-end, goal-oriented optimization** in a much larger modeling space. This creates two roles for AI.

The Blessing: AI as a Tool

- **Tackling complexity:** Deep RL for hard objectives (e.g., portfolio optimization).
- **Counterfactuals:** Generative models (GANs, diffusion) for realistic what-ifs.
- **Structural estimation:** Data-driven solutions to previously intractable questions.

The Curse: AI as an Agent

- **New agents:** Not just human proxies; AI agents have distinct behaviors.
- **Bias at scale:** Human-like biases can be embedded and scaled.
- **Interactions:** Rich dynamics among AIs and humans need study.

Our Roadmap

First: the *Blessing* (AI as a tool). Then: the *Curse* (AI as agent).

Part 1: The Blessing

AI in Quant Finance: Pricing and Portfolio Choice (based on the work of Yongxin Yang @ QMUL EECS)

Option Pricing: Accuracy & Speed

- **Normalizing Flows**: $\sim 10\times$ lower pricing error.
- **Neural surrogates**: $\sim 300\times$ faster evaluation.
- **Hypernetworks**: full surface from ~ 9 samples.

Portfolio Choice: Sparse Replication

- Track a large index with a small, fixed subset (non-convex).
- **Stochastic NNs**: $\sim 10\times$ faster than traditional solvers.
- **Meta-learning**: learns how to learn from history.

Limitation

These are computational wins; the underlying asset models are often standard.

Part 1: The Blessing

Case: Bridging Data Silos with Federated Learning (based on the work of Ahmed Sayed @ QMUL EECS)

The Problem

Collaboration in finance is needed, but data is siloed and private.

- Examples: anti-money laundering, credit risk.
- Laws prevent direct data sharing.

Federated Learning (FL)

Train models across institutions *without sharing raw data*.

- Private FL systems for AML (e.g., industry pilots).
- FL for credit scoring: reported AUC improvements in deployments.

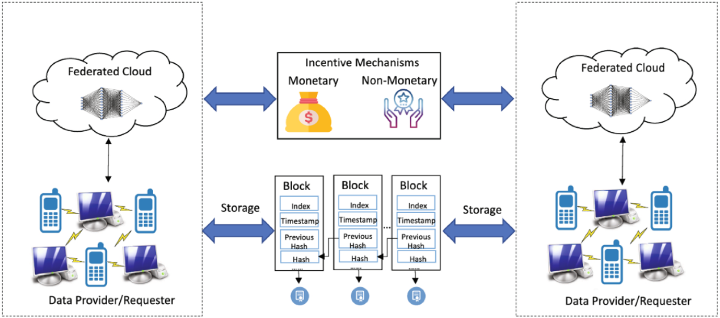

Source: Data Marketplaces (Springer, 2024)

From Trading Assets to Trading Knowledge

The Economic Shift Behind AIxMarkets

From Assets to Knowledge

Markets now trade not just goods and securities, but **knowledge itself** — data, models, and algorithms.

- These are hard to value *ex ante* because value emerges only in use or through combination.
- They are often **complementary**: one firm's model + another's dataset create value only together.
- Revealing enough information to price knowledge may destroy its private value (**Arrow's paradox**).

The Challenge

We need new market mechanisms to enable the exchange of knowledge — while preserving privacy and incentives.

Designing a Market for Knowledge (MARKSE)

From Concept to Mechanism

Mechanism Sketch: Combinatorial Double Auction (CDA)

- **Combinatorial:** allows bids on **bundles** (handles complementarities).
- **Double auction:** two–sided price discovery (reduces uncertainty through competition).

The Role of AI

- **AI as a solver:** speeds up winner determination / clearing (MILP + neural heuristics).
- **AI as an agent:** populates the market with MARL traders to test price discovery and efficiency.

Punchline

Use AI to both **solve** realistic market designs and **stress–test** them with learning agents.

