Connect or Clique: The Role of Agent Collaboration Networks in the Housing Market

Ruiqi Wang* Hefan Zheng†

October 2026

Abstract

Online real estate platforms make property listings widely accessible, yet buyers' effective search opportunities remain shaped by agents' informal collaboration networks. Using detailed records of online browsing, offline property showings, and transactions in the Beijing housing market, we recover agents' collaboration networks from their historical interactions with an unsupervised machine learning method. We find that after contacting an agent, buyers shift their online search toward locations with greater coverage by their agent's network, and network-linked properties are more likely to receive offline visits, even conditional on prior awareness and local market characteristics. We estimate a structural model of buyer awareness, offline search, and purchase that distinguishes two channels through which agent networks operate: shaping which listings buyers become aware of and steering which known listings they visit. Counterfactuals show that the two channels work in opposite directions: network-based information expands the set of listings buyers know and improves their match quality, while steering diverts their limited visits toward partners' listings and offsets much of this gain. Broadening collaboration breaks down the silos that agent cliques impose on buyers' search, substantially improving match quality with only a modest increase in search effort.

Keywords: Consumer search, agent collaboration networks, information frictions, steering, real estate brokerage

1 Introduction

Digital platforms have made product information widely accessible, yet consumers in many markets continue to rely on intermediaries to navigate their choices. Professionals such as real estate agents,

*Department of Economics, National University of Singapore.
†East China Normal University.

insurance and mortgage brokers, and financial advisers help consumers screen large choice sets, obtain information that standardized disclosures cannot convey, and coordinate the steps preceding a transaction. In carrying out these tasks, intermediaries exchange information and coordinate with one another, developing collaborative relationships with particular counterparts over time. These relationships, in turn, may direct consumer attention and shape how they search the market.

In this paper, we study how collaboration networks among intermediaries shape consumers' search and purchase outcomes in the housing market. Properties are highly heterogeneous, while online platforms provide primarily standardized listing information. Buyers therefore rely on agents to gather soft information and screen alternatives. To perform these tasks, agents draw not only on their own knowledge and the shared listing platform, but also on information and assistance obtained through collaboration with other agents. The brokerage firm we study provides a setting in which all agents have access to the same listing inventory, operate under uniform commission-sharing rules, and are encouraged to cooperate across shops. Yet collaboration remains selective: agents repeatedly work with particular partners. Moreover, these networks are predetermined from the buyers' perspective. Buyers cannot observe agents' existing collaboration relationships, and shop managers assign buyers to agents largely based on agent availability, limiting buyers' ability to select agents according to their network connections.

Our central question is how collaboration networks among agents influence the information buyers acquire and the allocation of their limited offline attention, and what these effects imply for their eventual purchases. To address this question, we first combine data on buyers' complete search process, from online browsing to offline visits and purchase, with agents' working histories. Using an unsupervised machine learning method, we identify agents' collaboration networks from their past cooperation patterns, constructing each agent's network from the preceding year's interactions. To relate these networks to buyer search, we classify a property as network-linked if its listing agent belongs to the buyer agent's recovered network.

We find that agents' collaboration networks are related to buyer search at two margins. After agent-mediated search begins, buyers reallocate online browsing toward areas where their agent's network has greater listing coverage. Within the same buyer's search, a business circle[1] with network coverage 10 percentage points higher gains 1.02 percentage points more of her browsing share. The geographic composition of the information buyers receive thus shifts toward the reach of their agent's network. Furthermore, network-linked listings receive additional offline attention even when buyers knew about them before contacting an agent. Among these previously browsed listings, network-linked properties are 3.96 percentage points more likely to receive an offline visit than other

[1]A business circle is a geographically defined cluster of residential communities with a sufficiently large properties to support the firm's local operations. Each business circle contains multiple brokerage shops, each of which is assigned an exclusive service area and staffed by a group of agents.

listings in the same business circle. These patterns tie agents' networks to both the information buyers receive and the listings they choose to inspect, but they do not reveal how either margin affects the buyers' ultimately purchase.

We therefore develop a structural model of buyer awareness, offline search and purchase that incorporates both roles of the agent's network. In the awareness stage, a buyer learns about available properties through online browsing and information provided by her agent, and network linkage can shift the probability that a listing enters her awareness, which we refer to as the information channel. The buyer then searches among the properties she is aware of. To learn her match value with a property, she must visit it in person, which takes time and effort, so she decides how many properties to visit by weighing the expected gain against these costs. Which properties she visits, however, also depends on her agent's recommendations and showing coordination, which may reflect relational incentives toward network partners. Through this steering channel, network linkage can shift the weight a listing receives in determining which properties the buyer actually visits, in a way she does not anticipate. Both channels ultimately affect which property she purchases from the set she has visited. We estimate the awareness equation by maximum likelihood and use search and purchase behavior to estimate preferences, showing priority, and search costs by generalized method of moments.

The estimates indicate that collaboration networks play a substantial role in both stages of buyer search. Network linkage significantly increases the likelihood that a property enters a buyer's awareness. It continues to matter once a listing is known: holding property attractiveness fixed, a network-linked listing receives a showing weight 8.72 times that of an otherwise comparable non-linked listing. Shutting down each channel in turn while keeping the recovered networks fixed, we find that the two channels have opposing effects of comparable magnitude. Without the information channel, the match quality of the purchased property falls by 2.50 percent. Buyers save on visit cost but lose more from a narrower set of alternatives. Removing steering while preserving information raises the match quality by 2.08 percent at unchanged search depth and awareness, as visits shift from partners' listings toward more attractive listings buyers already know. Agent networks connect buyers to a broader set of listings, but they also act as cliques that divert buyers' limited visits.

We further examine how alternative collaboration structures affect buyer outcomes. Confining networks to agents within the same shop reduces expected awareness size by 4.31 percent and lowers match quality. Expanding networks to all agents within a district, by contrast, substantially enlarges buyers' awareness sets and raises match quality by 21.50 percent. Buyers search more under this broader cooperation, with expected search depth rising by 5.68 percent, but the expected search welfare, which nets out search costs, increases by 7.98 percent. A common digital inventory therefore does not by itself equalize buyers' search opportunities. Platforms can improve buyer

benefits by collaboration and information sharing, while underscoring the importance of limiting network-based steering. These results point to two principles for platforms that promote agent cooperation: encourage collaboration and information sharing beyond selective partnerships, and limit the influence of network ties on which properties buyers visit. In practice, platforms could loosen the boundaries between agent networks by using commission splits to reward cooperation beyond an agent's usual pattern and by offering a shared calendar through which any agent can book showings directly. Offline search recommendations from the platform would also give her an independent reference for assessing her agent's showing suggestions.

