Discussion of “AI Copilots in Real Estate: Evidence from China”

Lin William Cong, Xiang Hui, and Yushan Zhou

Yu Zhang
Guanghua School of Management, Peking University

ABFER Annual Conference 2026

What I take away from the paper

A very nice setting: an intermediary-facing AI copilot introduced in a large, two-sided resale housing market.

Two empirical designs:

- A neighborhood-level field experiment in Jinan.
- An 18-city exposure DiD exploiting pre-rollout app usage.

Main empirical message:

- Listings sell faster; buyer and seller time outcomes improve.
- Prices rise; ratings improve.
- Speed gains are two-sided; price effects load mainly on buyer-side exposure.

My reading

This paper is not only about AI productivity. It is a shock to the *production function of real-estate intermediation*.

Main suggestion and road map

My main suggestion: sharpen the economic interpretation of what the copilot changes.

$$\underbrace{\text{AI copilot}}_{\text{technology shock}} \quad \Rightarrow \quad \underbrace{\text{agent production function}}_{\text{search, matching, pricing, communication}} \quad \Rightarrow \quad \underbrace{\text{market outcomes}}_{\text{price, speed, ratings}}$$

1. **Start from the paper's** 2×2 **framework:** re-label outcomes by economic margins.
2. **Use a random-matching model:** interpret price and time-on-market together.
3. **Then connect to finance papers on intermediaries:** what agent functions can AI scale?

Comment 1: What exactly is the 2×2 matrix?

The paper's treatment-side labels are very useful:

$$\underbrace{E^B}_{\text{buyer-side AI exposure}} = \underbrace{\text{recommendations} + \text{buyer communication}}_{\text{search, screening, match quality}}$$

$$\underbrace{E^S}_{\text{seller-side AI exposure}} = \underbrace{\text{price-adjustment advice} + \text{listing text}}_{\text{price discovery, presentation}}$$

But the outcome-side labels are less sharp. Price, sale probability, and ratings are not purely buyer-side or seller-side objects.

Suggested refinement

Re-label outcomes by the economic margin they inform:

$$\underbrace{Y^{search}}_{\text{buyer search / screening}}, \quad \underbrace{Y^{pricing}}_{\text{seller price discovery / listing strategy}}, \quad \underbrace{Y^{joint}}_{\text{match surplus, clearing, incidence}}.$$

A more model-based 2×2 interpretation

	Buyer-side AI E^B	**Seller-side AI** E^S
Buyer search / screening Y^{search}	Own-side: fewer wasted viewings, shorter buyer search, better recommendation quality	Cross-side: better priced/described listings make inventory easier to evaluate
Seller price discovery $Y^{pricing}$	Cross-side: better buyer demand signals may discipline seller pricing	Own-side: more frequent, smaller price revisions; faster convergence to market-clearing prices
Joint outcomes Y^{joint}	Price can rise if buyer–house match quality improves	Sale probability and speed can rise; price effect is ambiguous ex ante

Why this helps

It avoids treating price mechanically as a "seller-side" outcome. In a search-and-bargaining model, price is a joint surplus-and-incidence object.

Comment 2: One random-matching model for price and speed

In the spirit of Ngai and Tenreyro (AER 2014), suppose buyers randomly arrive to a listed house.

$$\underbrace{v_{ij}}_{\text{buyer } i\text{'s value for house } j} = \underbrace{H_j}_{\text{common house value}} + \underbrace{\kappa\varepsilon_{ij}}_{\text{buyer–house match value}}$$

$$\underbrace{\text{Trade}}_{\text{sale occurs}} \iff \underbrace{v_{ij}}_{\text{buyer valuation}} \geq \underbrace{R_j}_{\text{seller reservation value}}$$

$$\underbrace{\mu_j}_{\text{sale hazard}} = \underbrace{\lambda_j}_{\text{arrival rate of relevant buyers}} \cdot \underbrace{\Pr(v_{ij} \geq R_j)}_{\text{probability an arrival clears reservation}}, \qquad \underbrace{TOM_j}_{\text{expected time on market}} = \frac{1}{\mu_j}$$

$$\underbrace{P_j}_{\text{expected transaction price}} = \underbrace{R_j}_{\text{seller reservation}} + \underbrace{\theta}_{\text{seller surplus share}} \underbrace{E[v_{ij} - R_j \mid v_{ij} \geq R_j]}_{\text{accepted-match surplus}}$$

Punchline

A better agent production function can raise prices and reduce time on market in the same model: it brings more relevant buyers and improves the quality of accepted matches.

How the copilot maps to model primitives

AI module	Model primitive	Predicted outcome
Buyer recommendations / communication	$\lambda_j \uparrow$, or distribution of ε_{ij} shifts right	$P \uparrow$, $TOM \downarrow$
Seller price-adjustment advice	$R_j \downarrow$ or faster convergence of the pricing threshold	$TOM \downarrow$, P ambiguous
Listing title / description	$\lambda_j \uparrow$ or better buyer screening	$TOM \downarrow$; $P \uparrow$ if match quality rises

Why this helps interpret the results

Buyer-side exposure should be the most natural source of price increases; seller-side exposure can mainly improve clearing and price discovery without necessarily raising prices.

