AI Copilots in Real Estate: Evidence from China *

Lin William Cong
Cornell University,
ABFER, & NBER

Xiang Hui
Washington University
in St. Louis

Yushan Zhou
Dalian University of
Technology & DEFT Lab

December 5, 2025

Abstract

We study an agent-facing AI copilot introduced on a nationwide Chinese real estate platform. Because the copilot is accessible only via the firm's app, we identify its effects using a difference-in-differences design that interacts the rollout with pre-treatment district app penetration, and we corroborate the results with a field experiment that randomizes AI access across neighborhoods. The copilot increases the probability a listing sells within three months, shortens time-on-market for both sellers and buyers, and raises buyer and seller ratings, even as prices rise. To probe mechanisms, we proxy exposure to the copilot's buyer- vs. seller-facing workflows using pre-period district shares of buyers and sellers transacting via the app. This decomposition suggests strong cross-side effects——exposure on one side speeds outcomes on the other——and that the price increase loads on buyer-side exposure. Gains are larger for less-experienced agents and in thinner markets. The combination of higher buyer trust, faster transaction speed, and buyer-side-driven price increases is consistent with improved match quality and lower search frictions, rather than pure surplus extraction. Overall, our findings suggest that agent-facing AI improves matching efficiency in two-sided markets.

*

1 Introduction

The advancement of artificial intelligence (AI) technology has driven its rapid adoption across sectors (Maslej et al., 2025). While empirical evidence on AI's effect is growing, its implications in two-sided markets like real estate remain underexplored. In theory, AI can lower search costs and improve matching, raising total surplus. However, by equipping one side with better data and recommendations, it could also redistribute surplus in ways that discourage participation on the other side, offsetting efficiency gains. Whether AI improves market efficiency, and does so without harming the other side, is therefore an open question.

In this paper, we provide empirical evidence on these questions by studying the introduction of an AI copilot for real estate agents on China's largest real estate platform. Embedded directly in the agent-client chat workflow, the copilot advises agents on four tasks: for buyers, it recommends suitable properties and reply scripts; for sellers, it suggests price adjustments and property titles and descriptions.

We use a proprietary dataset from 18 cities for roughly one year before and after the copilot's introduction. The data contains a weekly panel of over 13 million listings, tracking on-site viewings and price adjustments. This sample includes about 230,000 completed transactions, for which we observe final prices and time on market for buyers and sellers. A subset of 90,000 transactions are rated, and for these transactions we observe overall ratings, review text, and structured feedback labels. Lastly, we have access to pre-rollout app penetration data at the district level (analogous to U.S. counties). The richness of this data allows us to measure market outcomes and study the mechanisms on both the buyer and seller sides.

Our primary design is a difference-in-differences (DiD) that exploits cross-district variation in pre-existing app penetration. This variation determines effective exposure to the treatment because the copilot is accessible only through the firm's app. Our specification includes neighborhood fixed effects (more granular than a district) and city-by-month fixed effects. We discuss the identification assumptions, including parallel trends, exclusion restric-

tion, potential concurrent changes, spillover, and no multiple versions of treatment in Section 3.3. To validate our findings and relax the DiD's identifying assumptions, we complement this analysis with a field experiment in Jinan, which randomized AI access across neighborhoods. This experiment provides a direct causal estimate, helping to rule out confounding factors such as potential concurrent platform changes delivered through the app.

We find that the AI copilot increases the probability a listing sells within three months, reduces time on market for both buyers and sellers, and raises post-transaction ratings from both sides, even as final prices increase. Replacing the overall penetration ratio with pre-period buyer- and seller-side app shares, the price effect loads on the buyer-side share, while both shares are associated with shorter time-on-market on the opposite side, consistent with cross-side effects. Event-study estimates, which are run for all main outcome variables, and for both the overall ratio and separately for the buyer- and seller-side ratios, show parallel pre-trends and immediate, persistent post-rollout effects, which strengthens a causal interpretation of these results. Field experiment estimates on the overall effects align in sign and magnitude.

We then explore the specific mechanisms driving these outcomes. On the demand side, buyers undertake fewer on-site viewings and report more "accurate recommendations", better communication, and higher trust, all consistent with lower search costs and improved matching. On the supply side, sellers adopt more agile pricing (more frequent, smaller adjustments) and report more "reasonable price adjustment" and better articulation of selling points. These same-side improvements produce cross-side benefits: tools aimed at one side shorten time-on-market and raise satisfaction on the other side.

What happens to surpluses? While sellers clearly benefit from higher prices and faster sales, the effect on consumer surplus is ambiguous. We do not attempt to estimate the full surplus split. Instead, we ask whether buyers appear worse off. Our evidence suggests that they do not: First, buyer ratings and trust improve. Second, the combination of higher prices and faster sales is consistent with an outward shift in demand. Third, the price effect loads

on buyer-facing tools rather than seller-facing tools that would plausibly raise bargaining power. Lastly, buyers' search process seems to become more efficient with fewer on-site viewings. Taken together, the evidence is most consistent with AI creating higher match value and lower search frictions, not pure surplus extraction.

Finally, we show that the effects are larger for less-experienced agents and in thinner markets, suggesting the copilot acts as a "leveler" for intermediaries and is most valuable where frictions are highest. These patterns help reconcile the net efficiency gains with the possibility of surplus redistribution: to the extent any redistribution happens, it does not seem to deter the other side enough to offset the matching improvements we document.

Overall, our results strongly suggest that the copilot raises market efficiency: sale probability rises, time-on-market falls, and participants report higher satisfaction. These gains arise through both same-side and cross-side mechanisms and are stronger where intermediary frictions are higher (less-experienced agents, thinner markets). While prices increase, the evidence points to higher match quality and reduced search costs rather than pure surplus redistribution.

The remainder of the paper proceeds as follows. Section 2 describes the setting. Section 3 describes data and empirical strategy. Section 4 presents the main results. Section 5 unpacks the underlying mechanisms. Section 6 reports the results from our field experiment. Section 7 concludes.

1.1 Related Literature

Our paper is closest to work showing that AI chatbot improves service performance and "levels up" lower-skilled workers. In customer support and sales settings, AI guidance raises resolution rates and customer satisfaction, with the largest gains for underperforming agents (Luo et al., 2019; Schanke et al., 2021; Brynjolfsson et al., 2025). Our work contributes by studying an agent-facing AI copilot in a high-stakes, two-sided market, where the copilot helps real-estate agents' search, communication, and pricing advice, and we trace both same-

side and cross-side consequences in equilibrium outcomes. To our knowledge, this specific use of an AI assistant has not been studied, primarily due to data limitations.

A second stream studies consumer-facing algorithmic price signals such as Zillow's Zestimate. These papers document how algorithmic prices affect listing and selling outcomes and can generate feedback loops between human decisions and algorithmic recommendations (Fu et al., 2022; Malik and Manzoor, 2023), and can mitigate certain disparities (Yu, 2020; Fu et al., 2025). Our work differs in both mechanism and outcomes: we study an *agent-facing* copilot embedded in workflow, rather than a consumer-facing price signals, and we move beyond prices and time-on-market to analyze cross-side effects and perceived service quality.

More broadly, our study connects to the literature on AI-driven recommendation and dynamic pricing tools in other contexts, including ridesharing platforms (Hall et al., 2015; Liu et al., 2023; Castillo, 2023), online rental markets (Zhang et al., 2021b; Pan and Wang, 2021; Huang, 2024), and e-commerce platforms (Zhang et al., 2021a; Fong et al., 2023; Shi et al., 2024). Our paper also contributes to the growing literature documenting AI's rapid diffusion and its economic effects across a wide range of domains, including writing (Noy and Zhang, 2023; Wiles et al., 2025; Hui et al., 2024; Demirci et al., 2025), translation (Brynjolfsson et al., 2019; Zhu and Walker, 2024), coding (Cowgill et al., 2020; Peng et al., 2023), and professional services (Zhou and Lee, 2024; Peukert et al., 2024; Goldberg and Lam, 2025; Dell'Acqua et al., 2023). Similarly, AI is increasingly used in high-stakes decisions in fields such as risk assessment (Cong et al., 2024; Bertomeu et al., 2025), customer engagement (Luo et al., 2019; Schanke et al., 2021; Brynjolfsson et al., 2025), education (Mollick and Mollick, 2023), and healthcare (Agrawal et al., 2024). These findings inform the broader economics of AI adoption (Brynjolfsson et al., 2018; Agrawal et al., 2019). Our contribution is to show how an intermediary-facing AI tool affects matching efficiency and welfare-relevant outcomes in a large real estate marketplace.

2 Setting

2.1 China's Real Estate Market

Scale: The Chinese real estate market is among the largest globally, reaching total commercial housing sales of 17.36 trillion RMB in 2020.[1] Real estate is the primary form of household wealth in China, with an urban homeownership rate of 96% and residential housing assets accounting for about 69% of total household assets.[2]

Data Infrastructure: Unlike the U.S., where the Multiple Listing Service (MLS) provides industry-wide data sharing, Chinese brokerages typically own and guard their listing data. This fragmentation limits transparency and potentially inefficient matching. Consumers consequently rely heavily on in-person intermediation by agents. Additionally, this data fragmentation reinforces a dense network of physical real estate stores, as consumers must rely heavily on in-person interactions with agents due to incomplete online information.

Governance Structure: Large firms in China directly employ agents (who earn salary plus commissions) and often intermediate both sides of a deal (dual agency). This firm-centric model encourages keeping buyers and sellers within the same organization and gives the firm control over tools, training, and workflows. This is different from the U.S., where licensed brokers supervise agents who are independent contractors.

Platforms: Starting in the late 2010s, leading firms built integrated online–offline platforms. Unlike U.S. platforms (e.g., Zillow) that only aggregate listings and provide user-friendly interface, these Chinese platforms control both data and transactions, combining MLS-style listing management, consumer-facing, Zillow-like search, and brokerage services in one place. This vertically integrated model has allowed dominant platforms to leverage network effects, attract more buyers and sellers, and expand through acquisitions of smaller firms.

[1] Source: https://www.stats.gov.cn/sj/zxfb/202302/t20230203_1900967.html

[2] Source: https://m.163.com/dy/article/FB0AF3P20539ADBX.html; https://stock.hexun.com/2023-08-24/209822345.html

2.2 Empirical Context

We study one of the largest nationwide platforms. By 2019 it operated thousands of offline stores across more than one hundred cities and employed on the order of one hundred thousand agents, facilitating millions of annual transactions. Sellers and small agencies list on this platform to access its verified, up-to-date inventory and buyer traffic. Buyers use it for trusted, up-to-date listings and professional intermediation. As user base grows, the platform have better data, leading to better matches, which in turn would attract even more users—a reinforcing loop that increases on-platform share of transactions.

Transactions happen through both offline (walk-ins to a store, phone) and online (app/web) entry points. About 40% remain offline end-to-end, and the remaining on the firm's app launched in 2015. The app is a central workflow when clients engage via the app, agents message, recommend homes, manage listings, and schedule viewings there. The AI copilot is available *only within the app's agent-client chat*.

There are clear service boundaries on this platform. Agents serve clients within a local neighborhood (platform-defined commercial or residential areas). Districts are strictly nested within cities, but neighborhoods can straddle district borders. Company policy restricts agents from operating outside their local area: if a client wishes to view a property in another district, the case is handed off to the destination district's agent. This limits cross-district spillovers, which we detail in Section 3.3.

The typical transaction process on the platform is summarized in Figure A1. Buyers initiate in-person or via the app. The agent elicits preferences, recommends listings, schedules on-site viewings, and brokers negotiation. A seller lists through an agent in-person or via the app. The agent helps set an initial price, writes titles and descriptions, monitors interest, and recommends price adjustments. If a deal is reached, the platform records transaction date and price and both sides submit ratings and reviews and structured labels about the agent's service.

2.3 Policy Change: Introduction of AI Copilot for Agents

On January 16, 2019, the firm introduced an app-exclusive AI copilot to agents, which became popular quickly.[3] The AI copilot gives recommendations targeted four main agent challenges. For buyers, it recommends suitable properties and reply scripts. For sellers, it suggests price adjustments and property titles and descriptions. Importantly, AI recommendations are advisory and visible only to agents, who retain full discretion over adoption decisions without managerial pressure.

Figure A1 illustrates how the AI copilot integrates into transaction processes. For buyer-facing workflow, the copilot recommends to agents suitable listings based on conversation text and buyers' browsing histories; see Figure 1(a) for an example. Additionally, the AI suggests professional reply scripts to improve communication; see Figure 1(b) for an example. For seller-facing workflow, the copilot suggests dynamic price adjustments using comparable listings and market information (Figure 1(c), and proposes optimized titles and descriptions to highlight key selling points (Figure 1(d)).

Our setting offers several advantages for studying AI's effect on matching in a two-sided market. First, it offers a sharp and exogenous shock, i.e., the unanticipated introduction of the AI copilot. Second, the institutional feature that copilot is accessible only through the agent app allows us to proxy for exposure using pre-treatment district app penetration. Third, the rich data on detailed property characteristics and transaction outcomes, in particular, ratings and reviews, allows us to link mechanisms to outcomes while being able to control for fine property and geographic heterogeneities. We assess identification threats (e.g., spillovers, parallel trends) in Section 3.3 and benchmark magnitudes with a neighborhood-level field experiment in Section 6.

[3]By late 2019, about 60% of AI recommendations were adopted by agents.

3 Data and Empirical Framework

3.1 Data

We use a proprietary dataset from a nationwide real estate platform covering 18 major Chinese cities from late February 2018 to early January 2020. The AI copilot was introduced in 16 January 2019, roughly at the mid-point of our window. The data links property listings, weekly behaviors while a listing is active, transactions outcomes, and post-sale evaluations.

Listings and Behavioral Panels: For all listings that are active in our window, we observe a weekly panel on two key behaviors: on-site buyer viewings and seller price adjustments. We also observe detailed, time-invariant property characteristics (e.g., size, rooms, floor, age, decoration). These panels span about 13 million unique property listings during the sample period. At the listings level, our main conversion outcome is an indicator for sale within 3 months of listing.[4]

Transaction and Feedback: For listings that successfully sell, we observe a transaction record (closing date and price, last listed price, and time-on-market for buyers and sellers), detailed property characteristics, and agent attributes (gender, age, tenure, performance). The transaction sample contains roughly 230,000 closings; we drop around 20 observations for missing fields. For a subset of roughly one-third of these closings, we also observe post-transaction feedback from both buyers and sellers, comprising 5-point overall ratings, text reviews, and structured labels to evaluate specific aspects of agent's service (e.g., "accurate recommendation", "trustworthy", "reasonable price adjustment", "selling-point accuracy")[5].

Treatment Intensity and Covariates: Our treatment intensity is based on *pre-treatment* app penetration provided by the firm at the district level: the share of transactions con-

[4]We define the conversion outcome this way to handle the usual truncation issue as some listings enter before or exit after our window. We chose 3 months because 54% of properties that are eventually sold are sold within 3 months. We also use a four-month selling window as a robustness check and find that the results remain robust.