Related literature. This paper contributes to three strands of literature. First, we add to research on relationships among real estate agents. Existing studies examine how professional networks form and influence agent performance and transaction outcomes (Han and Hong, 2016; Smith et al., 2019; Lopez, 2021; Xie, 2022; Scofield and Xie, 2023; Deng and Zhang, 2025; Shen and Sun, 2025; Fang et al., 2026). We extend this literature by tracing how agents' collaboration relationships affect the consumers they serve, from the properties buyers become aware of to those they visit and eventually purchase. By distinguishing information provision from network-based steering, we quantify how the same relationships can both improve and distort buyer–property matches. Second, we contribute to research on the role of intermediaries in shaping consumer decisions. Existing work documents the value of real estate intermediation (Hendel et al., 2009; Aiello et al., 2026) and shows that agents' incentives and recommendations can redirect consumer choices (Barwick et al., 2017; Li et al., 2023; Barry et al., 2024; Zhang et al., 2026). Beyond housing, platforms shape consideration through prominence (Lee and Musolff, 2021), and broker networks govern how information diffuses in financial markets (Di Maggio et al., 2019). In these studies, the influence of intermediaries stems from formal compensation or platform incentives. We identify a distinct source arising from agents' collaboration networks and examine how it shapes both the information buyers acquire and the allocation of their limited offline attention. Third, we contribute to the empirical literature on consumer search and limited consideration (De los Santos et al., 2012; Honka, 2014; Murry and Zhou, 2020; Moraga-González et al., 2023), including housing search(Genesove and Han, 2012; Fan et al., 2025). Building on the awareness, consideration and choice structure of Honka et al. (2017) and models of endogenous consideration through simultaneous search Hortaçsu and Syverson (2004), we incorporate an intermediary who influences both awareness and the composition of the consideration set. Because buyers do not observe network-based steering, their perceived selection process differs from the actual one, allowing agents to influence which properties buyers evaluate without necessarily changing how extensively they search.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting,

data and constructs agents' collaboration networks. Section 3 presents reduced-form evidence on network-mediated search. Section 4 develops the model, and Section 5 discusses identification and estimation. Section 6 reports the estimates, Section 7 presents the counterfactual analysis, and Section 8 concludes.

2 Institutional Setting, Data, and Agent Networks

2.1 Institutional Setting

We study a leading real estate brokerage firm with operations in 28 cities across China. The firm organizes its agents into more than 8,000 local brokerage shops, which serve as its primary offline units for customer reception, daily operations, and local market coverage. It also maintains proprietary online platforms through which customers can browse property listings, obtain housing information, and locate agents. This combination of offline shops and online platforms forms an integrated system that supports customers throughout the entire process of buying or selling a property.

Customers must enter into a formal agency agreement with the firm before receiving personalized brokerage services. Throughout the whole process, a customer is typically represented by a single agent. A property enters the platform after the seller contracts with a *listing agent*, who verifies and maintains the listing information, communicates with prospective buyers' agents, and coordinates access to the unit. Buyers can initially browse listings online without engaging an agent. Once a buyer decides to begin a formal offline search, she needs to enter into an agency relationship with the firm and work with a *buyer agent*, who guides her through the remainder of the search and transaction process.

The platform provides buyers with standardized listing information, such as price, location, size, and basic property attributes. However, the properties are highly heterogeneous. A posted listing does not fully convey features such as a unit's interior condition, noise exposure, natural light, or the condition of its residential compound and surroundings. Soft information from the seller side, such as the urgency of the sale and the practical scope for negotiation, is rarely disclosed on the platform. Moreover, buyers commonly search across multiple business circles as they trade off price, location, and other characteristics, and sellers similarly seek the broadest possible exposure for their listings. Serving these needs requires an agent to draw on information and resources that extend beyond her own shop.

The firm provides a formal structure for this kind of cross-agent and cross-shop cooperation through its Agent Cooperation Network (ACN). Under the ACN model, agents from different shops can perform distinct roles within the same transaction and receive corresponding shares

of the commission.[2] Although these role-based arrangements are designed to promote resource sharing across agents and shops, the repeated interactions they enable can also give rise to persistent and selective patterns of cooperation. The firm does not assign agents to fixed collaboration groups; instead, these patterns accumulate over time through agents' prior working interactions. In Section 2.3, we use the resulting patterns of cooperation to construct agents' *collaboration networks*. Because networks differ in their composition and reach, agents associated with different networks may draw on different sources of information about properties, sellers, and local market conditions despite sharing the same formal listing inventory. Buyer agents can use this information to supplement platform listings when recommending properties, arranging visits, and assisting with negotiation and transaction procedures.

2.2 Data

2.2.1 Data Description

This study focuses on the housing market in Beijing, where the firm holds the largest market share among real estate brokerage firms. The dataset covers consumers registered with the firm who were actively searching between January 1, 2023 and December 31, 2023. For each of these consumers, we observe their complete online browsing records, along with their offline property visit and purchase decisions. This is one of the most novel features of our data: we can trace a buyer's entire search process, from initial information gathering through her offline visits to an eventual purchase decision. By aligning this search history with when a buyer first engages an agent, we can separately examine how agent involvement shapes the information she acquires and the decisions she subsequently makes.

The dataset also covers all listings posted on the platform, together with property characteristics. For each listed property, we observe its list price, address, property size, number of bedrooms, number of living rooms, number of bathrooms, orientation, parking availability and whether the property is located in a school district. Our sample consists of 286,817 properties and 485,817 active buyers.

A further distinctive feature of our data is that we observe agents' complete working records: the customers each agent represents, the identities of the agents she interacts with during property visits, negotiations, and transactions. These detailed interaction histories allow us to recover recurring patterns of collaboration and construct agents' underlying collaboration networks. Over the sample period, we track the complete work histories of 19,111 real estate agents operating across Beijing's

[2]The ACN model divides the transaction process into specialized tasks that can be performed by different agents or shops. Seller-side activities include sourcing sellers, maintaining property listings, and negotiating commissions, while buyer-side activities include acquiring buyers, matching buyers to properties, and facilitating transactions. An agent may perform multiple roles within a transaction.

six urban districts.

Our listing and agent records cover January 1, 2022 through December 31, 2023, extending one year before the buyer sample period. We extend this window for two reasons. First, it ensures that our data cover the search records of all buyers throughout the sample period. Second, the longer history of agent activities allows us to construct collaboration networks using relationships established before the focal buyer begins her search. The resulting network measure is therefore predetermined relative to the buyer's search process and is not mechanically affected by her subsequent property visits or purchase. See Appendix A.1 for details on sample construction.

2.2.2 Summary Statistics

Properties. Table 1 reports summary statistics for the properties in our sample. The average property has a size of 84.5 square meters and is listed at 5.62 million CNY, equivalent to 49,866 CNY per square meter, the most commonly referenced price metric in the Chinese housing market. Residential properties in Beijing are predominantly apartments. The modal unit in our sample has two bedrooms, one living room and one bathroom. Approximately 62.8% of the properties are south-facing and 39.6% are located within designated school districts.

Of the 286,817 properties listed during our sample period, 16.1% are sold. Among properties that sell, the average time on market is 13 weeks, and the average transaction price is 5.17 million CNY, or 48,698 CNY per square meter.