CDA Intuition: Bundle Bids

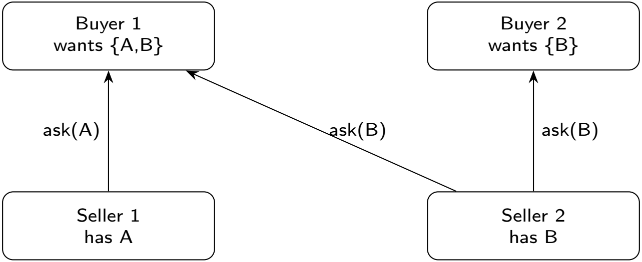

The market clears bundles to maximize surplus subject to feasibility.

Next Step

We now switch from this general CDA picture to a minimal laboratory: the BBDA.

Roadmap

Part 2: The Curse

When the Algos Run the Asylum

Double Auctions Are Everywhere

- **Financial & Commodities:** stock, bond, carbon markets.
- **Energy:** day-ahead/real-time electricity, smart grids.
- **Digital Platforms:** online ad auctions, ride-sharing.
- **Compute & Data:** cloud compute, emerging data/model markets.

The Problem

AI agents are already pervasive participants in many of these real markets.

The Core Function: What is Price Discovery?

Most Important Job of a Market

Aggregate and impound dispersed private information into a public signal: **the price**.

Why It Matters

- Decentralized identification of the rational-expectations value.
- Critical for financial, energy, and data/AI-model markets.
- **Informational efficiency:** how much information price reveals $\rightarrow$ Price Discovery
- **Allocational efficiency:** how optimally are goods allocated $\rightarrow$ Resource Allocation
- Tests how strategic behavior impacts efficiency.

Takeaway

We assess market health by *informational and allocational efficiency*. We test this in a simple double auction.

Our Laboratory: The Buyer's Bid Double Auction (BBDA)

Buyers' Bids & Sellers' Asks, Uniform Price

Setup for Experiments

We now test price discovery by replacing rational humans with RL agents inside this simple double auction.

- m buyers (want 1 unit) and n sellers (have 1 unit).
- Everyone simultaneously submits bids (buyers) and asks (sellers).
- Order all quotes: $s_{(1)} \leq \cdots \leq s_{(m+n)}$.
- Market clears at uniform price $p = s_{(m+1)}$.
- Buyers with bid $\geq p$ trade; sellers with ask $< p$ trade.

Why This Mechanism?

Foundational for studying allocational efficiency, **informational efficiency** (how private signals map into price), and strategic behavior.

The Model We Test: Correlated Private Values (CPV)

Strategic Bidding with Full Information

Model: Agents See Their True Value

This is a **Correlated Private Value (CPV)** model.

- Common Value: μ
- Private Value: ε_i
- **Agent Input:** The agent *knows* their true value: $z_i = \mu + \varepsilon_i$
- **Strategic Problem:** Even knowing their value, buyers must bid strategically ("shade" their bid) to get a better price
- Note that sellers bid truthfully *by design*

This is the Model We Test

Our experiment (next slides) uses this CPV setup to test if AI agents can find the optimal strategic "offset" from their true value z_i.

Baseline Theory (Humans)

In this simple double auction, with rational agents:

- Practical and robust vs. "optimal" mechanisms (Wilson critique, ex-post losses, budget balance).
- Allocational efficiency is high; losses vanish quickly as markets grow.
- Strategic effects are *small* relative to sampling noise.

Question

What changes when we replace humans with AI agents?

The Curse 1: Failure to Generalize

- Theory: buyer offset halves as market size doubles; seller offset ≈ 0.
- RL matches theory *in each fixed market*, but treats each size as a new task.

Market	Buyer Offset (Theory)	Seller
2×2	-0.6896	0
4×4	-0.3398	0
8×8	-0.1639	0
16×16	-0.0805	0

Rule of Thumb

Buyer's optimal offset scales with market size: doubling the market halves the offset. RL fits each size separately and fails to infer this scaling law.

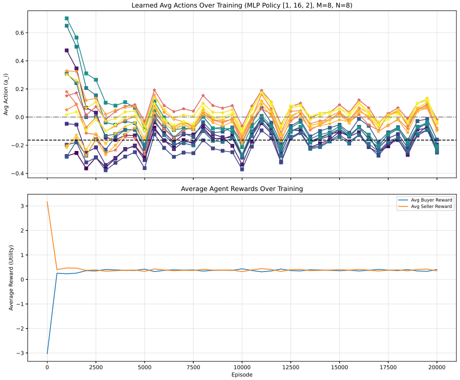

Example: 8×8 buyers/sellers

Lesson

A powerful solver is not necessarily a good extrapolator.