The results can be read through this random-matching lens

Empirical pattern

Buyer-side AI raises prices and reduces search frictions; seller-side AI mainly improves price-discovery and clearing margins. The price and speed effects should be interpreted together.

Model object	**Empirical moment in the paper**
$E[v_{ij} - R_j \mid trade]$	Price effect loads mainly on buyer-side exposure
$\lambda_j \cdot \Pr(v_{ij} \geq R_j)$	Buyer and seller time-on-market fall; sale probability rises
Buyer screening efficiency	Buyer time falls; on-site viewings per transaction fall
Seller threshold / price discovery R_j	Seller-side exposure improves price-adjustment workflow; seller tools can recommend price reductions
Where frictions bind most	Larger gains for less-experienced agents and thinner markets

Additional moments to sharpen what the agent production function does

Moment	What it helps pin down
Closing probability per viewing	Search efficiency / match efficiency in the buyer-recommendation workflow.
Sale-to-list ratio and list-price path	Separates seller price discovery from final transaction-price effects.
Heterogeneity by hard-to-match houses/buyers	Match-quality channel should be strongest where baseline matching is hardest.

Emphasis

These moments would help readers map the reduced-form results into the agent production function: search, matching, and price discovery.

Framing comment: which frictions can this AI directly touch? (un?)codifiable?

The paper motivates the setting partly through soft information:

> *"resale housing remains a high-friction market because of soft, or uncodifiable, information about property attributes...."*

My reading of the modules:

Module	**Most direct friction affected**
Property recommendation	codifiable matching / shortlist construction
Reply scripts	communication and trust-building
Price-adjustment advice	price discovery using observables and comparables
Listing title/description	presentation of existing information

Suggested nuance

The copilot seems to directly improve codifiable matching, communication, and price discovery. It may help with soft local information indirectly, by freeing agent attention or helping agents deploy tacit knowledge more effectively.

Empirical comment: sharpening the "leveling" interpretation

The leveling result is important: effects are larger for less-experienced agents, lower-performing agents, and thinner markets.

A few institutional details could make this interpretation sharper:

- **Denominator:** junior agents may have lower baseline performance, so relative gains may look larger even for similar absolute gains.
- **Team structure:** Lianjia agents often work in teams; the relevant knowledge stock may be team-level or team-leader experience.
- **Within-team selection:** high-value clients may be routed to senior agents, while junior agents handle different clients.

Useful additions: show both absolute and relative effects; add team fixed effects or team-leader tenure; exploit within-team variation in agent experience.

Punchline: this would make the "AI substitutes for missing human experience" interpretation more transparent.

Empirical comment: Hot markets in 2019?

The exposure DiD uses pre-period app penetration. A natural reader question is whether high-app districts were also places with stronger local housing demand in 2019.

The paper already has neighborhood fixed effects and city-by-month fixed effects, and controls for income, GDP, unemployment, and population. One possible addition would be to absorb or test differential trends by:

- district age structure or younger-population share;
- pre-period price momentum;
- pre-period transaction momentum;
- pre-period online adoption growth.

Why this matters: the paper notes that app usage is related to younger and more digitally engaged users; those local markets may also be especially dynamic in 2019.

Comment 3: Connect to finance papers on intermediaries

Now the finance-intermediary literature fits naturally: it tells us what parts of the *agent production function* has been proved important and may be reshaped by AI.

- **Aiello–Garmaise–Nadauld (RFS 2026):** intermediary attention affects sale probability, speed, and price. ordering effects. neighboring listing loses.
  AI paper: intermediary technology improvement across the platform – so the interpretation is less about redistribution across listings and more about raising intermediary productivity.
- **Agarwal–He–Sing–Song (JFE 2019):** agents have information advantages that can affect prices and bargaining.
  AI paper: AI may compress human information advantages, but may also change surplus division.
- **Agarwal–Choi–He–Sing (RFS 2019):** agents and networks shape matching.
  AI paper: recommendation AI may codify and scale the matching skill of good agents.

Contribution statement

Previous finance papers show that agents matter. This paper asks what happens when AI changes what agents can do.

A suggested re-framing of the contribution

Current framing: AI improves agent productivity in a two-sided market.

Stronger framing:

Version I would like

This paper studies an intermediary-facing AI tool in a high-friction asset market. The tool changes the production function of human brokers. The results show that AI can improve market clearing, and the side-specific evidence suggests that buyer-facing AI operates through search and match quality while seller-facing AI operates through price discovery and listing execution.

This gives the paper three audiences:

- AI at work: productivity and leveling.
- Housing search: random matching, match quality, and time-on-market.
- Finance intermediation: what agents do, and how technology changes surplus creation versus surplus division.

To conclude

- I like the paper a lot: great setting, rich data, clean institutional variation.
- My main suggestion is to make the economics of the 2×2 matrix more model-based.
- Buyer-side AI looks like a search/match-quality technology.
- Seller-side AI looks like a price-discovery/listing-execution technology.
- The finance-intermediary literature can help the paper speak beyond AI productivity.

Looking forward to citing Cong, Hui, and Zhou!