Table 1. Summary Statistics of Property Characteristics

	Mean	SD
List price (million CNY)	5.622	3.280
List price per m^2 (CNY)	49,866	26,399
Property size (m^2)	84.526	31.598
Bedrooms	2.120	0.687
Living rooms	1.159	0.366
Bathrooms	1.018	0.296
South facing	0.628	0.483
Is school district house	0.396	0.489
Indicator of being sold	0.161	0.368
Weeks on market	13.064	16.187
Transaction price (million CNY)	5.173	2.874
Transaction price per m^2 (CNY)	48,698	26,084

Consumer Search and Purchase Behavior. In our sample, 485,817 buyers register with the firm and begin a search during the sample period. For each buyer, we observe her complete online

browsing history, all of her offline property visits, and the property she purchases. Specifically, we are able to track the complete timeline of a buyer's search process.

As discussed in Section 2.1, a buyer must enter into an agency contract with the firm before beginning a formal offline search, and all offline property visits in our data are consequently conducted with an agent. Figure 1 illustrates how we use this fact to reconstruct a buyer's search process. We use the date of the buyer i's first offline visit as the observable onset of agent-mediated search and denote it by τ_i. We denote by B_i the complete set of properties the buyer i browses online during her observed search spell. We refer to the properties she browses before τ_i as B_{i0}, and to the properties she browses after τ_i as $B_i^{\text{new}} = B_i \backslash B_{i0}$. The buyer then combines information acquired through online browsing with recommendations from her agent to decide which properties to visit in person. We denote this set of offline visits by N_i. Among the properties she visits in person, the buyer ultimately selects one to purchase.

Figure 1. Buyer Search Timeline

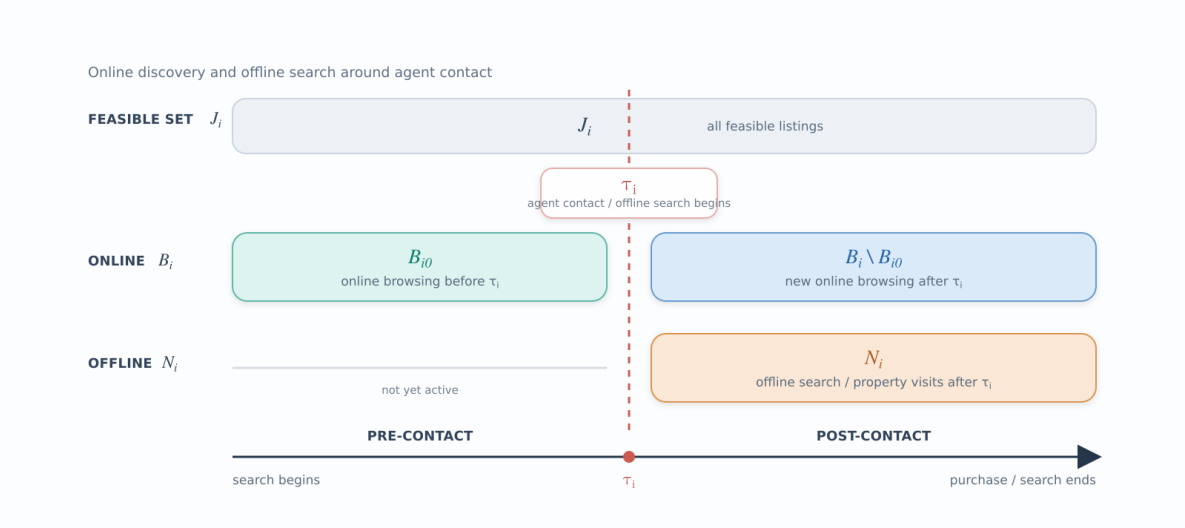


Table 2 reports summary statistics on buyer search behavior. On average, a buyer's search lasts 270 Days from her first online browsing activity to her eventual purchase. Of this period, 209 days occur before the observable onset of agent-mediated search τ_i, and 61 days occur thereafter. Buyers browse 329 properties online over the course of their search on average. 223 are viewed during the pre-contact stage, and 106 are viewed after the buyer begins her offline search with the assistance of an agent. In terms of offline search, the average buyer visits 13.3 properties in person before making a purchase. Most buyers' searches are relatively concentrated at the district level. 92.6% of buyers conduct all their online browsing and offline visits within a single district. Within that district, the average buyer searches across five business circles.

Table 2. Buyer Search Behavior

	Mean	SD	Median
Panel A: Search Duration (weeks)			
Total search duration	38.548	24.612	38.714
Pre-contact duration	29.828	24.025	27.071
Post-contact duration	8.720	13.237	6.429
Panel B: Properties Browsed Online			
Total (B_i)	329.150	213.624	276
Pre-contact (B_{i0})	223.211	184.950	150
Post-contact ($B_i \setminus B_{i0}$)	105.939	136.817	94
Panel C: Offline Visits			
Number of listings visited offline (N_i)	13.264	12.871	9

Agents. Table 3 reports summary statistics on agent activity and direct interaction patterns. Panel A shows that the average agent represents 52 clients, is involved in 178 property visits, and closes 2 transactions per year.

Panel B examines the extent and structure of agents' direct interactions. On average, 56.88% of an agent's property visits and 43.71% of her closed transactions involve an agent from a different shop, indicating that direct cooperation routinely extends beyond an agent's own shop. Although the average agent interacts directly with 174 distinct colleagues over the sample period, her top five most frequent collaborators account for 30.54% of her total interactions. This concentration indicates that agents' direct interactions are not evenly spread across all counterparts. Instead, they cluster around a group of partners, suggesting that agents' cooperation follows a structured pattern rather than occurring at random.

2.3 Agent Network Construction

An agent's collaboration network is inherently latent. No party in this market, including the firm itself, explicitly reports the informal professional relationships that determine who an agent trusts, relies on, or systematically turns to for information. Our administrative data capture the universe of instances in which agents jointly participate in property visits, negotiations, or transactions over the sample period. The granularity and coverage of our records allow us to infer this latent group structure using an unsupervised machine-learning method for community detection. Rather than equating a network with any single pairwise tie or an institutional proxy such as shop affiliation, this approach identifies collaboration networks from the overall pattern and intensity of interactions

Table 3. Agent Activity and Interaction Patterns

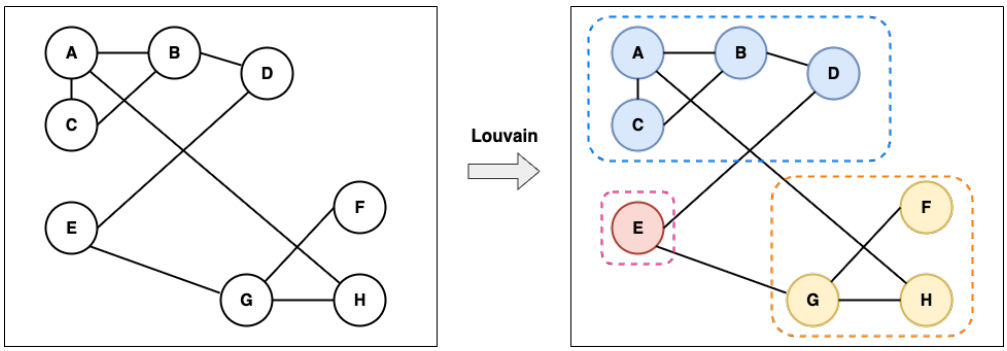
	Mean	SD	Median
Panel A: Agent Activity			
Number of clients represented	52.169	41.945	43
Number of property visits involved	177.988	119.787	159
Number of transactions closed	1.939	2.037	1
Panel B: Interaction patterns			
Share of visits involving a different-shop agent (%)	56.882	17.083	57.820
Share of transactions involving a different-shop agent (%)	43.712	21.113	41.007
Number of distinct agents interacted with	174.211	104.667	171
Share of interaction with top five collaborators (%)	30.536	17.114	26.223

Notes: The unit of observation is an agent-year. An interaction refers to a property visit or transaction in which two agents jointly participate.

among agents.