The Curse 2: Need for Informed Priors

Why Economic Structure Still Matters

Economic Insight

Our theory shows that optimal bidding is *independent* of an agent's private value z_i. Embedding this as an **architectural prior** (so the policy ignores z_i) greatly improves learning stability.

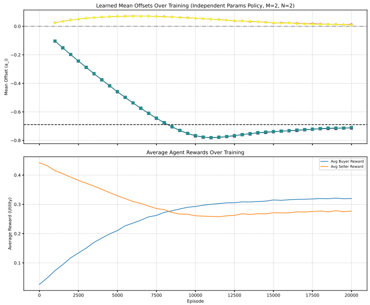

Informed AI (stable convergence): policy ignores z_i and learns a constant offset.

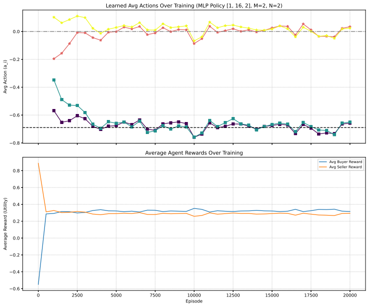

Uninformed AI: policy sees irrelevant z_i and must learn to ignore it (slower, less stable).

Lesson

Raw AI is brittle without **economic priors** and careful mechanism design.

The Curse 3: Sensitivity to Tuning

- Outcomes hinge on exploration vs. exploitation schedules (designer choices).

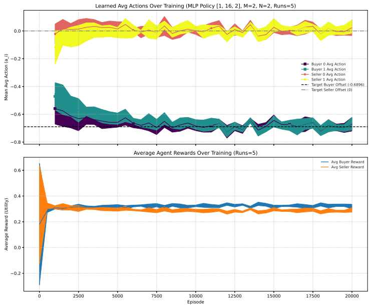

Less exploration (converges).

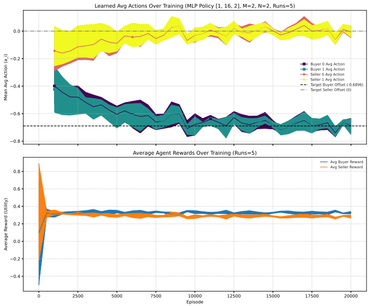

More exploration (slower, less stable).

The Verdict: Blessing or Curse?

The Blessing: AI as a Powerful Tool

- Solves complex pricing and portfolio problems.
- Enables privacy-preserving collaboration (FL).
- Shines when paired with **realistic** economic models & mechanisms.

The Curse: AI as a Brittle Agent

- **Brittleness:** fails to extrapolate simple patterns.
- **Inefficiency:** needs *informed priors* to learn efficiently.
- **Sensitivity:** heavily tuned by design choices.

Take-Home

The future isn't *less* AI in markets – it's **AI guided by economic theory and robust market design**.

Engineering Takeaways
What to Build, Test, and Watch Out For

Design Principles

- Markets are systems: mechanism → agents → price (feedback).
- Add economic priors to agents; do not start from scratch.
- Measure informational and allocational efficiency, not only performance.
- Stability over short-run gains; tune exploration schedules.

Where to Extend

- BBDA → multi-asset CDA with bundles.
- MARL coordination and equilibria under market rules.
- Theory shapes how AI learns—as an inductive bias—and helps check when it fails—as a debugging tool).

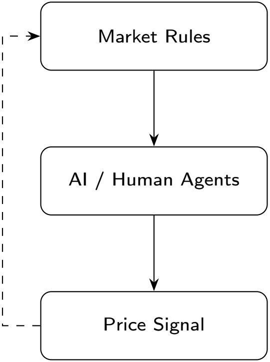

Engineering Testbench

Tune the market–AI feedback loop; **Measure** informational and allocational efficiency.

Thank You

Connect: **LinkedIn** (QR code)