[5]These post-transaction feedback entries include 75,596 from buyers and 20,336 from sellers. The larger number of buyer feedback entries is due to the fact that, during the sample period, companies sent text message or phone reminders to buyers to complete their evaluations.

ducted through the app, and separately, the buyer- and seller-side shares. We use these as treatment-intensity measures (details in Section 3.2). We complement these with district-level demographics from the firm's internal sources (population, income, GDP, unemployment). Districts are strictly nested within cities (prefecture-level or direct-administered municipalities). The finest geographic unit in our regressions is the platform-defined "neighborhood", for which we will include fixed effects.

Table 1 presents the summary statistics for the variables used in the main outcomes. We use the natural logarithm of one plus days on market for sellers, days on market for buyers and transaction prices as our dependent variables. The mean of Log_days_seller is about 3.99 (or 55 days on market for sellers). The mean of Log_days_buyer is about 2.66, corresponding to an average time on market of about 15 days for buyers. The mean of Log_price is 10.27, corresponding to an average transaction price of 28,800 RMB per square meter, or approximately 4,122 USD per square meter. Appendix Table A1 also presents the summary statistics for the additional variables used in the demand-side and supply-side analyses.

3.2 Empirical Strategy

Our primary objective is to examine how the adoption of AI copilot for supporting agents affects outcomes in the housing market. In our context, the AI's functionality is deeply integrated with the firm's mobile app and depends on customer-agent interactions within that channel. That is, the AI's features are designed to be used within the app's agent-client chat interface. This distinction between app-based and offline clients forms the basis of our identification strategy. Critically for our research design, the proportion of app-based versus offline clients varies significantly across geographic regions.

Our identification strategy exploits the differential proportion of users (buyers and/or sellers) using the app for transactions across districts during the pre-treatment period. We

estimate the effect using a continuous diff-in-diffs regression:

$$\mathrm{y}_{i,n,d,c,t} = \beta_1 \mathrm{Post}_t + \beta_2 \mathrm{Ratio_tran}_d + \beta_3 \mathrm{Post}_t \times \mathrm{Ratio_tran}_d + X_{i,d,t}\delta \ + \xi_n + \tau_{ct} + \epsilon_{i,n,d,c,t} \quad (1)$$

Our dependent variables, $y_{i,n,d,c,t}$, are whether the listed property is sold within three months, the time on market for sellers, time on market for buyers, and housing transaction prices for transaction i in neighborhood n district d city c in year-month t. Post_t indicates whether the transaction occurred after the deployment of the AI copilot. $\mathrm{Ratio_tran}_d$ indicates the proportion of housing transactions at the district level that were conducted through the app during the pre-treatment period. The specification also includes neighborhood fixed effects ξ_n to account for time-invariant neighborhood-specific confounders, city-month fixed effects τ_{ct} to control for seasonal factors and the effect of city-specific shocks on the housing market (e.g., changes in real estate policies). $X_{i,d,c,t}$ represents two sets of control variables: (1) housing characteristics such as building size, number of rooms, etc., and (2) four key district-level variables—income, unemployment, GDP, population. In the robustness analyses in online appendix Table A4, we also control for the interaction of these district-level variables with Post in the analyses.

We use three geographic levels: city (prefecture-level or direct-administered), district (county-level units nested within cities; this is where treatment intensity varies), and platform-defined neighborhoods (the finest unit used for location fixed effects). Neighborhoods are smaller than districts but can occasionally straddle district borders[6]; we address potential cross-border spillovers in Section 3.3.

Since our treatment group is defined based on the proportion of housing transactions conducted through the app at the district level, we control for district-level characteristics ($X_{i,d,t}$) and include city-by-month fixed effects (τ_{ct}) to capture broader geographic characteristics and time trends. The location fixed effects (ξ_n) are defined at the finer, neighborhood

[6]For example, in Beijing, the Dongzhimen neighborhood spans across both Dongcheng District and Chaoyang District.

level to account for localized neighborhood-specific attributes that may affect housing market outcomes at a more granular level.

Our coefficient of interest is β_3, which measures the effect of different levels of access to the AI copilot on key outcomes in the housing market following the deployment of the AI copilot. In our main specifications, standard errors are clustered at the district level, which is conceptually the level of random assignment (Abadie et al., 2023). In an alternative specification, we replace Ratio_tran_d with Ratio_buyer_d and Ratio_seller_d, which represent the proportion of buyers and sellers using the app, both interacted with the *Post* dummy. Because the copilot's features are embedded in distinct chat workflows—agent–buyer vs. agent–seller, we use Ratio_buyer_d and Ratio_seller_d to proxy for exposure to the copilot's buyer-facing and seller-facing workflow.[7]

3.3 Discussion of Identification Assumptions

We discuss several threats to the validity of our identification strategies: differential secular trends, exclusion restriction violation, Stable Unit Treatment Value Assumption (spillovers and consistency) violations.

Differential Secular Trends: The concern is that high-penetration districts may be on different trends (digital adoption, income growth, agent quality), so Post × Ratio could pick up those trends rather than the AI copilot. This concern is fundamentally about the parallel trends assumption, which requires that districts with high and low app penetration have similar time trends in the outcome variables. While this assumption cannot be proved, we test for parallel pre-trends by estimating a flexible diff-in-diffs event-study specification:

$$\text{y}_{i,n,d,c,t} = \alpha\text{Period}_k + \gamma\text{Ratio_tran}_d + \sum_{k\neq -1} \beta_k \text{ Period }_k \times \text{Ratio_tran}_d + X_{i,d,t}\delta + \xi_n + \tau_{ct} + \epsilon_{i,n,d,c,t} \tag{2}$$

[7] We do not observe function-by-function tool usage. These pre-period shares are proxies for differential exposure because the copilot's buyer- (seller-) facing features live in the agent–buyer (agent–seller) chat flows.

We can inspect whether the parallel trends assumption holds by checking if β_k is non-zero when k is less than 0, which corresponds to the period before the AI copilot deployment.

Exclusion Restriction Violations: The exclusion restriction requires that pre-existing app penetration only affects housing outcomes through access to AI after the rollout. While testing for parallel pre-trends based on Equation 2 largely mitigate this concerns, one remaining concern is that high penetration in a district could correlate with latent traits (richer clientele, high-price segments, storefront density) in that district that also correlate with responsiveness to AI. For example, if app penetration is higher in district with richer clientele, and richer people are more responsive to technological changes, then a comparison in the outcome variables across richer vs. poorer district may not only reflect the different app penetration, but also differential responsiveness from different types of people, for a given level of app penetration. We address this in several ways:

First, we control for a rich set of fixed effects in Equation 1. Time-invariant neighborhood factors, e.g., the average wealth of households in a neighborhood, is controlled for by the neighborhood fixed effects. Additionally, differential time trends across different cities, e.g., differential demographic changes by city, are controlled for by the city-by-month fixed effects. Second, we include rich property-level covariates $X_{i,n,d,t}$ (e.g., building age, rooms, etc), and time-varying district-level covariates (income, unemployment, GDP, population).

Third, we investigate the nature of this geographic variation of app penetration. Although smartphone penetration in China's first-tier cities is nearly saturated, the share of second-hand housing transactions conducted through the app varies across regions.[8] In an interview with a senior executive at the company, we learned that the proportion of consumers using the app to contact agents ranges from about 25% to 70%, depending on three factors. First, younger buyers and sellers are more likely to communicate with agents via the app rather than visiting physical stores. Second, for high-priced properties, both buyers and sellers tend to be more cautious and prefer in-person meetings with agents over app-based interactions.

[8] As of June 2019, 99.1% of Chinese internet users accessed the internet via smartphones.

Third, in areas where offline stores are sparse or less visible, a higher proportion of consumers rely on the app to connect with agents. We control for the potential varying elasticity based on groups with different latent selection traits, including the three factors listed above. In our robustness checks (Appendix Table A4), we additionally control for interactions between the Post dummy and property-level and district-level controls (e.g., building age, number of rooms, unemployment, population, income, and GDP). The results show that our main treatment effect estimates (on Post×Ratio) remain stable and statistically significant with the inclusion of the additional interaction terms, which alleviates this concerns of differential elasticity along selection variables.

Concurrent Changes: Another concern is the possibility of other concurrent changes on the platform that could confound our results. For example, if the firm rolled out a new user interface or a separate marketing campaign at the same time as the AI copilot, these changes could also affect market outcomes. Note that our empirical design already mitigates this threat by controlling for city-by-month fixed effects. In addition, we run a series of placebo tests on outcomes that the AI should not effect, as will be discussed in Section 4.1, and find null results there. Lastly, the neighborhood field experiment (Section 6) that randomized copilot availability gives qualitatively same effects, suggesting that our main results capture the AI copilot effect, rather than other feature changes.

SUTVA - Spillover Component: Lastly, we need to assume no spillovers from treated to untreated units. The primary concern is whether an agent, after being equipped with the AI copilot, could indirectly benefit their offline clients. We consider four potential channels for such a spillover. First, an agent could attempt to directly use the AI for an offline client. We view this as impractical because the AI copilot is not a standalone tool. Its core functions are deeply integrated into the app's chat interface between the agent and a specific online client, relying on that client's data to generate real-time recommendations.

Second, an agent could try to convert an offline client into an app user to enable the AI's functions. This was prevented by a strict institutional policy. During our sample period, the

company explicitly prohibited agents from requesting that customers download or switch to the app for communication. Third, an agent could learn general principles from their AI-assisted interactions over time and apply this knowledge to their offline clients. While this form of learning is possible, the event study graphs (e.g., in Figures 2 and 3) show an immediate and sharp effect in the first month following the AI's introduction, suggesting that learning over time is not the primary driver of our results.

A final concern is potential spillover from neighborhoods that straddle district boundaries. We address this threat both institutionally and empirically. Institutionally, company policy restricts agents to their local, store-based area. If a buyer wants to view properties in another district, the case must be formally transferred to a local agent, which limits direct cross-district interference. Empirically, we address neighborhood-level spillovers by excluding transactions in neighborhoods that straddle district borders and re-estimating our main specification. The resulting coefficients are statistically indistinguishable from, and if anything, slightly larger than, our baseline estimates (Appendix Table A2). The stability of our results suggests that this type of spillover is not driving our findings.

Therefore, while we acknowledge these possibilities, the institutional context and the pattern of our results suggest that any such spillovers are likely limited in scope. To the extent that some rapid learning does occur and benefit the control group, our results should be interpreted as a conservative lower bound of the AI's true impact on market efficiency.

SUTVA - Consistency Component: The consistency component of the SUTVA requires that there is no different versions or forms of treatment, as the treatment intensity $Ratio_d$ is a continuous variable between 0 and 1. Following the recommendation in (Callaway et al., 2024), we replace the continuous *Ratio* variable with an indicator variable for treatment to verify the consistency of our results under a simplified, binned treatment definition. Binning collapses nearby treatment doses into a few bins that are more plausibly homogeneous versions of treatment and reduces reliance on functional-form assumptions and measurement error in $Ratio_d$. The results are reported in Appendix Table A7.

4 Main Results

In this section, we present the main findings of our study. We first show the overall effect of the AI copilot on key market outcomes: whether the listed property is sold within three months, time on market for sellers and buyers, transaction prices, and buyer and seller satisfaction ratings. We then decompose this effect to understand how the AI's distinct buyer-facing and seller-facing functions contribute to these changes and generate cross-side effects.

Before going into our main estimates, we briefly summarize the validation checks that support their causal interpretation. The event study plots confirm parallel pre-trends, a series of placebo and composition tests are null, and the results are robust to excluding potentially contaminated "straddling" neighborhoods. Importantly, these quasi-experimental findings on overall effects are strongly corroborated by our randomized controlled trial (Section 6), providing strong support for a causal interpretation.

4.1 Overall Effect

Table 2, Panel A, presents the key parameter estimates using Equation 1 (see coefficient estimates for all control variables in Appendix Table A3). Overall, results show that access to AI copilots significantly increases the likelihood that a listed property is sold within three months, decreases the time on market for both sellers and buyers, increases housing transaction prices, and increase satisfaction ratings from both buyers and sellers.

More specifically, in column 1, we observe that for a district at mean pre-treatment app penetration (0.446), the likelihood that a listed property is sold within the first three months after it is listed increases by 18.8 percentage points (0.4215*0.446) compared to the control group (agents without access to AI). Column 2 shows that for a district at mean pre-treatment app penetration, the time on market for sellers is lower by 0.198 (–0.4434*0.446) compared to the control group. Relative to a baseline pre-treatment mean of 3.757, this

represents a 5.27% decrease. In column 3, we see that access to AI copilots reduces the time on market for buyers by 0.111 (-0.2484*0.446), or by 4.04% relative to the baseline mean of 2.747. Column 4 shows that access to AI copilots also has a positive and statistically significant effect on unit transaction prices. The estimated coefficient indicates an increase in transaction prices by 0.045 (0.10×0.446), relative to the control group, a 0.43% increase from the pre-treatment mean of 10.535.[9]

Although prices rise, both buyer and seller post-transaction ratings increase. Panel A, columns (5) and (6) shows that at mean pre-period app penetration, buyer ratings rise by 9.84% (0.2206 × 0.446), and ratings from sellers by 15.49% (0.3474 × 0.446). Together with shorter time-on-market, this is more consistent with improved matching and lower search frictions than with pure surplus reallocation. We elaborate on surplus in Section 4.2.

We conducted a few validation and robustness checks. Event–study estimates based on Equation 2 show flat pre-trends and immediate, persistent post-introduction effects for the treatment ratio (Figure 2). This suggests a causal relationship between AI adoption of the outcome variables.

We also investigate whether composition changes. If effects reflected shifts in the property/agent/client mix or unrelated platform shocks, we should see movements in listing characteristics and in outcomes the copilot does not target. We do not see these in the data. First, Appendix Table A10 shows no significant change in the characteristics of properties (rooms, size, age, floor, decoration).[10] Second, we use a series of placebo tests on outcomes the AI should not affect. Appendix Table A8, columns (3) and (4) show that the AI has virtually no effect on sellers' perceptions of an agent's communication skills.

Lastly, the results are robust to adding interactions of district covariates with *Post* (Appendix Table A4), alternative clustering at the district–week level and the city level (Ap-

[9]Given that the average unit price of housing in our sample is 47,351 RMB and the average building size is 84.28 square meters, the adoption of AI copilot leads to an average increased in price by about 17,160 RMB, or about 2,400 USD (0.43% * 47,351 RMB * 84.28).

[10]The only exception is column 4, where the outcome variable is the level of decoration of the property. While the estimated coefficient on Post×Ratio_tran is negative, the magnitude is small and only marginally significant.

pendix Tables A5 and A6), and a binary treatment in place of the continuous ratio (Appendix Table A7).