We formalize agents' interactions as an undirected, weighted graph $G = (V, E)$. Each node $v \in V$ represents a real estate agent. An undirected edge $(v_i, v_j) \in E$ connects two agents if they co-participate in at least one property visit, and its weight reflects the frequency of such co-visits between the two agents over the relevant period. We apply the Louvain algorithm to this graph to partition agents into communities that maximize within-group interaction intensity relative to what would be expected under random assignment. Each resulting community constitutes an agent's ***collaboration network***: agents assigned to the same community are considered part of the same network, regardless of whether they have interacted directly.

Figure 2. Agent Weighted Network Graph



Our approach captures collaboration networks that reflect agents' realized patterns of cooperation, rather than imposing definitions based on administrative affiliation or isolated pairwise ties. Two such simpler alternatives illustrate the tradeoff. Prior research has shown that relationships among agents within the same brokerage shop can shape agent behavior and transaction outcomes, making shop

membership a natural benchmark for defining professional networks. In our setting, however, shop membership provides an incomplete representation of agents' working relationships. A substantial share of agent interactions occurs across shops, while agents affiliated with the same shop may exhibit different patterns of collaboration and draw on distinct sources of information. A second natural benchmark is an agent's ego network, the set of agents with whom she has co-participated in at least one direct interaction. Although this direct-contact definition is commonly used in network analysis, it can give disproportionate weight to incidental interactions while excluding potentially important indirect relationships sustained through common collaborators. By aggregating information across the entire interaction graph rather than relying on any single tie, our community-based definition instead allows two agents to share a network without a direct edge and permits a direct edge to connect agents assigned to different networks. Indeed, we find that referral behavior is stronger between agents who belong to the same network but are not directly connected than between directly connected agents assigned to different networks (Appendix A.2). This pattern suggests that our network structure better captures the channels through which agents actually acquire, transmit, and mobilize information.

We construct a separate collaboration network for each year in our data, using that year's interaction graph. For a buyer who first meets an agent in year t, we assign her the agent's network as identified in year $t-1$. Thus, the network is constructed entirely from interactions that occurred before the buyer's search began, it predates the buyer–agent relationship and captures the effect of predetermined agent relationships on the buyer, rather than any network structure that could be mechanically shaped by the buyer's own search or purchase. Table 4 summarizes the resulting collaboration networks. Networks vary substantially in size and exclusivity, with the median network spanning 7 shops. Appendix A.3 further examines the factors associated with the formation of these collaboration networks.

Table 4. Collaboration Network Characteristics

	Mean	SD	Median
Network size (number of agents)	57.670	35.503	52.000
Share of internal visits	0.474	0.160	0.446
Share of internal transactions	0.147	0.103	0.122
Number of distinct shops affiliated	7.702	4.545	7.000
Share of female agents (%)	38.539	11.152	37.848

Notes: Statistics are computed at the network-year level.

3 Empirical Evidence on Network-Mediated Search

This section documents how buyer search changes once an agent becomes involved and how these changes relate to the reach of the agent's network. Our linked online and offline records allow us to trace how listings enter buyers' awareness and advance through the search process toward offline visits and purchase. We use these records to provide evidence on network-mediated search and to motivate the structural framework developed in Section 4.

3.1 Multi-Stage Search around Agent Contact

Searching for a home involves a sequence of increasingly costly decisions. Buyers first acquire preliminary information about available listings and become aware of a subset of the market. They then decide which of these listings merit closer consideration through an offline visit, which requires time, coordination, and travel. After inspecting a smaller set of properties, they finally decide which property to purchase. We treat online browsing as an observable trace of the awareness stage[3] and offline visits as the buyer's consideration set.

The observed search behavior in our data is consistent with these stages being empirically distinct. We find that online and offline search unfold over time rather than occurring as a single simultaneous decision. 93.59 percent of buyers browse online before their first offline visit, and among these buyers, online browsing begins on average 209 days before offline search. The scale of online search is only weakly related to the scale of offline search. The correlation between the number of listings a buyer browses online and the number she visits offline is 0.179. Moreover, buyers typically visit multiple properties before purchasing at most one. These patterns indicate that awareness, offline consideration, and purchase arise from distinct decisions in the housing search process, and collapsing them into a single choice stage would misrepresent how buyers search. We therefore model them as separate stages in Section 4.

We next examine how buyers' search evolves once agent-mediated search begins. A natural expectation is that the search narrows as it progresses. If buyers learn about the market and refine their preferences and budgets, they may eliminate unsuitable alternatives and concentrate on a smaller range of prices, sizes, and locations. Buyers in our sample browse for 209 days on average before their first offline visit, so the onset of agent-mediated search follows a nontrivial period of independent search. Nevertheless, buyers' online browsing frequently expands or is substantially

[3]Three features of our setting make online browsing a close measure of awareness. First, the firm is the largest brokerage in Beijing, and by our sample period its brokerage process was fully digitized: every listing enters the online platform once the seller contracts with a listing agent, and buyers must log into a registered account to view complete listing information. Second, listings that agents share with buyers through the app are recorded in the browsing data once the buyer opens them, so awareness generated by the agent also appears in our records. Third, 92.71 percent of the listings buyers visit offline also appear in their online browsing records.

reorganized after agent-mediated search begins. Buyers become aware of listings outside their previously revealed price and size ranges. 6.76 percent of the listings they newly browse fall outside the price range revealed before, most of them above its upper bound, and 12.96 percent fall outside their earlier size range. Buyers also redirect their attention across geographic areas. After contracting with an agent, approximately 30 percent of buyers expand their search to a broader set of locations, while another 45 percent reallocate search attention across locations without expanding their geographic scope.

This reorganization in observed awareness is not mirrored in buyers' offline consideration. The listings buyers visit offline remain similar in price and location to those they browsed before agent contact. This contrast suggests that the post-contact change in awareness does not primarily reflect a broad shift in buyers' budgets or underlying preferences. Buyers encounter a broader and differently composed set of listings, but their costly offline choices remain anchored to the preferences revealed before agent contact. This divergence is further evidence that awareness and consideration should be treated as distinct stages of search. It also indicates that buyers' own learning alone does not account for the change in their awareness, since learning that changed what buyers want would also change the listings they choose to visit. The timing makes agent involvement a natural candidate mechanism, as buyer agents may influence which listings enter awareness through the information they can access and transmit. The next two subsections examine whether the listings entering buyers' awareness and their subsequent offline visit decisions are systematically aligned with the predetermined reach of the buyer agent's collaboration network.

3.2 Buyer Awareness Shifts toward the Agent's Network

Listings entering buyers' observed awareness after agent contact are more likely to belong to the buyer agent's network. Let L_{ij} equal one if the listing agent of property j belongs to the collaboration network of buyer i's agent, in which case we refer to j as network-linked for buyer i. The network is identified from agent interactions in the year before the buyer's search begins (Section 2.3), so L_{ij} does not depend on the buyer's own record. We compare the incidence of network-linked listings between B_{i0}, the set of listings buyer i browses before τ_i and B_i^{new}, the set she browses after τ_i. On average across buyers, the share of network-linked listings rises from 20.15 percent in B_{i0} to 28.96 percent in B_i^{new}. To make this comparison within the same buyer, we estimate

$$L_{ij} = \alpha_i + \beta \text{PostNew}_{ij} + \gamma_1 \text{PriceDistance}_{ij} + \gamma_2 \text{PreArea}_{ij} + \epsilon_{ij} \tag{1}$$

where PostNew_{ij} equals one if $j \in B_i^{\text{new}}$ and zero if $j \in B_{i0}$, α_i is a buyer fixed effect, $\text{PriceDistance}_{ij}$ is the distance between the listing's price and the median list price the buyer revealed before τ_i, and

PreArea_{ij} indicates whether j is located in business circles the buyer browsed before τ_i.

Without controls $X_{ij} = (\text{PriceDistance}_{ij}, \text{PreArea}_{ij})$, listings entering a buyer's awareness after agent contact are 3.54 percentage points more likely to be network-linked than listings she browsed before ($\hat{\beta} = 0.0354$; Table 5, column 1), about 18 percent above the pre-contact share. This increase could arise mechanically if network-linked listings cluster near buyers' earlier price levels and in their primary districts and buyers narrow their search toward them. Controlling for $X_{ij} = (\text{PriceDistance}_{ij}, \text{PreArea}_{ij})$, Column (2) in Table 5 shows that $\hat{\beta}$ raises to 6.62 percentage points. So the shift does not reflect buyers concentrating their search around their earlier price levels and main districts. And when we replace buyer fixed effects with buyer × business circle fixed effects and compare listings only within the same buyer and business circle, the estimate $\hat{\beta}$ falls to 0.14 percentage points (Table 5, column 3). The sharp attenuation suggests that post-contact network alignment operates primarily through changes in the geographic composition of search, with at most a small additional shift toward network listings within a fixed local market.

Table 5. Network Alignment of Buyer Awareness after Agent Contact

	(1)	(2)	(3)	(4)
PostNew_{ij}	0.0354***	0.0662***	0.0014	0.0015
	(0.0029)	(0.0032)	(0.0014)	(0.0014)
$\text{PriceDistance}_{ij}$		-0.0196***		-0.0020
		(0.0026)		(0.0017)
PreArea_{ij}		0.1675***		Absorbed
		(0.0064)		
Buyer FE	Yes	Yes	No	No
Buyer × Business circle FE	No	No	Yes	Yes
Observations	6,676,812	6,676,812	6,676,812	6,676,812

Notes: Standard errors clustered at the level of the buyer's agent are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Consistent with this, buyers shift their attention toward business circles where the agent's network has stronger coverage. Let $\text{NetworkCoverage}_{ig}$ denote the share of listings in business circle g listed by agents in the network of buyer i's agent $a(i)$, constructed from the same predetermined network. For each buyer and each business circle she browses before or after τ_i, let PreShare_{ig} and PostShare_{ig} denote the circle's shares of her browsing before and after τ_i. We estimate

$$\Delta\text{SearchShare}_{ig} = \alpha_i + \delta\text{NetworkCoverage}_{ig} + \epsilon_{ig}, \tag{2}$$

$$\Delta\text{SearchShare}_{ig} = \text{PostShare}_{ig} - \text{PreShare}_{ig}, \tag{3}$$

where the buyer fixed effect restricts the comparison to business circles within the same buyer's search. We estimate $\hat{\delta} = 0.1022$, with a 95% confidence interval [0.0811, 0.1233]. For a given buyer, a business circle with network coverage that is 10 percentage points higher experiences an increase in browsing share that is 1.02 percentage points larger after contact. Thus, the geographic reorganization documented in Section 3.1 is systematically aligned with the reach of the buyer agent's collaboration network. Buyers' information search shifts toward business circles in which their agents' collaboration networks have greater coverage.

Taken together, these results show that the listings entering buyers' awareness after agent contact shift toward the agent's collaboration network, mainly through the business circles buyers browse. We interpret this pattern as agents gathering information about listings through their network connections and passing it on as buyers form their awareness of the market. Through this channel, the agent's network may shape the information set from which buyers subsequently choose which listings to visit offline. The browsing data do not distinguish listings buyers find on their own from those shared by their agents, and the evidence does not identify causal effects of the network. Buyers exploring on their own, however, is unlikely to explain this alignment, since the network formed before their search began and is not observable to them.[4] Motivated by these evidence, the model in Section 4 allows the agent's network to affect which listings enter buyers' awareness.

3.3 Network-Linked Listings Advance to Offline Visits

Section 3.2 shows that the listings entering buyers' awareness are aligned with the agent's network. We now ask whether the agent's network also shapes which of the listings buyers already know advance to offline visits. We restrict attention to listings in B_{i0}, which entered the buyer's awareness before the onset of agent-mediated search. Listings in B_i^{new} are initially excluded because the composition of this awareness set may itself have been shaped by the agent's network, as documented in Section 3.2. This restriction allows us to examine the relationship between network affiliation and offline visits conditional on awareness formed before agent contact, without conflating it with the network's role in shaping awareness. In the structural model in Section 4, we still model that buyers make their offline visit decisions from their full awareness sets.