4.2 A Two-Sided Story

To further understand how the AI copilot functions, we replace the single treatment measure with two pre-period proxies for exposure to the copilot's workflows: the district share of buyers transacting via the app (*Ratio_buyer*) and, separately, the district share of sellers transacting via the app (*Ratio_seller*.[11]) Re-estimating Equation 1 with both interactions yields three findings (Table 2, Panel B).

First, efficiency gains are two-sided. Column 1 shows that both $Post \times Ratio_buyer$ and $Post \times Ratio_seller$ increase the probability a listing sells within three months, by 6.77% (0.1744*0.388) and 10.91% (0.4117*0.265), respectively. They also reduce time-on-market for both buyers and sellers, as shown in columns (2) and (3). This shows a strong spillover effect on market efficiency, as tools designed for one side translate into faster transactions for the other side.

Second, price effects are driven by the buyer side. As shown in column 4, the increase in transaction prices is loaded on $Post \times Ratio_buyer$; the $Post \times Ratio_seller$ coefficient is near zero and not statistically significant. This pattern is consistent with buyer-facing features (recommendations, reply scripts) improving match quality and willingness-to-pay, rather than seller-facing features extracting surplus through bargaining.

Third, satisfaction rises on both sides, including cross-side. Columns (5) and (6) show that buyer-facing exposure increases seller ratings and seller-facing exposure increases buyer ratings, alongside own-side gains. This suggests that making the process work better for one side appears to make the entire transaction more satisfactory for the other.

As tests for the parallel trends assumption, event-study specifications for both *Ratio_buyer*

[11] We do not observe tool usage by function. These pre-period shares proxy exposure to the buyer-facing (agent-buyer chat) and seller-facing (agent-seller chat) workflows because the copilot is embedded in those app flows.

and *Ratio_seller* show flat pre-trends and immediate, persistent post-introduction change (Figure 3), supporting a causal interpretation of these two-sided and cross-side patterns.

5 Mechanisms

We study how the copilot changes the agent's workflow and, through the agent, buyer and seller behavior. Note that we do not observe function-level usage; instead, "exposure" to buyer- vs. seller-facing workflows is proxied by pre-period district shares of buyers and sellers transacting via the app (Section 3.1). We first study demand-side mechanisms, then supply-side mechanisms, and conclude with two-sided dynamics (cross-side effects, heterogeneity, brief feedback-loop evidence) and surplus interpretation.

5.1 Demand-Side Mechanisms

Buyer Feedback on Agent's Service: We find that exposure to the buyer-facing workflow is associated with higher perceived recommendation quality and better communication, as shown in Table 3. First, buyers are more likely to select the "high-quality recommendation" label: For a district with average app penetration, the likelihood increases by a remarkable 0.176 (0.3944×0.446), or 42% relative to the pre-treatment mean (column 1).

Second, the quality of agent-buyer communication improves substantially. This is evidenced both by a sharp decrease in buyers reporting "communication inefficiency" (column 2) and, based on our NLP analysis of review texts,[12] a significant increase in explicit positive evaluations of the agent's communication skills (column 3).

Third, buyer trust in agents also rises. Column 4 shows a significant increase in the likelihood that a buyer chooses the "trustworthy" label when evaluating agents. As shown

[12]Guided by prior research (Niu et al., 2025), we prompt GPT-3.5 with the following instruction: "You are a business researcher. Please identify whether the following review contains a positive evaluation of the agent's communication skills. If it does, assign an indicator value of 1; otherwise, assign 0. We also cross-checked a subset of the classifications manually and found a high level of consistency with the prompt-based labels.

in Panel B of Table 3, all of these positive effects on the buyer experience are driven by the agent's use of the AI's buyer-facing functions.

We conduct a series of placebo tests. We find that the AI has no effect on an agent's *on-site* presentation skills or perceived appearance and temperament—attributes that AI copilot is not designed to assist with. Furthermore, the AI's communication tools, which are exclusively designed for buyer-agent interactions, have no effect on ratings given by sellers. The consistently null results of these tests, which can be found in Appendix Table A8, increase our confidence that the effects we document are attributable to the AI's specific functions.

Consequences for Buyer Search. Consistent with more efficient matching, we find that on-site viewings per transaction decline when buyer-facing exposure is higher (Table 3, column 5). For a district with average buyer app penetration, AI's buyer-facing functions reduces on-site viewings by 0.138 (-0.3548×0.388), representing a 8.57% decrease from the pre-treatment mean of 1.6. Additionally, we find that the aggregate transaction volume is stable (Appendix Table A8), suggesting limited scope for mechanical composition effects.

Buyers also end up purchasing homes with higher last listed prices. Column 6 shows that average access to AI copilot increases Log_listed_price by 0.043 (0.0963×0.446), a 0.41% increase from the pre-treatment mean of 10.57. This is consistent with the overall price increase documented earlier and suggests that a more efficient search process helps buyers find properties that better fit their needs. The event studies graphs for these outcomes are reported in Figure A2 in the Appendix and are consistent with the validity of the parallel trends assumption.

5.2 Supply-Side Mechanisms

Seller Feedback on Agent's Service. We find that exposure to the seller-facing workflow raises the probability that sellers choose the "reasonable price adjustment recommendations" label. For a district with average app penetration, the likelihood increases by 0.11

(0.2465×0.446), representing a 36.7% increase relative to the pre-treatment mean, according to Table 4, column 1. Additionally, the exposure also raises the probability that sellers choose the "accurately capturing the property's key selling points" label. Column 2 shows that the AI increases this likelihood by 14.6%. Panel B of the table confirms that both of these improvements in perceived service quality are driven specifically by the agent's use of the AI's seller-facing functions. Event studies again support parallel pre-trends (Fig. A4).

Consequences for Seller Pricing. We find that pricing adjustments become more frequent, smaller, consistent with faster convergence to market-clearing prices, according to columns 3 and 4 in Table 4. The decomposition exercise in Panel B, column 3 shows that average access to the supply-side AI function increases the weekly frequency of price adjustments by 0.032 (0.0820 × 0.388), or a 20% increase from its pre-treatment mean of 0.16. Meanwhile, this reduced the average magnitude of those adjustments by 0.308 (-0.7933 × 0.388), a 7.3% reduction from its pre-treatment (column 4). The event study graphs in Figure A4 are consistent with the validity of the parallel trends assumption.

Additional checks show the frequency increase is driven by more downward adjustments (Appendix Table A9; event studies in Figure A3). Composition checks on listing characteristics are small and mostly insignificant (Appendix Table A10), and outcomes not targeted by the seller-facing tools (e.g., sellers' communication-skill labels) are null (Appendix Table A8). We therefore attribute the supply-side effects to the copilot's functions rather than composition changes.

5.3 Two-Sided Dynamics and Heterogeneity

The AI's effect is best understood through the lens of two-sided market dynamics. We find that the AI generates reciprocal benefits for efficiency and satisfaction, and that these effects are most pronounced for less-experienced agents and in thinner markets. We did not find strong evidence on feedback loop caused by the AI copilot.

Cross-Side Efficiency and Satisfaction. Decomposing exposure, buyer-facing tools

shorten sellers' time on market (Table 2, Panel B, column 2, and seller-facing tools shorten buyers' time (column 3). Cross-side improvements extend to ratings: seller-facing tools raise buyer ratings and buyer-facing tools raise seller ratings (columns 5 and 6). Seller-facing features that sharpen listing quality also make recommendations appear more "optimal" to buyers (Table 3, Panel B, column 1).

While the evidence points to cross-side effects, we do not find strong evidence on feedback loop. Exploratory triple interactions between DiD exposure and sellers' price adjustment frequency or magnitude provide limited evidence of direct buyer–seller behavioral loops. The only one specification (more rational downward price adjustments) shows a significant reduction in buyers' time on market (Appendix Table A15, column 4), consistent with the AI copilot helping sellers make more rational and effective price adjustments, thereby reducing days on market for buyers.

Heterogeneity by Agent Experience and Market Thickness. We find that the AI's ability to generate these positive cross-side effects is most pronounced for less-experienced agents, but the effect do not seem to differ by their gender or education. Specifically, we include a triple interaction term to study how agents' tenure (representing the number of years agents have been in the industry pre-treatment) mediates the effects in Table 5. The Tenure variable has been standardized. column 2 shows that the cross-side effects from the buyer-focused AI onto a seller's time on market is significantly larger for agents with less tenure with the platform. We find similar (but only marginally significant) cross-side pattern from seller-focused AI on buyer's time on market too in column 3. We see similar patterns for ratings (columns 5 and 6)

The findings are robust to proxying experience with agent performance, which is an internal score used by the platform to evaluate agents calculated during the pre-treatment period, as an alternative proxy for experience (Appendix Table A11). Lastly, we find no meaningful heterogeneity by agent gender and education (Appendix Tables A12 and A13).

Lastly, we find that buyer-facing exposure is especially valuable in thinner markets. We

defined market thickness as the number of listed properties in the neighborhood with the same number of rooms as the transacted property during the pre-treatment period, which is then normalized. Effects on sellers' time on market and on both sides' ratings are larger where pre-period local supply is thin (Table A14), consistent with AI alleviating matching frictions when high-quality matches are harder to find.

5.4 Surplus Implications

What do higher prices imply for surplus? While higher prices and lower time-to-sell should benefit sellers, the effect on buyers is not immediately obvious. On the one hand, holding matching quality the same, a higher transaction price implies a lower buyer surplus. On the other hand, the use of AI may increase matching quality by allowing buyers to make better decisions and lower search costs. While we cannot directly recover willingness-to-pay, we provides suggestive evidence that buyers are not made worse off.

First, and most directly, we find a significant increase in buyers' stated satisfaction and trust following the AI's introduction (Table 3). It is unlikely that trust and satisfaction would increase if buyers felt they were being systematically disadvantaged in the transaction.

Second, the main results in Table 2 are consistent with an outward shift in the demand curve. The combination of a higher average transaction price and a shorter time on market (a proxy for increased quantity sold in a given period) suggests that overall demand has increased, which suggests that buyers are deriving more value from the properties they find.

Third, our decomposition analysis shows that the price effect loads on buyer-facing exposure, whereas seller-facing exposure has no effect on price (Table 2, Panel B, column 4). This pattern is inconsistent with a pure seller-bargaining story and more consistent with buyers finding higher-value matches.

Lastly, buyers seem to have a more efficient search process, as evidenced by fewer on-site viewings (Table 3) and shorter search time (Tables 2).

Taken together, the evidence points to value creation via improved matching and lower

search costs. Higher prices reflect, at least in part, the surplus generated by higher matching value rather than extraction. We emphasize that these are suggestive surplus inferences, consistent across feedback, behavior, and decomposition, but not a full decomposition surplus redistribution.

6 Field Experiment

We complement our difference-in-differences design with a randomized field experiment that provides causal validation without relying on the parallel trends assumption. The partner firm ran the experiment in Jinan (a city outside our main sample) during 2019, randomizing AI copilot access at the neighborhood level. In treated neighborhoods, the copilot was available in the agent app; in control neighborhoods, it was disabled (agents were told the feature was under maintenance). Half (34) of the 68 neighborhoods in Jinan were randomized to the control group. The experiment ran for roughly eleven months through the end of 2019.

Randomization checks show no meaningful pre-treatment differences across conditions (Appendix Table A16). We estimate intent-to-treat (ITT) effects at the neighborhood level, clustering standard errors by neighborhood, on the same main outcomes as in our DiD: sale within three months, seller and buyer time on market, transaction prices, and ratings from buyers and sellers.

Table 6 shows effects that qualitatively mirror our DiD estimates. Relative to control neighborhoods, AI access increases the probability a listing sells within three months by about 11 percentage points, reduces time on market for sellers and buyers (about 10% and 19%, respectively), and raises prices by roughly 0.8%. These results reinforce that the improvements we document stem from the copilot rather than other contemporaneous platform changes.

We use the field experiment as corroborative evidence for three reasons: the single-city scope and smaller sample, the lack of buyer– and seller–facing exposure proxies, and

the absence of weekly behavior panels in this dataset. Nonetheless, the consistency of the experimental results with our main DiD estimates increases our confidence in the primary identification strategy and the overall conclusions of the paper.

7 Conclusion

In this paper, we study the introduction of an agent-facing AI copilot in a large, two-sided real estate market. We find that the copilot improves market efficiency: listings are more likely to sell within three months, and time on market falls for both buyers and sellers. Decompositions using pre-period buyer- and seller-side app penetration indicate that these gains arise from both *same-side* and *cross-side*: AI tools aimed at one side of the market generate tangible benefits for the other. Effects are larger for less-experienced agents and in thinner markets, consistent with the technology helping intermediaries execute the core tasks of managing a two-sided search and negotiation process.

While transaction prices increase, we present several pieces of evidence: higher buyer satisfaction and trust, changes in market outcomes consistent with an outward shift in demand, and a decomposed price effect driven by buyer-side tools. The evidence suggests that the higher price reflects improved matching quality and lower search frictions, rather than with pure surplus reallocation. We view these surplus implications as suggestive, since we do not observe willingness-to-pay directly.

Our study is subject to several limitations. First, our post period covers about one year, so we may not capture longer-run equilibrium adjustments as agents and clients adapt and as the platform iterates on design. Second, our welfare evidence relies on ratings and search measures rather than direct surplus estimates. A more direct quantification of consumer surplus remains a valuable avenue for future research.

Third, external validity also merits discussion. Our setting features firm-employed agents who often intermediate both sides of a transaction, which differs from single-sided represen-

tation common in some other countries, e.g., the United States. We expect the efficiency and "leveling" effect to generalize broadly, as these address fundamental issues of reducing information friction and augmenting professional skills that exist in any market. However, the net price effect is likely context-specific and depends on relative adoption and effectiveness of buyer- versus seller-facing tools, as well as institutional details that shape agent incentives.

Overall, the evidence indicates that agent-facing AI can enhance matching efficiency in complex, two-sided marketplaces by augmenting intermediary performance and generating beneficial cross-side effects. For platform designers, our results highlight the value of embedding AI within both buyer- and seller-facing workflows and of targeting support to less-experienced intermediaries.

References

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, 138(1):1–35.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation policy and the economy*, 19(1):139–159.

Agrawal, A., Gans, J., Goldfarb, A., and Tucker, C. E. (2024). *The economics of artificial intelligence: Health Care Challenges*. University of Chicago Press.

Bertomeu, J., Lin, Y., Liu, Y., and Ni, Z. (2025). The impact of generative ai on information processing: Evidence from the ban of chatgpt in italy. *Journal of Accounting and Economics*, page 101782.

Brynjolfsson, E., Hui, X., and Liu, M. (2019). Does machine translation affect international trade? evidence from a large digital platform. *Management Science*, 65(12):5449–5460.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, page qjae044.

Brynjolfsson, E., Rock, D., and Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The economics of artificial intelligence: An agenda*, pages 23–57. University of Chicago Press.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. (2024). Difference-in-differences with a continuous treatment. Technical report, National Bureau of Economic Research.