For each buyer–listing pair, define $\text{OfflineVisit}_{ij} = \mathbb{1}(j \in N_i)$ as an indicator equal to one if buyer i subsequently visits listing j offline. Network-linked listings are substantially more likely

[4]Buyers have limited ability to select agents based on their network reach. Agent profiles on the platform display only the five residential compounds in which an agent most frequently conducts showings, which typically lie within a single business circle. They do not disclose the agent's collaboration network or its coverage across business circles. The listings represented by an agent provide similarly little information about her network, because listings are assigned largely according to agent availability when sellers contract with the firm. Moreover, although buyers may request a particular agent, buyer–agent matches are ultimately determined by shop managers, primarily based on agents' availability and existing workloads within the shop.

to advance to an offline visit. Among listings in B_{i0}, 6.16 percent of network-linked listings are subsequently visited offline, compared with 0.61 percent of listings outside the buyer's network. This difference could reflect variation across buyers in their propensity to visit properties, geographic differences in offline search, or persistent differences in the attractiveness of individual listings. We therefore estimate

$$\text{OfflineVisit}_{ij} = \alpha_{i\times Biz(j)} + \gamma_j + \beta L_{ij} + \delta \text{PriceDistance}_{ij} + \epsilon_{ij},\ j \in B_{i0} \tag{4}$$

where $\alpha_{i\times Biz(j)}$ is a buyer × business circle fixed effect and γ_j is a listing fixed effect. The specification therefore compares listings already known to the same buyer within the same business circle, while absorbing listing characteristics that are common across buyers. The estimate is $\hat{\beta} = 0.0396$, with 95% confidence interval [0.0364, 0.0429]. Conditional on prior awareness, network-linked listings are 3.96 percentage points more likely to receive a subsequent offline visit than other listings considered by the same buyer within the same business circle. The association therefore cannot be attributed to buyers' preferences for particular area or price levels, or to characteristics that make some listings attractive to buyers in general. It indicates that the agent's network is related not only to the information buyers acquire, but also to which of the listings they already know receive costly offline attention.

The contrast with Section 3.2 is informative. Within a business circle, the network is nearly unrelated to which listings enter buyers' awareness, yet it predicts which of the listings buyers already know advance to offline visits. The two sets of results thus locate the network at different margins of search: the reallocation of awareness across business circles and the selection of listings within a business circle for offline consideration. The second pattern is consistent with agents giving priority to network-linked listings when recommending or arranging visits. Several features of collaboration networks could generate such priority. Agents may hold better information about listings represented by their network partners, such as the seller's urgency or the scope for negotiation. Agents also face relational incentives. Directing buyers toward partners' listings helps sustain collaborations that generate reciprocal referrals and future joint transactions, consistent with the stronger referral ties we find within networks (Appendix A.4). Coordinating visits with familiar listing agents may also be less costly. Our evidence does not distinguish among these mechanisms, however, the network affects visit decisions through a channel distinct from awareness. This motivates a second role for the network in our model: shaping which listings in a buyer's awareness set advance to offline visits.

Overall, the evidence in this section points to the two sides of agent networks. Networks connect buyers to listings and areas beyond their independent search, and they can also operate as cliques, channeling buyers' costly offline visits toward listings held by the agent's partners. Our model captures both sides by allowing the network to affect which listings enter a buyer's awareness set

and which of these listings advance to offline visits.

4 Model

In this section, we develop a model of housing search in which the agent's collaboration network shapes buyers' awareness, consideration, and purchase outcomes. We first provide an overview of the model and outline its key assumptions. We then model the awareness stage, followed by the buyer's search and purchase decisions.

4.1 Model Overview

We develop a housing search model in which the buyer moves through three stages: awareness, consideration, and purchase. A listing can enter the buyer's awareness through browsing on the platform or information shared by her agent. We treat awareness acquisition as costless in the model, i.e., the buyer does not exert any costly effort to become aware of a property. We model revealed awareness as a function of listing characteristics, buyer characteristics, buyer-listing fit, and network affiliation, rather than as the outcome of an optimizing decision. Consideration involves costly offline visits that require time, coordination, and travel and the buyer visits properties to learn her match value with them. The buyer's consideration set is thus modeled as the outcome of a simultaneous search given her awareness set. And finally, the buyer purchases the property that yields the highest utility among those visited, without incurring further search costs.

This separation is motivated by the evidence in Section 3. Online browsing and offline visits differ in timing, scale, and composition. The three stages also play different roles in the search process. Awareness captures the acquisition of preliminary property information, consideration reflects the costly decision about which known properties to inspect in person, and purchase takes place among the properties she has visited. Moreover, the network's role differs across stages. At the awareness stage, the network is associated primarily with geographic reallocation; at the consideration stage, network-linked properties remain more likely to receive visits within the same buyer and business circle cell. Separating the stages therefore allows us to distinguish the network's role in information exposure from its role in offline consideration and to trace how both shape purchase outcomes.

Although housing search unfolds over time, we adopt a static representation to focus on the extent and composition of search over each buyer's search spell. We aggregate browsing and visits and retains the sequence from the awareness set A_i to the number and set of visits (n_i, N_i) and to the purchase decision y_i. Online browsing and offline visits may be interleaved in the observed search history, so the model's stages describe its information and decision structure rather than distinct

calendar periods. The static formulation abstracts from the timing of individual visits, sequential belief updating, and stopping decisions. But this formulation preserves what our question requires: the information buyers acquire and the search decisions through which the network shapes their purchases.

The agent's network influences both awareness and consideration through distinct activities. Collaboration with network partners facilitates the acquisition and transmission of listing information, affecting which properties become known to the buyer. We refer to this as the information channel, captured by η. Among known properties, the agent's recommendations and coordination of visits can give priority to listings represented by network partners, affecting which properties the buyer evaluates in person. We refer to this as the steering channel, captured by ρ. We assume that network links affect information exposure and selection into visits without directly changing the buyer's utility from a property. The network's effects on purchase therefore operate through the alternatives the buyer becomes aware of and subsequently evaluates.

At the consideration stage, our specification separates the buyer's choice of search depth from the realization of the visited set. The buyer bears the time and effort costs of visits and chooses search depth by weighing expected benefits against these costs. We assume that the exact set of properties she visits also depends on factors outside her control, including property access, sellers' availability and her agent's recommendations and coordination. As in Hortaçsu and Syverson (2004) and Fan et al. (2025), the buyer thus chooses the number rather than the set of properties to visit. She expects the visited properties to favor those with higher mean utility but cannot observe her agent's routine collaboration patterns or the priorities they may create in recommendations. Motivated by this information asymmetry, we assume that the buyer chooses search depth based on beliefs that exclude any network component, while the actual selection of visits allows for steering toward properties listed by the agent's network partners.

4.2 Revealed Awareness

There are N buyers indexed by $i = 1, ..., N$. We use J_i to denote the set of properties available to buyer i during her search spell. Buyer i is represented by agent $a(i)$, and her agent-mediated search begins at τ_i, the date of her first offline visit. We specify buyer i's latent awareness score of property j as

$$a_{ij} = h_{0j} + \eta L_{ij} + \boldsymbol{\gamma_1 D_i} + \boldsymbol{\gamma_2 X_{ij}} + \epsilon_{ij}, \ \forall j \in J_i \backslash A_{i0}, \tag{5}$$

where h_{0j} is a property fixed effect, L_{ij} equals one if the listing agent $a(j)$ of the property j belongs to the collaboration network of $a(i)$, $\boldsymbol{D_i}$ is a vector of observed buyer characteristics (gender, acquisition channel, etc.), $\boldsymbol{X_{ij}}$ measures buyer–listing match, and ϵ_{ij} is an error term. We measure buyer–listing match using the preferences revealed by buyer i's browsing before her first contact

with the agent at τ_i. $\boldsymbol{X_{ij}}$ includes the distance between the list price of property j and the median list price of the properties she browses before τ_i, as well as an indicator for whether property j is located in a business circle represented in her pre-contact browsing. These controls account for the possibility that network-linked properties enter the buyer's awareness set because they align more closely with her pre-contact search preferences. Let $\theta_1 = (\{h_{0j}\}, \eta, \boldsymbol{\gamma_1}, \boldsymbol{\gamma_2})$ denote the parameters of the awareness stage, of which η captures how network linkage shifts the probability that a property enters the buyer's awareness set.