Castillo, J. C. (2023). Who benefits from surge pricing? *Available at SSRN 3245533*.

Cong, L. W., Guo, Y., Zhao, X., and Zhou, W. (2024). Writing quality and soft information in the genai age: Evidence from online credit markets. *Available at SSRN 4959535*.

Cowgill, B., Dell'Acqua, F., Deng, S., Hsu, D., Verma, N., and Chaintreau, A. (2020). Biased programmers? or biased data? a field experiment in operationalizing ai ethics. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 679–681.

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S.,

Krayer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

Demirci, O., Hannane, J., and Zhu, X. (2025). Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science*.

Fong, J., Manchanda, P., and Song, Y. (2023). How effective is suggested pricing?: Experimental evidence from an e-commerce platform. *Experimental Evidence from an E-Commerce Platform (November 10, 2023)*.

Fu, R., Huang, Y., Mehta, N., Singh, P. V., and Srinivasan, K. (2025). Unequal impact of zestimate on the housing market. *Marketing Science*.

Fu, R., Jin, G. Z., and Liu, M. (2022). Does human-algorithm feedback loop lead to error propagation? evidence from zillow's zestimate. Technical report, National Bureau of Economic Research.

Goldberg, S. and Lam, H. T. (2025). Generative ai in equilibrium: Evidence from a creative goods marketplace. *Available at SSRN 5152649*.

Hall, J., Kendrick, C., and Nosko, C. (2015). The effects of uber's surge pricing: A case study. *The University of Chicago Booth School of Business*.

Huang, Y. (2024). Pricing frictions and platform remedies: the case of airbnb. *Available at SSRN 3767103*.

Hui, X., Reshef, O., and Zhou, L. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6):1977–1989.

Liu, M., Tang, X., Xia, S., Zhang, S., Zhu, Y., and Meng, Q. (2023). Algorithm aversion: Evidence from ridesharing drivers. *Management Science*.

Luo, X., Tong, S., Fang, Z., and Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, 38(6):937–947.

Malik, N. and Manzoor, E. (2023). Does machine learning amplify pricing errors in the housing market?–the economics of machine learning feedback loops. *arXiv preprint arXiv:2302.09438*.

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., et al. (2025). Artificial intelligence index report 2025. *arXiv preprint arXiv:2504.07139*.

Mollick, E. and Mollick, L. (2023). Assigning ai: Seven approaches for students, with prompts. *arXiv preprint arXiv:2306.10052*.

Niu, Y., Wu, J., Jiang, S., and Jiang, Z. (2025). The bullwhip effect in servitized manufacturers. *Management Science*, 71(1):1–20.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Pan, Q. and Wang, W. (2021). Costly price adjustment and automated pricing: The case of airbnb. *Available at SSRN 4077985*.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.

Peukert, C., Abeillon, F., Haese, J., Kaiser, F., and Staub, A. (2024). Strategic behavior

and ai training data. *arXiv preprint arXiv:2404.18445*.

Schanke, S., Burtch, G., and Ray, G. (2021). Estimating the impact of "humanizing" customer service chatbots. *Information Systems Research*, 32(3):736–751.

Shi, A., Chan, T., and Zhang, D. (2024). Using algorithmic scores to measure the impacts of targeting promotional messages. *Available at SSRN 4262156*.

Wiles, E., Munyikwa, Z., and Horton, J. (2025). Algorithmic writing assistance on jobseekers' resumes increases hires. *Management Science*.

Yu, S. (2020). Algorithmic outputs as information source: The effects of zestimates on home prices and racial bias in the housing market. *Available at SSRN 3584896*.

Zhang, M., Sun, T., Luo, L., and Golden, J. (2021a). Consumer and ai co-creation: When and why human participation improves ai creation. *USC Marshall School of Business Research Paper Sponsored by iORB, No. Forthcoming*.

Zhang, S., Mehta, N., Singh, P. V., and Srinivasan, K. (2021b). Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? an analysis in the context of airbnb. *Marketing Science*, 40(5):813–820.

Zhou, E. and Lee, D. (2024). Generative artificial intelligence, human creativity, and art. *PNAS nexus*, 3(3):pgae052.

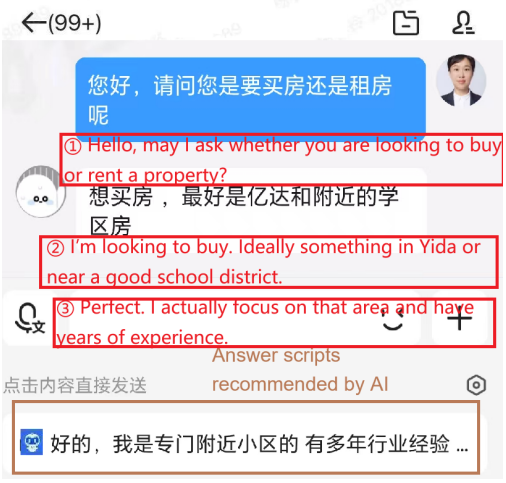
Zhu, K. and Walker, D. (2024). The promise and pitfalls of ai technology in bridging digital language divide: Insights from machine translation on wikipedia. *Available at SSRN 4708614*.

Figures and Tables

Figure 1: Sample AI Output: The App Interface for Agents



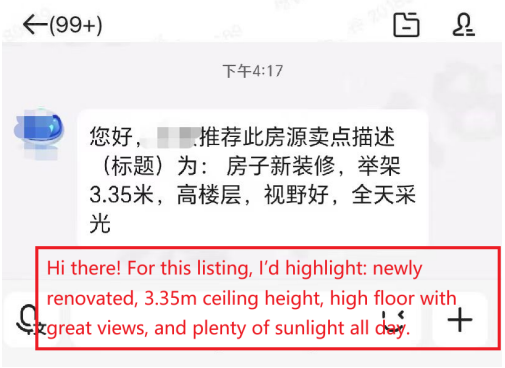
(a) AI-recommended Listings (Functioning for buyers)



(b) AI-recommended Dialogue (Functioning for buyers)



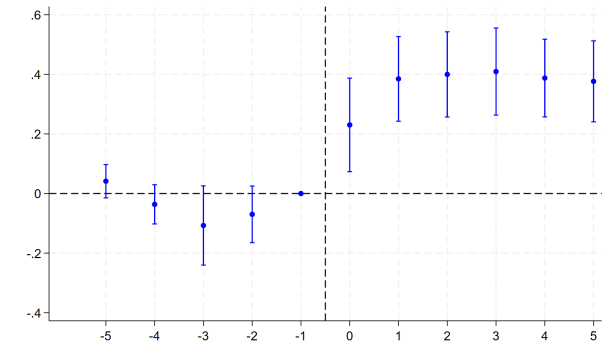
(c) AI-recommended Price adjustments (Functioning for Sellers)



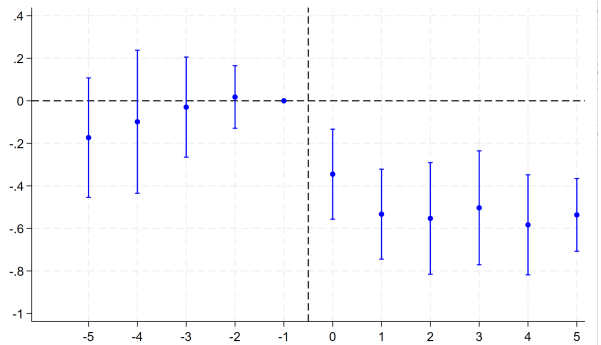
(d) AI-recommended Selling-point Descriptions (Functioning for Sellers)

Notes. This figure, as an example, displays AI-generated recommendations for agents and the corresponding app interface for agents. These recommendations are respectively targeted at buyers and sellers. Panel (a) displays the specific properties recommended by the AI based on buyers' needs. Panel (b) shows an example where the buyer expressed interests in properties in highly ranked school zones, the AI recommended the agent to use a friendly, professional reply highlighting the agent's experience and expertise in the neighborhood. Panel (c) shows the AI's suggestion for a recent price adjustment to the listed property. Panel (d) provides appealing, optimized property title suggestions designed to attract buyers and improve conversion rates. It shows an example of a recommended title: "Beautifully renovated home with 3.35-meter ceilings, high floor with great views, and all-day sunlight."

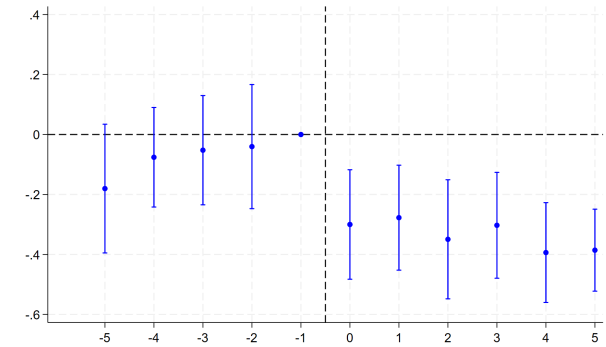
Figure 2: Event Studies for Main Outcomes: Total Effects of AI Adoption



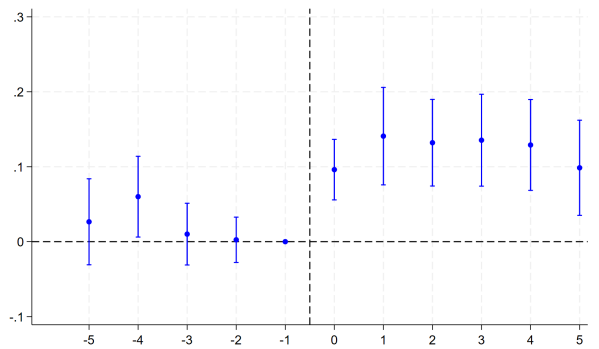
(a) Sold<3month (using Ratio_tran)



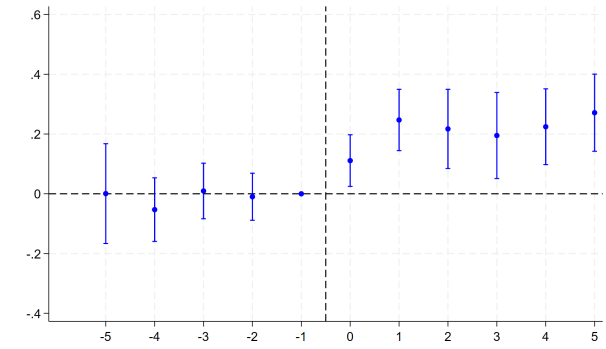
(b) Log_days_seller (using Ratio_tran)



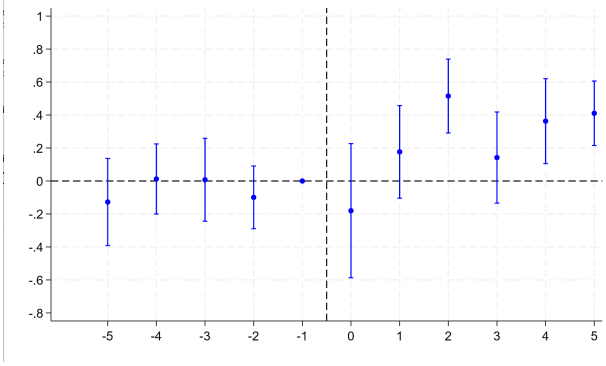
(c) Log_days_buyer (using Ratio_tran)



(d) Log_price (using Ratio_tran)



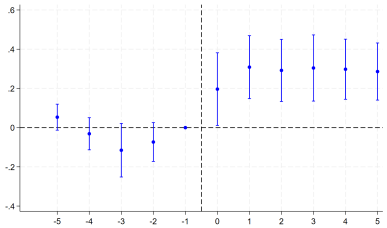
(e) Buyers' ratings (using Ratio_tran)



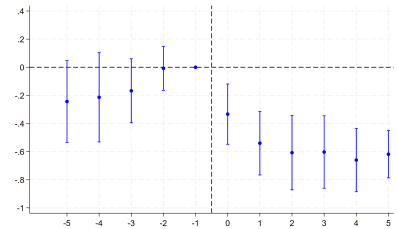
(f) Sellers' ratings (using Ratio_tran)

Notes. This figure plots the bimonthly coefficients and 95% confidence intervals from event study regressions (Equation 2) on the following outcomes: whether the listed property is sold within three months, time on market for sellers, time on market for buyers, unit transaction prices and buyers' and sellers' ratings. Standard errors are clustered at the district level.

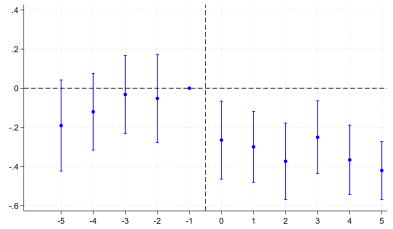
Figure 3: Event Studies for Main Outcomes: Demand-side and Supply-side Effects of AI Adoption



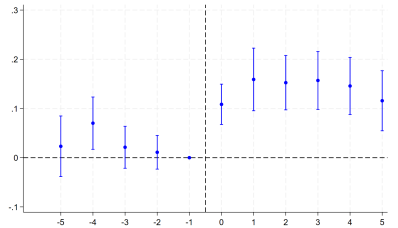
(a) Sold<3month (Ratio_buyer)



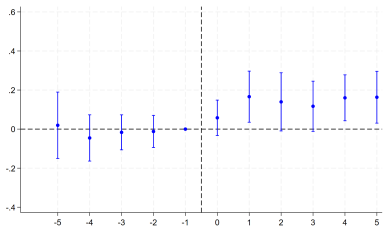
(b) Log_days_seller (Ratio_buyer)



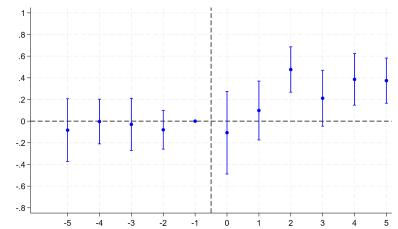
(c) Log_days_buyer (Ratio_buyer)



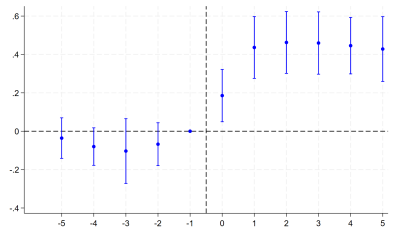
(d) Log_price (Ratio_buyer)



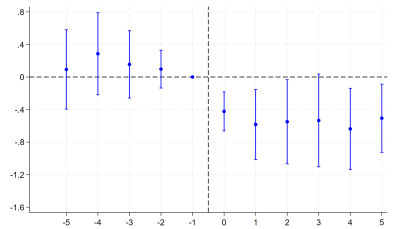
(e) Rating_buyer (Ratio_buyer)



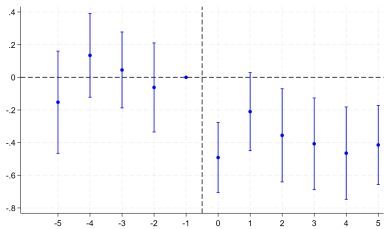
(f) Rating_seller (Ratio_buyer)



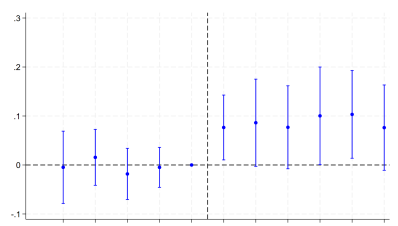
(g) Sold<3month (Ratio_seller)



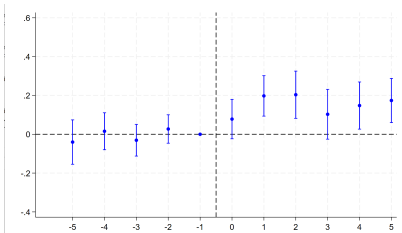
(h) Log_days_seller (Ratio_seller)



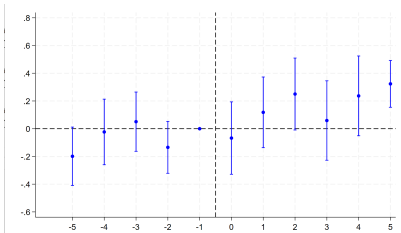
(i) Log_days_buyer (Ratio_seller)



(j) Log_price (Ratio_seller)



(k) Rating_buyer (Ratio_seller)



(l) Rating_seller (Ratio_seller)

Notes. This figure plots the bimonthly coefficients and 95% confidence intervals from event study regressions (Equation 2) regarding demand-side and supply-side on the following outcomes: whether the listed property is sold within three months, time on market for sellers, time on market for buyers, and unit transaction prices and buyers' and sellers' ratings. Ratio_buyer and Ratio_seller are the pre-treatment district shares of buyers and sellers transacting via the app, respectively. These proxies map to exposure to buyer-facing vs seller-facing workflows because those features live in the agent–buyer vs agent–seller chat flows. Standard errors are clustered at the district level.