We assume that buyer i is aware of property j if her awareness score for property j is larger than 0 and unaware otherwise. With ϵ_{ij} following a standard logistic distribution independently across properties, the probability that buyer i is aware of property j takes the binary logit form

$$P(a_{ij} > 0) = \frac{\exp(\bar{a}_{ij})}{1 + \exp(\bar{a}_{ij})}, \tag{6}$$

where $\bar{a}_{ij} = h_{0j} + \eta L_{ij} + \boldsymbol{\gamma}_1 \boldsymbol{D}_i + \boldsymbol{\gamma}_2 \boldsymbol{X}_{ij}$ is the deterministic component of the awareness score.

Note that we exclude A_{i0} from the model, where A_{i0} denotes the set of listings buyer i is already aware of before her agent becomes involved. Ideally, A_{i0} would include every listing buyer i knows about before τ_i, including those she learns about outside this platform, for example by passing by a property or hearing about it from acquaintances, etc. Our data, however, do not capture awareness acquired outside the firm's platform, so we measure A_{i0} by B_{i0}, the set of listings she browses online before τ_i. The gap between the two is limited in our setting. More than 80% of properties for sale in Beijing are listed on the firm's platform, and buyers must log in to view complete listing information.

4.3 Search and Purchase

Utility function. Let u_{ij} denote the utility that buyer i derives from property j. We specify the utility as

$$u_{ij} = \alpha p_j^e + \boldsymbol{\beta x_j} + \xi_j + \nu_{ij}, \tag{7}$$

where p_j^e is the expected transaction price of the property j, $\boldsymbol{x_j}$ is a vector of observed property characteristics such as size, number of bedrooms and property age. ξ_j captures property characteristics that are known to consumers but are unobservable to researchers, such as the style of interior renovation shown in listing photos. The term ν_{ij} is the match value of product j for buyer i, which she learns only after visiting the property in person. It captures buyer-specific taste for the property along dimensions that a listing cannot fully convey, such as a particular level of natural light and a particular amenity in the neighborhood. We assume that ν_{ij} is i.i.d. and follows a normal distribution with mean 0 and variance σ^2.

We rewrite the utility function as

$$u_{ij} = \delta_j + \nu_{ij}, \tag{8}$$

where $\delta_j = \alpha p_j^e + \boldsymbol{\beta x_j} + \xi_j$ is the mean utility, collecting the property j-specific terms.

Search stage. Given her awareness set A_i, buyer i undertakes offline search to learn about properties. Before visiting a property, she knows its mean utility δ_j and the distribution of match values, but not her property-specific match value ν_{ij}. An offline visit resolves this uncertainty, allowing her to assess the property's utility before making a purchase decision. Property visits entail time, coordination, and travel costs. The buyer trades off the expected gain from visiting more properties, a better eventual match, against these costs when deciding how many properties to visit. We model this decision as simultaneous search: buyer i optimally decides on the number of properties to visit (denoted by n_i) before learning property-specific match values. Conditional on this number, the visited set N_i is drawn from her awareness set. This specification allows uncertainty about which properties are ultimately visited while preserving systematic selection toward properties with higher mean utility.

The expected gain from visiting n properties depends on the buyer's beliefs about which properties she will visit. We assume that she expects selection to reflect property mean utility and does not anticipate any additional factor. Let $C^n(A_i)$ denote the collection of all subsets of A_i containing n properties. For a given number of visits n, her perceived probability of visiting a particular set $N \in C^n(A_i)$ is

$$P^{\text{belief}}(N \mid A_i, n) = \frac{\prod_{j \in N} \exp(\delta_j)}{\sum_{F \in C^n(A_i)} \prod_{j \in F} \exp(\delta_j)}. \tag{9}$$

Under this perceived selection rule, properties with higher mean utility are more likely to be visited. Specifically, the benefit of visiting n properties is that buyer i learns her match values with the visited properties, which then form her choice set at the purchase stage. Given buyer's beliefs, the expected benefit of visiting n properties is

$$EB_i(n) = \sum_{N \in C^n(A_i)} E_{(\{\nu_{ij}\}_{j \in N})}[\max_{j \in N} \delta_j + \nu_{ij}] \times P^{\text{belief}}(N \mid A_i, n), \tag{10}$$

where the term $E_{(\{\nu_{ij}\}_{j \in N})}[\max_{j \in N} \delta_j + \nu_{ij}]$ is the expected value of searching a sampled set N.

We specify the cost of searching n properties as $C_i(n) - \epsilon_{in}$, where $C_i(n)$ is the deterministic component and ϵ_{in} is a search cost shock. The shocks are i.i.d. across buyers and search intensities and follow a type-I extreme value distribution with location parameter zero and scale parameter κ.

We parameterize the deterministic component as

$$C_i(n) = c_0 n + c_1 n^2, \tag{11}$$

where c_0 captures the linear component of search costs and the quadratic term c_1 allows the marginal cost of a visit to vary with the number of visits.

When deciding on her search intensity, buyer i chooses n_i to maximize her net gain. The optimization problem at search stage is as follows:

$$\max_{0 \le n \le \overline{n}} T_i(n) = EB_i(n) - C_i(n) + \epsilon_{in}, \tag{12}$$

where $\overline{n}$ is the maximum number of visits. Buyer i chooses n_i after observing information in awareness set A_i and search cost shocks. Given that ϵ_{in} follows a type-I extreme value distribution with scale parameter κ, the probability of the search decision that buyer i visits n_i properties takes the logit form

$$P(n_i \mid A_i; \theta) = \frac{exp\{[EB_i(n_i) - C_i(n_i)]/\kappa\}}{\sum_{n=0}^{\overline{n}} exp\{[EB_i(n) - C_i(n)]/\kappa\}} \tag{13}$$

The properties the buyer actually visits, however, depend on the buyer's interaction with her agent. In the housing market, visits are arranged with the agent, whose recommendations and coordination influence which properties enter the buyer's consideration set. Network relationships may lead the agent to give priority to properties listed by network partners. We capture this influence by allowing network links to affect selection in addition to mean utility. Specifically, each property $j \in A_i$ receives a showing weight

$$\omega_{ij} = \exp(\delta_j + \rho L_{ij}), \tag{14}$$

where L_{ij} is the network linkage indicator defined in the awareness stage. Conditional on the chosen number of visits n_i, the probability of realizing a particular visited set $N_i \in C^{n_i}(A_i)$ is

$$P^{\text{actual}}(N_i \mid A_i, n_i; \theta) = \frac{\prod_{j \in N_i} \omega_{ij}}{\sum_{F \in C^{n_i}(A_i)} \prod_{j \in F} \omega_{ij}} \tag{15}$$

The parameter ρ captures the additional weight placed on network-linked properties, holding mean utility constant. The distinction between the perceived and actual selection rules reflects our assumption that the buyer does not anticipate the additional network component in her agent's recommendations. She chooses search intensity using P^{belief}, while the realized consideration set follows P^{actual}. Accordingly, ρ affects the composition of actual visits but does not enter the perceived selection probabilities used to calculate $EB_i(n)$.