Table 1: Summary Statistics

Variable	Mean	SD	P10	Median	P90
Sold<3month	0.546	0.498	0	1	1
Log_days_seller	3.982	1.256	2.197	4.094	5.533
Log_days_buyer	2.664	1.656	0	2.639	4.905
Log_price_unit	10.274	0.763	9.326	10.267	11.275
Rating (buyer)	4.884	0.392	5	5	5
Rating (seller)	4.693	0.632	4	5	5
Log_listed_price	10.308	0.767	9.357	10.305	11.315
Discount	0.032	0.042	0	0.028	0.067
Ratio_tran	0.446	0.222	0.110	0.39	0.747
Ratio_buyer	0.388	0.208	0.117	0.353	0.716
Ratio_seller	0.265	0.162	0.085	0.276	0.444
Building size	84.28	34.178	49	80	128
Building age	3.347	1.507	1	3	5
Floor	2.118	0.801	1	2	3
Decoration	2.475	0.642	2	3	3
Room	2.183	0.794	1	2	3
School	0.478	0.5	0	0	1
Subway	0.525	0.499	0	1	1
Population	167.032	107.259	59	123	361
Unemployment	2.065	0.889	0.720	2	3.1
Gdp	12.432	7.129	5.268	9.989	19.284
Income	5.49	1.356	3.955	4.978	7.555
Education	0.092	0.288	0	0	0
Gender	0.684	0.465	0	1	1
Performance	3.43	1.36	1.130	4.02	4.71
Tenure	4.879	3.278	1	4	10

Notes. Table presents summary statistics for the 231,680 transacted properties. *building size* represents the indoor area in square meters. *building age* is an indicator for the building's age: 1 for 0-5 years, 2 for 5-10 years, 3 for 10-15 years, 4 for 15-20 years, 5 for 20-30 years, and 6 for over 30 years. $floor$ indicates the property's floor level: 0 for low, 1 for middle, and 2 for high. *decoration* is an indicator for decoration status: 1 for un-decorated, 2 for lightly decorated, and 3 for well-decorated. *room* is an indicator of number of rooms: 1 for one bedroom, 2 for two bedrooms, 3 for three bedrooms, 4 for four bedrooms, and 5 for more than four bedrooms. *school* is a dummy variable indicating proximity to a school: 1 for close, 0 otherwise. *subway* is a dummy variable indicating proximity to a subway: 1 for close, 0 otherwise. *Population* is the district's population in ten thousand people; $Income$ is the district-level per capita net income in ten thousand RMB; *Gdp* is the district's per capita GDP in ten thousand RMB; and $Unemployment$ is the district-level unemployment rate. "Rating (buyer)" and "rating (seller)" refer to the 5-point scale ratings (from 1 to 5) provided by buyers and sellers, respectively.

Table 2: Main Effects of AI Exposure on Market Outcomes

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.0372	0.5006***	0.2087***	-0.0392***	-0.0636*	-0.1947***
	(0.037)	(0.065)	(0.025)	(0.012)	(0.035)	(0.051)
Ratio_tran	-0.0602	0.0906	0.3763***	-0.0715	-0.1111**	-0.2656
	(0.060)	(0.146)	(0.104)	(0.092)	(0.047)	(0.192)
Post×Ratio_tran	0.4215***	-0.4434***	-0.2484***	0.1005***	0.2206***	0.3481***
	(0.080)	(0.145)	(0.047)	(0.029)	(0.053)	(0.061)
Panel B: Decomposed Effects and Cross-Side Spillovers						
Post	-0.0263	0.5234***	0.2202***	-0.0411***	-0.0567	-0.1897***
	(0.041)	(0.068)	(0.025)	(0.012)	(0.035)	(0.046)
Ratio_buyer	0.1221*	0.0687	0.3275***	-0.0197	-0.0424	0.0536
	(0.072)	(0.139)	(0.123)	(0.068)	(0.056)	(0.263)
Post×Ratio_buyer	0.1744*	-0.2376**	-0.1411**	0.1152***	0.1340***	0.2936***
	(0.100)	(0.107)	(0.055)	(0.028)	(0.041)	(0.056)
Ratio_seller	-0.3070***	0.2173	-0.0168	-0.1440	-0.0457	-0.1689
	(0.102)	(0.263)	(0.214)	(0.107)	(0.051)	(0.280)
Post×Ratio_seller	0.4117***	-0.4877**	-0.2550***	-0.0005	0.1525***	0.1354**
	(0.103)	(0.229)	(0.089)	(0.033)	(0.048)	(0.055)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.222	0.079	0.128	0.948	0.047	0.079
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Table presents the estimation results of Equation 1. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective region. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Demand-Side Mechanisms: Buyer Feedback, Trust, and Search Effort

	(1) Recomm_ optimal	(2) Talk_ inefficiency	(3) Talk_ comfort	(4) Trust	(5) Viewing	(6) Log_listed _price
Panel A: Overall Impact						
Post	-0.2835***	0.0510***	-0.0205*	-0.0278***	-0.0147	-0.0364***
	(0.049)	(0.015)	(0.012)	(0.010)	(0.085)	(0.012)
Ratio_tran	-0.1539	0.0242*	-0.0038	-0.0320	-0.1711	-0.0584
	(0.125)	(0.013)	(0.013)	(0.026)	(0.224)	(0.097)
Post×Ratio_tran	0.3944***	-0.0683***	0.0688**	0.0575**	-0.0341	0.0963***
	(0.092)	(0.024)	(0.029)	(0.023)	(0.143)	(0.029)
Panel B: Decomposed Effects						
Post	-0.2638***	0.0461***	-0.0151	-0.0266***	0.0287	-0.0384***
	(0.042)	(0.014)	(0.011)	(0.009)	(0.082)	(0.012)
Ratio_buyer	-0.3274**	0.0161	-0.0048	0.0057	0.1783	0.0067
	(0.145)	(0.011)	(0.013)	(0.029)	(0.219)	(0.068)
Post×Ratio_buyer	0.2324**	-0.0608***	0.0714***	0.0591**	-0.3548**	0.1057***
	(0.109)	(0.016)	(0.026)	(0.023)	(0.163)	(0.029)
Ratio_seller	0.2167	0.0133	-0.0051	-0.0614**	-0.6711*	-0.1752
	(0.149)	(0.013)	(0.012)	(0.026)	(0.360)	(0.110)
Post×Ratio_seller	0.2587**	-0.0129	-0.0019	0.0099	0.3094	0.0081
	(0.101)	(0.017)	(0.014)	(0.012)	(0.190)	(0.031)
Observations	75,596	75,596	75,596	75,596	13,581,455	231,680
Adjusted R^2	0.201	0.020	0.047	0.044	0.091	0.949
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table reports the regression results from Equation 1, focusing on the demand-side mechanisms. Columns (1)-(4) show effects on buyer feedback. Column 5 shows effects on weekly on-site viewing frequency. Column 6 shows effects on the last listed prices of sold properties. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Supply-Side Mechanisms: Pricing Advice and Listing Strategy

	(1) Pricing_ optimal	(2) Sellpoint_ Accuracy	(3) Total_adjust _times	(4) Total_adjust _mag
Panel A: Overall Impact				
Post	-0.2437***	-0.1999***	-0.0103	0.1355
	(0.040)	(0.042)	(0.007)	(0.087)
Ratio_tran	-0.0598	-0.0239	-0.0362***	0.5318**
	(0.084)	(0.083)	(0.012)	(0.261)
Post×Ratio_tran	0.2465***	0.3270***	0.0639***	-0.4054**
	(0.065)	(0.087)	(0.013)	(0.155)
Panel B: Decomposed Effects				
Post	-0.2262***	-0.2003***	-0.0115	0.0516
	(0.033)	(0.028)	(0.007)	(0.097)
Ratio_buyer	0.0652	0.0672	-0.0199	-0.2197
	(0.084)	(0.094)	(0.017)	(0.263)
Post×Ratio_buyer	0.0041	0.0746	0.0204	0.2791
	(0.070)	(0.051)	(0.027)	(0.204)
Ratio_seller	-0.1749**	-0.1448	-0.0359	0.9377*
	(0.081)	(0.102)	(0.023)	(0.513)
Post×Ratio_seller	0.3331***	0.4149***	0.0820***	-0.7933***
	(0.058)	(0.053)	(0.031)	(0.244)
Observations	20,336	20,336	13,581,455	2,171,883
Adjusted R^2	0.125	0.134	0.027	0.054
Controls	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes

Notes. This table presents regression results based on Equation 1 for the supply-side mechanisms. Columns (1) and (2) show effects on seller feedback labels. Columns (3) and (4) show effects on the frequency and magnitude of weekly price adjustments. All columns include location and City×Month fixed effects and standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 5: Heterogeneity in the Effect of AI Adoption by Agents' Tenure

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Moderation of Overall Effect						
Post×Ratio_tran	0.518***	-0.508***	-0.315***	0.101***	0.2954***	0.5759***
	(0.082)	(0.149)	(0.052)	(0.029)	(0.060)	(0.182)
Post×Ratio_tran×Tenure	-0.0341***	0.0222***	0.0249***	-0.0002	-0.0147**	-0.0683**
	(0.003)	(0.008)	(0.007)	(0.001)	(0.007)	(0.033)
Panel B: Moderation of Decomposed Effects						
Post×Ratio_buyer	0.258**	-0.297***	-0.180***	0.117***	0.2064***	0.6272***
	(0.102)	(0.110)	(0.058)	(0.028)	(0.041)	(0.162)
Post×Ratio_buyer×Tenure	-0.0308***	0.0229**	0.0158	-0.0005	-0.0144**	-0.0838**
	(0.004)	(0.009)	(0.010)	(0.001)	(0.007)	(0.034)
Post×Ratio_seller	0.448***	-0.505**	-0.323***	0.001	0.1420***	-0.0728
	(0.105)	(0.231)	(0.093)	(0.034)	(0.043)	(0.172)
Post×Ratio_seller×Tenure	-0.0109*	0.0053	0.0247*	0.0015	0.0021	0.0309
	(0.007)	(0.009)	(0.013)	(0.002)	(0.006)	(0.037)
Observations	267,882	231,680	231,680	231,680	75,595	20,336
Adjusted	0.224	0.079	0.128	0.948	0.049	0.114
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents heterogeneity analysis by agent's tenure (continuous). For brevity, only the highest-order interaction terms and their constituent two-way interactions are shown. All specifications include the full set of main effects and controls based on Equation 1 Standard errors clustered at the district level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Field Experiment: Intent-to-Treat Effects of Copilot Access

	(1) Sold< 3 month	**(2)** Log_days _seller	**(3)** Log_days _buyer	**(4)** Log_price _unit	**(5)** Rating (buyer)	**(6)** Rating (seller)
Post	-0.0034	0.4739***	0.4712***	-0.1084***	-0.0689**	-0.0632
	(0.029)	(0.064)	(0.090)	(0.015)	(0.033)	(0.074)
Post×Treat	0.1106***	-0.3704***	-0.4506***	0.0817***	0.1247***	0.2323**
	(0.040)	(0.096)	(0.121)	(0.017)	(0.041)	(0.096)
Building size	-0.0022***	0.0074***	0.0074***	-0.0003	-0.0002	0.0020
	(0.000)	(0.001)	(0.002)	(0.000)	(0.000)	(0.001)
Building age	0.0092	-0.0152	0.0329	-0.0239***	0.0033	-0.0108
	(0.006)	(0.018)	(0.023)	(0.007)	(0.015)	(0.061)
Floor	-0.0093	0.0838***	0.0049	-0.0245***	-0.0072	0.0445
	(0.010)	(0.027)	(0.032)	(0.004)	(0.011)	(0.027)
Decoration	0.0688***	-0.1246***	0.0531	0.0369***	0.0096	0.0383
	(0.013)	(0.031)	(0.043)	(0.005)	(0.019)	(0.061)
Room	0.0193	-0.0625	-0.0877*	0.0420***	0.0227*	-0.0808
	(0.018)	(0.050)	(0.050)	(0.014)	(0.011)	(0.057)
Observations	6,748	4,031	4,031	4,031	1,378	561
Adjusted R^2	0.110	0.100	0.033	0.650	0.128	0.031
District×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table presents the estimation results of the field experiment. The definitions of control variables are the same as in Table 2 and Appendix Table A3. Standard errors clustered at the neighborhood level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix A. Table and Figure

Figure A1: Flow chart

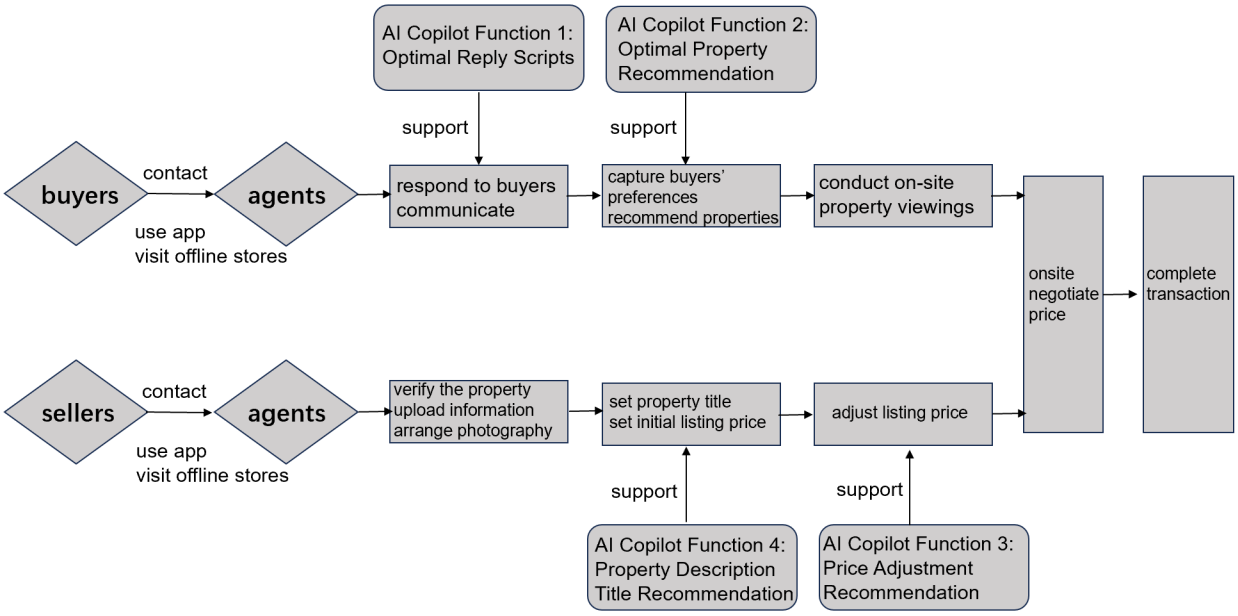


Table A3: Main Effects: Full Table [h]

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.0372	0.5006***	0.2087***	-0.0392***	-0.0636*	-0.1947***
	(0.037)	(0.065)	(0.025)	(0.012)	(0.035)	(0.051)
Ratio_tran	-0.0602	0.0906	0.3763***	-0.0715	-0.1111**	-0.2656
	(0.060)	(0.146)	(0.104)	(0.092)	(0.047)	(0.192)
Post×Ratio_tran	0.4215***	-0.4434***	-0.2484***	0.1005***	0.2206***	0.3481***
	(0.080)	(0.145)	(0.047)	(0.029)	(0.053)	(0.061)
Population	-0.0001	0.0002	0.0004	0.0001	-0.0000	0.0009*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Unemployment	-0.0043	-0.0709	0.1108	0.0045	-0.0166	0.1172
	(0.017)	(0.086)	(0.070)	(0.041)	(0.026)	(0.103)

—continued on next page—

—*continued from previous page*—

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Gdp	-0.0023	0.0001	0.0137**	0.0144***	-0.0016	0.0157***
	(0.002)	(0.008)	(0.006)	(0.003)	(0.002)	(0.006)
Income	-0.0052	0.0109	-0.0432	0.0592**	0.0257***	-0.1233***
	(0.014)	(0.051)	(0.027)	(0.024)	(0.009)	(0.018)
Building size	-0.0014***	0.0060***	0.0044***	-0.0009***	-0.0001	0.0003
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Building age	-0.1134***	0.0235***	-0.0413**	-0.0025	-0.0023	-0.0019
	(0.002)	(0.008)	(0.019)	(0.002)	(0.004)	(0.009)
Floor	0.0225***	0.0264***	-0.0043	-0.0086***	0.0015	-0.0025
	(0.003)	(0.007)	(0.003)	(0.002)	(0.002)	(0.006)
Decoration	0.0783***	-0.1307***	-0.0059	0.0511***	-0.0003	0.0124*
	(0.003)	(0.009)	(0.007)	(0.003)	(0.002)	(0.006)
Room	0.0160***	-0.0599***	0.0541***	0.0336***	0.0047*	-0.0123
	(0.005)	(0.013)	(0.008)	(0.007)	(0.003)	(0.008)
School	0.0399	-0.1612	0.1702	-0.1361**	0.0028	-0.0290*
	(0.049)	(0.176)	(0.149)	(0.065)	(0.004)	(0.016)
Subway	0.0072***	-0.0058	0.0098	0.0121**	0.0018	0.0058
	(0.002)	(0.008)	(0.010)	(0.005)	(0.003)	(0.009)
Constant	0.4951***	3.7643***	1.6696***	9.7200***	4.8407***	4.8859***
	(0.078)	(0.361)	(0.206)	(0.160)	(0.079)	(0.337)
Panel B: Decomposed Effects and Cross-Side Effects						
Post	-0.0263	0.5234***	0.2202***	-0.0411***	-0.0567	-0.1897***
	(0.041)	(0.068)	(0.025)	(0.012)	(0.035)	(0.046)
Ratio_buyer	0.1221*	0.0687	0.3275***	-0.0197	-0.0424	0.0536
	(0.072)	(0.139)	(0.123)	(0.068)	(0.056)	(0.263)
Post×Ratio_buyer	0.1744*	-0.2376**	-0.1411**	0.1152***	0.1340***	0.2936***
	(0.100)	(0.107)	(0.055)	(0.028)	(0.041)	(0.056)
Ratio_seller	-0.3070***	0.2173	-0.0168	-0.1440	-0.0457	-0.1689
	(0.102)	(0.263)	(0.214)	(0.107)	(0.051)	(0.280)
Post×Ratio_seller	0.4117***	-0.4877**	-0.2550***	-0.0005	0.1525***	0.1354**
	(0.103)	(0.229)	(0.089)	(0.033)	(0.048)	(0.055)
Population	-0.0001	0.0002	0.0004	0.0001	-0.0000	0.0009*

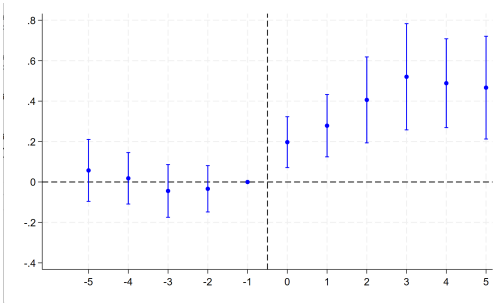
—*continued on next page*—

—*continued from previous page*—

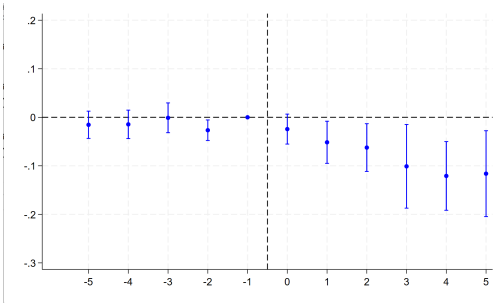
	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Unemployment	0.0040	-0.0644	0.1271*	0.0080	-0.0205	0.0972
	(0.018)	(0.087)	(0.073)	(0.041)	(0.025)	(0.125)
Gdp	-0.0025	-0.0002	0.0127**	0.0140***	-0.0015	0.0148*
	(0.002)	(0.008)	(0.006)	(0.004)	(0.002)	(0.008)
Income	-0.0048	0.0084	-0.0436	0.0585**	0.0250***	-0.1207***
	(0.014)	(0.052)	(0.027)	(0.024)	(0.009)	(0.024)
Building size	-0.0014***	0.0060***	0.0044***	-0.0009***	-0.0001	0.0003
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Building age	-0.1137***	0.0234***	-0.0413**	-0.0025	-0.0023	-0.0016
	(0.002)	(0.008)	(0.019)	(0.002)	(0.004)	(0.009)
Floor	0.0224***	0.0264***	-0.0043	-0.0086***	0.0015	-0.0024
	(0.003)	(0.007)	(0.003)	(0.002)	(0.002)	(0.006)
Decoration	0.0782***	-0.1307***	-0.0058	0.0511***	-0.0003	0.0124*
	(0.003)	(0.009)	(0.007)	(0.003)	(0.002)	(0.006)
Room	0.0160***	-0.0600***	0.0541***	0.0337***	0.0046	-0.0122
	(0.005)	(0.013)	(0.008)	(0.007)	(0.003)	(0.008)
School	0.0531	-0.1188	0.1876	-0.1170*	0.0025	-0.0289*
	(0.054)	(0.173)	(0.153)	(0.068)	(0.004)	(0.016)
Subway	0.0072***	-0.0059	0.0096	0.0120**	0.0017	0.0061
	(0.002)	(0.008)	(0.010)	(0.005)	(0.003)	(0.009)
Constant	0.4778***	3.7168***	1.6820***	9.7304***	4.8288***	4.8114***
	(0.086)	(0.361)	(0.211)	(0.153)	(0.078)	(0.390)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.222	0.079	0.128	0.948	0.047	0.080
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Table presents the full estimation results of Equation 1. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

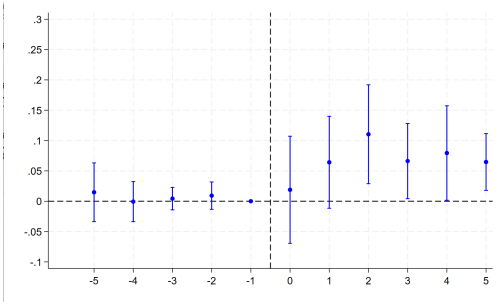
Figure A2: Event Studies for Demand-side Mechanisms



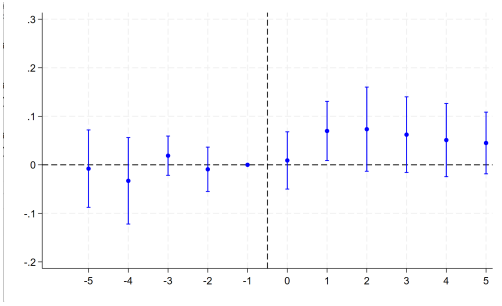
(a) recommend_optimal (using Ratio_tran)



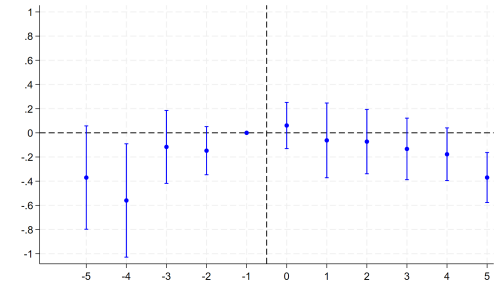
(b) Talk_inefficiency (using Ratio_tran)



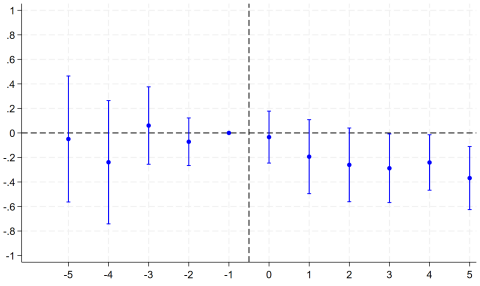
(c) Talk_comfort (using Ratio_tran)



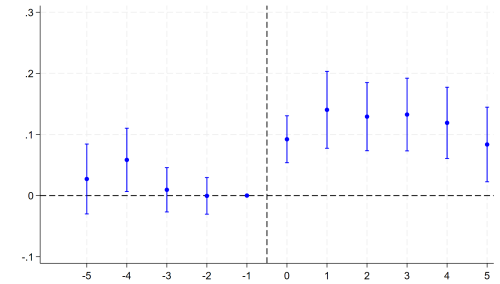
(d) Trust (using Ratio_tran)



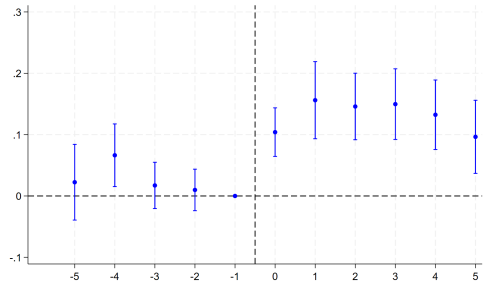
(e) Viewing (using Ratio_tran)



(f) Viewing (using Ratio_buyer)



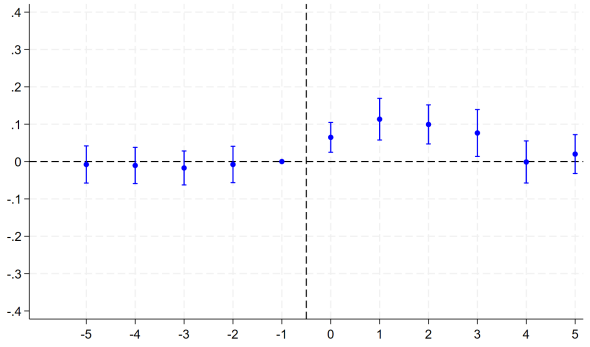
(g) Log_listed_price (using Ratio_tran)



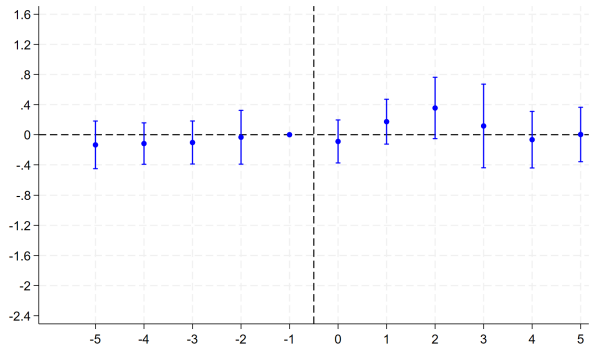
(h) Log_listed_price (using Ratio_buyer)

Notes. This figure plots the bimonthly coefficients and 95% confidence intervals from event study regressions (Equation 2) regarding demand-side outcomes. Panels (a)-(d) show effects on buyer feedback labels. Panels (e)-(h) show effects on weekly on-site viewing frequency and the last listed prices of sold properties, presenting results for both the overall (Ratio_tran) and decomposed (Ratio_buyer) specifications. *Ratio_buyer* is the pre-treatment district shares of buyers transacting via the app. Standard errors are clustered at the district level.

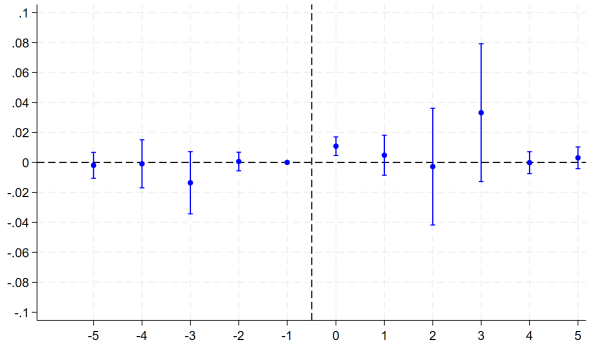
Figure A3: Event Studies for Upward and Downward Price Adjustments



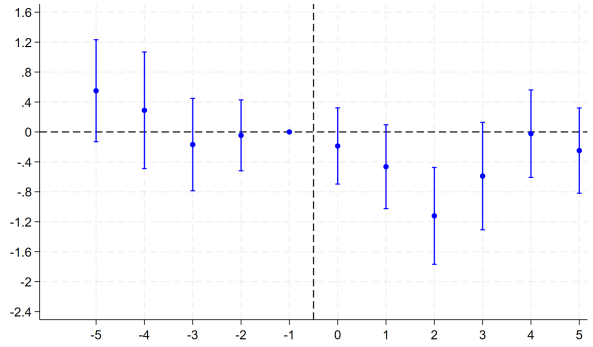
(a) Down_adjust_times (using Ratio_tran)



(b) Down_adjust_mag (using Ratio_tran)



(c) Up_adjust_times (using Ratio_tran)



(d) Up_adjust_mag (using Ratio_tran)

Notes. This figure plots the bimonthly coefficients and 95% confidence intervals from event study regressions (Equation 2) regarding supply-side price adjustment behaviors. Specifically, Down_adjust_times and Down_adjust_mag capture the weekly frequency and magnitude of downward price adjustments, while Up_adjust_times and Up_adjust_mag capture the weekly frequency and magnitude of upward price adjustments. Standard errors are clustered at the district level.

Table A1: Additional Summary statistics

	Mean	SD	P10	Median	P90
Demand side and supply side					
Total_adjust_times	0.187	0.533	0	0	1
Up_adjust_times	0.038	0.206	0	0	0
Down_adjust_times	0.149	0.469	0	0	1
Total_adjust_mag	4.151	3.628	1.150	3.16	8.29
Up_adjust_mag	4.75	4.684	1.120	3.33	10
Down_adjust_mag	3.96	3.325	1.090	3.09	7.8
Viewing	1.639	2.871	0	1	4
Log_volume	1.365	0.63	0.693	1.099	2.303
Num_list	233.495	228.641	42	166	495
Textual feedback (buyer)					
Recommend_optimal	0.368	0.482	0	0	1
Intro_well	0.097	0.291	0	0	0
Talk_inefficiency	0.01	0.097	0	0	0
Talk_comfort	0.012	0.108	0	0	0
Trust	0.042	0.201	0	0	0
Appearance	0.421	0.494	0	0	1
Textual feedback (seller)					
Pricing_optimal	0.255	0.436	0	0	1
Sellingpoint_accuracy	0.241	0.427	0	0	0
Talk_inefficiency	0.006	0.080	0	0	0
Talk_comfort	0.009	0.093	0	0	0
Trust	0.047	0.213	0	0	0
Appearance	0.252	0.434	0	0	1

Notes. This table presents summary statistics for the additional variables in the demand-side and supply-side datasets (which include 13,581,455 data points), as well as buyers' and sellers' evaluation data in the feedback system (which contain 75,596 and 20,336 data points, respectively).Specifically, Total_adjust_times refers to the number of price adjustments made to a property each week. Up_adjust_times refers to the number of upward price adjustments made to a property in a given week. Down_adjust_times refers to the number of downward price adjustments made to a property in a given week. Total_adjust_mag, Up_adjust_mag, and Down_adjust_mag respectively denote the total magnitude (absolute value) of price adjustments for the property during the week, the magnitude of upward price adjustments, and the magnitude of downward price adjustments. Viewing refers to the on-site viewing frequency of the property during the week. Volume represents the weekly transaction volume at the neighborhood level. Num_list represents the number of listed properties with the same number of rooms as the sold property in the neighborhood during the pre-treatment period. The textual feedback variables from Buyer and Seller indicate whether specific labels were selected in the feedback system. For instance, Recommend_optimal reflects that the property recommendations were precise, while Intro_well suggests that the agent's explanations were clear and effective. Talk_inefficiency denotes that the communication was ineffective, and Trust signals that the agent was perceived as trustworthy. Appearance captures positive evaluations of the agent's personal image and demeanor, whereas Pricing_optimal reflects that the pricing advice was considered appropriate. Moreover, Sellingpoint_accuracy indicates that the agent accurately highlight the key selling points of the property. In addition, we construct an extra variable, Talk_comfort, based on the review text to capture explicit positive evaluations of the agent's communication skills.

Table A2: Impact of AI Adoption: Estimation Results after Dropping Cross-District Neighborhoods

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.0399	0.5082***	0.2016***	-0.0453***	-0.0736**	-0.2132***
	(0.040)	(0.070)	(0.027)	(0.011)	(0.034)	(0.057)
Ratio_tran	-0.0373	0.2118	0.5335**	-0.1471	-0.1212**	-0.2631
	(0.114)	(0.237)	(0.221)	(0.157)	(0.054)	(0.201)
Post×Ratio_tran	0.4315***	-0.4740***	-0.2496***	0.1149***	0.2434***	0.3966***
	(0.083)	(0.159)	(0.051)	(0.026)	(0.052)	(0.073)
Panel B: Decomposed Effects and Cross-Side Spillovers						
Post	-0.0256	0.5312***	0.2113***	-0.0470***	-0.0702**	-0.2129***
	(0.044)	(0.073)	(0.027)	(0.011)	(0.034)	(0.050)
Ratio_buyer	0.1661	0.9074***	0.8604***	-0.0127	-0.0517	0.0453
	(0.116)	(0.307)	(0.243)	(0.196)	(0.060)	(0.268)
Post×Ratio_buyer	0.1712*	-0.2653**	-0.1588***	0.1249***	0.1535***	0.3653***
	(0.102)	(0.110)	(0.052)	(0.025)	(0.040)	(0.052)
Ratio_seller	-0.1673	-0.6021*	-0.4057	-0.2892*	-0.0701	-0.1536
	(0.144)	(0.345)	(0.354)	(0.153)	(0.049)	(0.283)
Post×Ratio_seller	0.4212***	-0.4964**	-0.2219***	0.0088	0.1741***	0.1358**
	(0.105)	(0.232)	(0.083)	(0.030)	(0.045)	(0.060)
Observations	225,206	194,133	194,133	194,133	61,012	16,385
Adjusted R^2	0.224	0.079	0.126	0.946	0.050	0.072
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1 after excluding neighborhoods that span across two districts. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective region. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A4: Robustness to Treatment Interacted with Controls

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.6380***	1.3622***	0.3519	-0.0917*	-0.0805	-0.3165
	(0.211)	(0.269)	(0.220)	(0.053)	(0.131)	(0.258)
Ratio_tran	0.0065	0.2368	0.4283***	-0.0858	-0.1013**	-0.2511
	(0.060)	(0.148)	(0.114)	(0.089)	(0.047)	(0.194)
Post×Ratio_tran	0.3113***	-0.6810***	-0.3340***	0.1201***	0.1906***	0.3243***
	(0.085)	(0.129)	(0.060)	(0.023)	(0.048)	(0.065)
Panel B: Decomposed Effects and Cross-Side Effects						
Post	-0.6448***	1.4293***	0.3840*	-0.0963*	0.1215***	0.2939***
	(0.204)	(0.272)	(0.217)	(0.054)	(0.039)	(0.066)
Ratio_buyer	0.1587**	0.1690	0.3635***	-0.0300	-0.0352	0.0609
	(0.076)	(0.138)	(0.129)	(0.066)	(0.057)	(0.265)
Post×Ratio_buyer	0.1209	-0.3880***	-0.1969***	0.1261***	0.1215***	0.2939***
	(0.102)	(0.076)	(0.061)	(0.024)	(0.039)	(0.066)
Ratio_seller	-0.2742***	0.2694	-0.0082	-0.1488	-0.0451	-0.1795
	(0.095)	(0.250)	(0.214)	(0.109)	(0.051)	(0.280)
Post×Ratio_seller	0.3337***	-0.5825***	-0.2745***	0.0104	0.1370***	0.1339**
	(0.093)	(0.189)	(0.087)	(0.028)	(0.049)	(0.060)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.229	0.081	0.128	0.948	0.049	0.081
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1, with additional controls (i.e., the interaction term between *Post* and four district-level variables). Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, Gdp, and their interaction terms with *Post* in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A5: Robustness to Alternative Clustering: Clustered at the District and Week Two-Way Level

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.0372	0.5006***	0.2087***	-0.0392***	-0.0636*	-0.1947***
	(0.037)	(0.066)	(0.025)	(0.013)	(0.035)	(0.052)
Ratio_tran	-0.0602	0.0906	0.3763***	-0.0715	-0.1111**	-0.2656
	(0.060)	(0.141)	(0.097)	(0.092)	(0.044)	(0.208)
Post×Ratio_tran	0.4215***	-0.4434***	-0.2484***	0.1005***	0.2206***	0.3481***
	(0.080)	(0.145)	(0.046)	(0.029)	(0.053)	(0.062)
Panel B: Decomposed Effects and Cross-Side Effects						
Post	-0.0263	0.5234***	0.2202***	-0.0411***	-0.0567	-0.1897***
	(0.053)	(0.069)	(0.026)	(0.012)	(0.035)	(0.047)
Ratio_buyer	0.1221*	0.0687	0.3275***	-0.0197	-0.0424	0.0536
	(0.063)	(0.139)	(0.096)	(0.068)	(0.059)	(0.252)
Post×Ratio_buyer	0.1744*	-0.2376**	-0.1411***	0.1152***	0.1340***	0.2936***
	(0.097)	(0.108)	(0.053)	(0.029)	(0.041)	(0.057)
Ratio_seller	-0.3070***	0.2173	-0.0168	-0.1440	-0.0457	-0.1689
	(0.096)	(0.256)	(0.210)	(0.106)	(0.060)	(0.290)
Post×Ratio_seller	0.4117***	-0.4877**	-0.2550***	-0.0005	0.1525***	0.1354**
	(0.104)	(0.231)	(0.092)	(0.033)	(0.048)	(0.056)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.222	0.079	0.128	0.948	0.047	0.079
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1, with standard errors clustered at both the week and district levels. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A6: Robustness to Alternative Clustering: Clustered at the City Level

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post	-0.0372	0.5006***	0.2087***	-0.0392**	-0.0682	-0.0682
	(0.071)	(0.100)	(0.050)	(0.014)	(0.055)	(0.055)
Ratio_tran	-0.0602	0.0906	0.3763***	-0.0715	-0.1265	-0.1265
	(0.082)	(0.281)	(0.087)	(0.155)	(0.103)	(0.103)
Post×Ratio_tran	0.4215***	-0.4434*	-0.2484***	0.1005*	0.2312**	0.2312**
	(0.130)	(0.225)	(0.064)	(0.049)	(0.081)	(0.081)
Panel B: Decomposed Effects and Cross-Side Effects						
Post	-0.0263	0.5234***	0.2202***	-0.0411***	-0.0596	-0.0596
	(0.086)	(0.099)	(0.048)	(0.013)	(0.053)	(0.053)
Ratio_buyer	0.1221	0.0687	0.3275*	-0.0197	-0.0115	-0.0115
	(0.094)	(0.269)	(0.165)	(0.109)	(0.132)	(0.132)
Post×Ratio_buyer	0.1744	-0.2376	-0.1411**	0.1152**	0.1511**	0.1511**
	(0.130)	(0.152)	(0.065)	(0.046)	(0.054)	(0.054)
Ratio_seller	-0.3070*	0.2173	-0.0168	-0.1440	-0.0292	-0.0292
	(0.165)	(0.281)	(0.297)	(0.137)	(0.118)	(0.118)
Post×Ratio_seller	0.4117**	-0.4877**	-0.2550***	-0.0005	0.1399**	0.1399**
	(0.172)	(0.216)	(0.086)	(0.027)	(0.062)	(0.062)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.222	0.079	0.128	0.948	0.047	0.079
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1, with standard errors clustered at city level. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Robustness: Using Binary Ratios

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Overall Impact of AI Copilot						
Post×Ratio_tran_high	0.1030*	-0.2845***	-0.1052***	0.0778***	0.1078***	0.0963***
	(0.059)	(0.087)	(0.037)	(0.015)	(0.015)	(0.027)
Post×Ratio_tran_low	-0.1365***	-0.0127	0.0712**	0.0110	-0.0167	-0.1390***
	(0.041)	(0.052)	(0.031)	(0.010)	(0.044)	(0.050)
Panel B: Decomposed Effects and Cross-Side Effects						
Post×Ratio_buyer_high	-0.0196	-0.1378	-0.0380	0.0569***	0.0338	0.1184**
	(0.052)	(0.088)	(0.058)	(0.017)	(0.023)	(0.049)
Post×Ratio_buyer_low	-0.0302	-0.0364	0.0374	0.0016	-0.0086	-0.0824
	(0.042)	(0.052)	(0.047)	(0.013)	(0.031)	(0.056)
Post×Ratio_seller_high	0.1251**	-0.1566	-0.0429	0.0215	0.0775***	0.0281
	(0.058)	(0.110)	(0.060)	(0.020)	(0.021)	(0.039)
Post×Ratio_seller_low	-0.1205**	0.0027	0.0163	0.0076	-0.0457	-0.0515
	(0.050)	(0.058)	(0.048)	(0.011)	(0.032)	(0.071)
Observations	267,882	231,680	231,680	231,680	75,596	20,336
Adjusted R^2	0.219	0.079	0.128	0.948	0.047	0.079
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1, with binary ratios based on quintile. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Placebo Tests and Additional Outcomes

	(1) Intro_well (Buyer)	(2) Appearance (Buyer)	(3) Talk_ineff (Seller)	(4) Talk_comfort (Seller)	(5) Trust (Seller)	(6) Discount	(7) Log_volume
Panel A: Overall Impact of AI Copilot							
Post	-0.1515***	-0.1358**	0.0139**	0.0028	0.0296**	0.0024	0.1115*
	(0.032)	(0.054)	(0.006)	(0.004)	(0.015)	(0.002)	(0.062)
Ratio_tran	-0.0527	0.0429	0.0550**	-0.0400***	0.0840***	-0.0121**	0.0421
	(0.070)	(0.129)	(0.025)	(0.015)	(0.023)	(0.005)	(0.102)
Post×Ratio_tran	0.0703	0.0408	-0.0080	-0.0026	-0.0176	-0.0033	0.1094
	(0.065)	(0.081)	(0.008)	(0.006)	(0.019)	(0.003)	(0.111)
Panel B: Decomposed Effects and Cross-Side Effects							
Post	-0.1468***	-0.1390***	0.0140***	0.0023	0.0242*	0.0023	0.1191*
	(0.029)	(0.048)	(0.005)	(0.004)	(0.014)	(0.002)	(0.065)
Ratio_buyer	0.0808	-0.1907	0.0542	-0.0553***	0.1024***	-0.0043	0.0831
	(0.081)	(0.153)	(0.036)	(0.020)	(0.030)	(0.006)	(0.110)
Post×Ratio_buyer	0.0348	0.0220	-0.0051	0.0018	0.0091	-0.0077*	-0.0010
	(0.065)	(0.063)	(0.007)	(0.006)	(0.023)	(0.004)	(0.115)
Ratio_seller	-0.2121***	0.3557**	-0.0430	0.0234	-0.0428	-0.0087	0.0677
	(0.067)	(0.151)	(0.043)	(0.019)	(0.045)	(0.009)	(0.171)
Post×Ratio_seller	0.0515	0.0455	-0.0063	-0.0050	-0.0248	0.0071	0.1660
	(0.052)	(0.052)	(0.006)	(0.006)	(0.022)	(0.005)	(0.118)
Observations	75,596	75,596	20,336	20,336	20,336	231,680	58,968
Adjusted R^2	0.168	0.207	0.017	0.049	0.120	0.089	0.612
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. This table reports the estimation results of Equation 1 for placebo tests and additional outcomes. Columns (1)-(2) show placebo tests on buyer feedback labels not targeted by the AI. Columns (3)-(5) show placebo tests on seller feedback for communication skills. Column 6 shows the effect on on-site discounts, defined as the difference between the transaction price and last listed price, divided by the last listed price. Column 7 shows the effect on weekly transaction volume. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A9: Analysis of Upward and Downward Price Adjustments

	(1) Down_adjust _times	(2) Down_adjust _mag	(3) Up_adjust _times	(4) Up_adjust _mag
Panel A: Overall Impact of AI Copilot				
Post	-0.0245***	-0.0884	0.0143***	0.5693***
	(0.009)	(0.102)	(0.005)	(0.143)
Ratio_tran	-0.0301**	0.4926**	-0.0024	0.6347
	(0.013)	(0.239)	(0.005)	(0.431)
Post×Ratio_tran	0.0533***	-0.3504**	0.0106	-0.6597***
	(0.014)	(0.152)	(0.011)	(0.232)
Panel B: Decomposed Effects and Cross-Side Effects				
Post	-0.0264***	-0.1740	0.0149***	0.4694***
	(0.008)	(0.113)	(0.005)	(0.135)
Ratio_buyer	-0.0070	-0.0954	-0.0058	-0.5533*
	(0.017)	(0.273)	(0.006)	(0.322)
Post×Ratio_buyer	0.0078	0.2421	0.0126	0.0849
	(0.025)	(0.196)	(0.011)	(0.245)
Ratio_seller	-0.0512**	0.7008	0.0087	1.6021*
	(0.023)	(0.487)	(0.007)	(0.843)
Post×Ratio_seller	0.0845***	-0.6345**	-0.0026	-0.9195***
	(0.031)	(0.256)	(0.011)	(0.245)
Observations	13,581,455	1,807,297	13,581,455	491,030
Adjusted R^2	0.026	0.059	0.015	0.047
Controls	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes

Notes. This table presents the estimation results of Equation 1 results for sellers' price adjustment behaviors. The outcome variables are the number of downward price adjustments, the magnitude of downward price adjustments, the number of upward price adjustments, and the magnitude of upward price adjustments, respectively. Panel A shows the overall effects. Panel B shows the impact decomposed by the AI's buyer- and seller-facing exposure. Post × Ratio_buyer, Post × Ratio_seller are pre-period district shares of buyers/sellers transacting via the app. Controls include apartment characteristics (e.g., building age, number of rooms, etc.) as well as unemployment, population, income, and Gdp in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A10: Effects on Composition of Listed Properties

	(1) Building size	**(2)** Building age	**(3)** Floor	**(4)** Decoration	**(5)** Room
Post	0.0253***	0.0093	-0.0027	0.0073**	-0.0086***
	(0.006)	(0.008)	(0.003)	(0.003)	(0.002)
Ratio_tran	-0.1961	0.2973	-0.0774***	-0.0178	0.0698
	(0.183)	(0.310)	(0.028)	(0.087)	(0.111)
Post×Ratio_tran	-0.0122	0.0189	0.0013	-0.0093*	0.0050
	(0.008)	(0.015)	(0.005)	(0.005)	(0.003)
Observations	13,581,455	13,581,455	13,581,455	13,581,455	13,581,455
Adjusted R^2	0.646	0.429	0.012	0.088	0.635
Controls	Yes	Yes	Yes	Yes	Yes
City×Month FE	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes

Notes: This table presents the estimation results of Equation 1 results for the composition of listed properties' characteristics. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Heterogeneity in the Effect of AI Adoption by Agent Performance

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Moderation of Overall Effect						
Post×Ratio_tran	0.435***	-0.478***	-0.297***	0.103***	0.3640***	0.6465***
	(0.081)	(0.150)	(0.054)	(0.029)	(0.057)	(0.121)
Post×Ratio_tran×Performance	-0.100***	0.110***	0.0784***	-0.0018	-0.0401***	-0.1010***
	(0.009)	(0.023)	(0.027)	(0.002)	(0.015)	(0.031)
Panel B: Moderation of Decomposed Effects						
Post×Ratio_buyer	0.183*	-0.284***	-0.180***	0.116***	0.2248***	0.5409***
	(0.102)	(0.109)	(0.060)	(0.028)	(0.049)	(0.128)
Post×Ratio_buyer×Performance	-0.0849***	0.0973***	0.0627**	-0.0012	-0.0248**	-0.0830**
	(0.010)	(0.028)	(0.012)	(0.003)	(0.012)	(0.038)
Post×Ratio_seller	0.463***	-0.551**	-0.320***	-0.003	0.2388***	0.1910*
	(0.106)	(0.233)	(0.094)	(0.034)	(0.049)	(0.108)
Post×Ratio_seller×Performance	-0.0427***	0.0725**	0.0473	-0.0017	-0.0243***	-0.0219
	(0.014)	(0.031)	(0.030)	(0.004)	(0.009)	(0.032)
Observations	267,882	231,680	231,680	231,680	75,595	20,336

Notes. This table presents heterogeneity analysis by agent's pre-treatment performance score (continuous). For brevity, only the highest-order interaction terms and their constituent two-way interactions are shown. All specifications include the full set of main effects and controls as in Table 2. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A12: Heterogeneity in the Effect of AI Adoption by Agent Gender (Binary)

	(1) Sold< 3 month	**(2)** Log_days _seller	**(3)** Log_days _buyer	**(4)** Log_price _unit	**(5)** Rating (buyer)	**(6)** Rating (seller)
Panel A: Moderation of Overall Effect						
Post×Ratio_tran×Gender	0.0149	0.0590	-0.0392	-0.0115	-0.0192	-0.0645
	(0.030)	(0.064)	(0.071)	(0.012)	(0.042)	(0.102)
Panel B: Moderation of Decomposed Effects						
Post×Ratio_buyer×Gender	-0.0261	0.1006	-0.0053	-0.0158	-0.0391	-0.0034
	(0.060)	(0.079)	(0.083)	(0.017)	(0.034)	(0.097)
Post×Ratio_seller×Gender	0.0750	-0.0534	-0.0224	-0.0028	-0.0041	-0.0232
	(0.105)	(0.177)	(0.090)	(0.035)	(0.036)	(0.091)
Observations	267,882	231,680	231,680	231,680	75,595	20,336

Notes. This table presents heterogeneity analysis by agent gender (1 if male). All specifications include the full set of main effects, controls, and constituent lower-order interaction terms, though only the highest-order terms are displayed for brevity. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A13: Heterogeneity in the Effect of AI Adoption by Agent Education

	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Moderation of Overall Effect						
Post×Ratio_tran×Education	-0.0312	0.1145	-0.0954	-0.0283	0.0961	-0.0535
	(0.056)	(0.094)	(0.097)	(0.019)	(0.069)	(0.160)
Panel B: Moderation of Decomposed Effects						
Post×Ratio_buyer×Education	0.0634	0.1621	-0.2051	-0.0098	0.0624	-0.1311
	(0.065)	(0.203)	(0.158)	(0.028)	(0.076)	(0.160)
Post×Ratio_seller×Education	0.0724	-0.0410	0.1368	-0.0341	0.1005	0.1334
	(0.105)	(0.337)	(0.222)	(0.037)	(0.062)	(0.135)
Observations	267,882	231,680	231,680	231,680	75,595	20,336

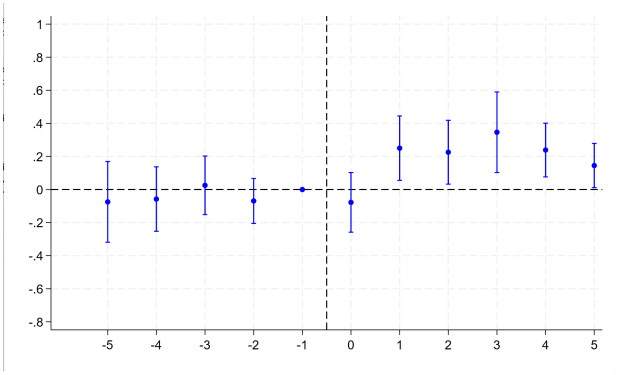
Notes. This table presents heterogeneity analysis by agent education (1 if from a prestigious university). All specifications include the full set of main effects, controls, and constituent lower-order interaction terms, though only the highest-order terms are displayed for brevity. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A14: Heterogeneity in the Effect of AI Adoption by Market Thickness

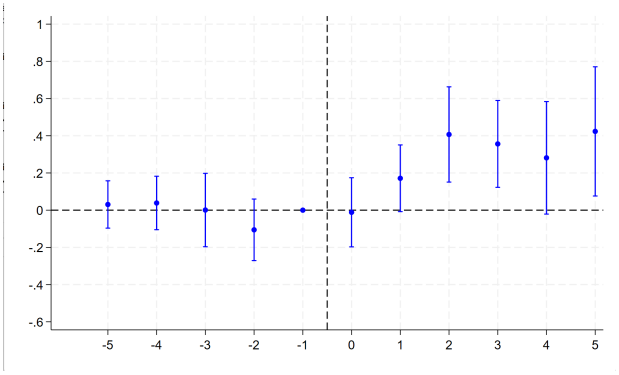
	(1) Sold< 3 month	(2) Log_days _seller	(3) Log_days _buyer	(4) Log_price _unit	(5) Rating (buyer)	(6) Rating (seller)
Panel A: Moderation of Overall Effect						
Post×Ratio_tran×Num_list	-0.0002	0.0007**	0.0003	-0.0001***	-0.0002**	-0.0007*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel B: Moderation of Decomposed Effects						
Post×Ratio_buyer×Num_list	0.0000	0.0009***	0.0002	-0.0002***	-0.0002**	-0.0006**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Post×Ratio_seller×Num_list	-0.0002	-0.0000	0.0001	0.0001	-0.0001	-0.0001
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	267,882	231,680	231,680	231,680	75,595	20,336

Notes. This table presents heterogeneity analysis by market thickness (the number of listings at the neighborhood level in the pretreatment period). For brevity, only the triple interaction terms are shown. All specifications include the full set of controls and lower-order interaction terms. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

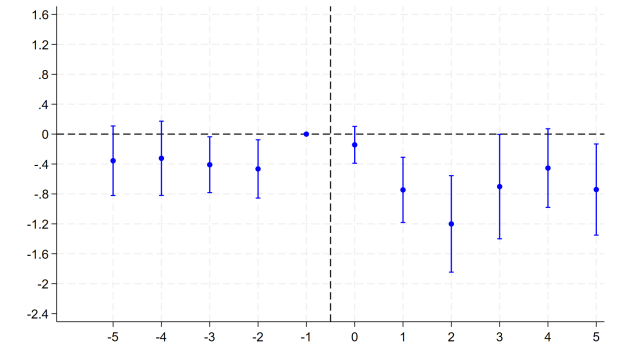
Figure A4: Event Studies for Effects on Seller Strategy



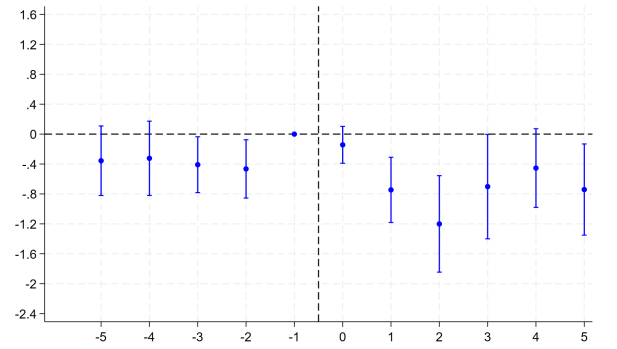
(a) Pricing_optimal (using Ratio_tran)



(b) Sellingpoint_accuracy (using Ratio_tran)



(c) Total_adjust_times (using Ratio_tran)



(d) Total_adjust_mag (using Ratio_tran)

Notes. This figure plots the bimonthly coefficients and 95% confidence intervals from event study regressions (Equation 2) regarding sellers' strategy. Pricing_optimal and Sellingpoint_accuracy indicate whether the corresponding labels—'appropriate pricing advice' and 'accurately capturing key selling points'—were selected in the seller feedback system. Total_adjust_times and Total_adjust_mag capture the total weekly frequency and magnitude of sellers' price adjustments, respectively. Standard errors are clustered at the district level.

Table A15: Impact of AI Copilot on Correlation between Days on Market for Buyers and Sellers' Price Adjustment Behaviors

	(1)	(2)	(3)	(4)
		Log_days_buyer		
Total_adjust_times×Post×Ratio_tran	0.0996			
	(0.936)			
Down_adjust_times×Post×Ratio_tran		0.6917		
		(1.119)		
Total_adjust_mag×Post×Ratio_tran			-0.1161	
			(0.087)	
Down_adjust_mag×Post×Ratio_tran				-0.1814**
				(0.077)
Observations	10,048	10,048	10,048	10,048
Adjusted R^2	0.291	0.293	0.292	0.291
City*Month FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes

Note: This table presents how AI copilot affects the correlation between district-level days on market for buyers and sellers' price adjustment behaviors. All variables are aggregated at the district level. Controls include unemployment, population, income, and Gdp in the respective district. Standard errors clustered at the district level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A16: Randomization Check

	Treatment Group	Control Group	p-Value of T-Test
Sold<3 month	0.5369	0.5267	0.28
Log_days_seller	3.7672	3.7777	0.43
Log_days_buyer	2.3362	2.3272	0.45
Log_price	9.8379	9.8473	0.21
Rating (buyer)	4.9261	4.9312	0.85
Rating (seller)	4.7232	4.7431	0.78
Building size	81.9303	82.0814	0.46
Building age	3.4686	3.4120	0.22
Floor	2.1543	2.1381	0.34
Decoration	2.4344	2.4463	0.35
Room	2.1380	2.1287	0.39