Purchase stage. After visiting the properties in N_i, buyer i learns her match values with them, and all uncertainty from her perspective is resolved. She then purchases the property that yields the highest utility among those she has visited:

$$j = \arg\max_{j \in N_i} u_{ij} \tag{16}$$

Unlike the buyer, we do not observe the match values ν_{ij}. The probability that buyer i purchases property $j \in N_i$ is therefore

$$P(y_i = j \mid N_i; \theta) = P(\delta_j + \nu_{ij} \geq \max_{k \in N_i} \delta_k + \nu_{ik}). \tag{17}$$

5 Identification and Estimation

5.1 Identification

The awareness parameters are identified by which available listings newly enter buyers' awareness after agent contact. The property fixed effects capture each listing's baseline visibility and are identified by how frequently it enters awareness among buyers for whom it is available. The effect of network linkage on awareness η is identified by comparing the same listing across buyers: a given listing is network-linked for buyers whose agent's network includes its listing agent and not for other buyers. Controlling for a listing's fit with the price and location preferences a buyer reveals before contacting her agent helps account for the possibility that network-linked listings are more closely aligned with what she was already searching for. The steering parameter ρ is identified by the additional tendency of network-linked listings to be visited relative to what their mean utilities predict, among listings in the buyer's awareness set. Because the same listing is network-linked for some buyers but not for others, comparing its visit frequency across buyers who are aware of it helps distinguish steering from the listing's baseline attractiveness.

The standard deviation of match values σ is identified by comparing the mean utility of the purchased property with those of the other properties the buyer visited. A larger gap between the mean utility of the purchased property and the highest mean utility among the visited properties implies a larger standard deviation of match values.

The search cost parameters are identified by the distribution of offline visit counts. The average number of visits helps identify the linear cost parameter c_0, while the variation in visit frequencies across different numbers of visits helps identify the quadratic cost parameter c_1. The scale parameter of search cost shocks κ is identified by how strongly the number of visits respond to expected

gains from additional visits. These gains vary across buyers because of differences in the size and composition of their awareness sets.

5.2 Estimation

We divide the model parameters to be estimated into three groups: the awareness parameters $\theta_1 = (\{h_{0j}\}, \eta, \boldsymbol{\gamma_1}, \boldsymbol{\gamma_2})$; the preference and steering parameters $\theta_2 = (\alpha, \boldsymbol{\beta}, \rho, \sigma)$ and the search cost parameters $\theta_3 = (c_0, c_1, \kappa)$. The awareness stage shares no parameters or unobservables with the search and purchase stages, so we estimate θ_1 separately.

5.2.1 Revealed Awareness

The outcome at this stage is whether each available listing enters the buyer's awareness after agent contact. Given the logit form in equation (6), the likelihood of buyer i's revealed awareness set A_i is

$$P_{iA_i} = \prod_{j}^{J_i} P(a_{ij} > 0)^{\phi_{ij}} (1 - P(a_{ij} > 0))^{1-\phi_{ij}}, \tag{18}$$

in which ϕ_{ij} equals 1 if buyer i is aware of property $j \in A_i$ and equals 0 otherwise. We estimate θ_1 by maximum likelihood. The search and purchase stages condition on the observed awareness sets, so the awareness estimates do not enter their estimation.

5.2.2 Search and Purchase Given Revealed Awareness

We estimate (θ_2, θ_3) sequentially, using moment conditions constructed from buyers' visit sets, purchase choices, and number of visits.

Unobserved property attributes ξ_j enter mean utility and are likely to be correlated with price. Following Petrin and Train (2010), we construct a control function for price and regress the expected price on the price instrument Z_j, the observed characteristics x_j, and location fixed effects $\boldsymbol{F_j}$:

$$p_j^e = \boldsymbol{\pi_z Z_j} + \boldsymbol{\pi_x x_j} + \boldsymbol{F_j} + \zeta_j. \tag{19}$$

The price instruments $\boldsymbol{Z_j}$ includes the average transaction price of properties sold in the same segment[5] during the three weeks preceding property j's listing, and DIV following Gandhi and Houde (2019). Let $\widehat{\zeta_j}$ denote the resulting residual. We incorporate this price control into mean

[5]We define mutually exclusive segments based on neighborhood and property size ranges, assigning each property to a unique segment.

utility:

$$\delta_j = \alpha p_j^e + \boldsymbol{\beta x_j} + \boldsymbol{\psi F_j} + \lambda \widehat{\zeta}_j. \tag{20}$$

We first construct two blocks of moments to estimate $\tilde{\theta}_2 \equiv (\theta_2, \boldsymbol{\psi}, \lambda)$. The visit moments compare the observed composition of N_i with the composition predicted by the actual selection rule in equation (15), conditional on A_i and n_i^{obs}. The purchase moments compare the observed purchase with the choices predicted by equation (17), conditional on the observed visited set N_i. Let $m_i^V(N)$ summarize the composition of a visited set and $m_i^P(y, N_i)$ summarize the purchase outcome within the observed visited set. Specifically, we have

$$\boldsymbol{m}_i^V(N_i) = \sum_{j \in A_i} \boldsymbol{q}_{ij}^V \mathbf{1}\{j \in N_i\}, \qquad \boldsymbol{m}_i^P(y_i, N_i) = \sum_{j \in N_i} \boldsymbol{q}_{ij}^P \mathbf{1}\{y_i = j\}, \tag{21}$$

where $\boldsymbol{q}_{ij}^V$ includes observed property characteristics, expected price, the control function residual, location indicators and the network indicator L_{ij} and $\boldsymbol{q}_{ij}^P$ includes observed property characteristics, the control function residual, and L_{ij}. The model-implied counterparts are

$$\mathbb{E}_{\tilde{\theta}_2}\left[\boldsymbol{m}_i^V \mid A_i, n_i\right] = \sum_{j \in A_i} \boldsymbol{q}_{ij}^V \Pr(j \in N_i \mid A_i, n_i; \tilde{\theta}_2), \tag{22}$$

$$\mathbb{E}_{\tilde{\theta}_2}\left[\boldsymbol{m}_i^P \mid N_i\right] = \sum_{j \in N_i} \boldsymbol{q}_{ij}^P \Pr(y_i = j \mid N_i; \tilde{\theta}_2). \tag{23}$$

Stacking the two blocks across buyers gives the sample moments

$$\overline{g}_2(\theta_2, \psi, \lambda) = \frac{1}{N} \sum_{i=1}^{N} \begin{pmatrix} m_i^V(N_i) - \mathbb{E}_{\theta_2}\left[m_i^V(N) \mid A_i, n_i^{\text{obs}}\right] \\ m_i^P(y_i, N_i) - \mathbb{E}_{\theta_2}\left[m_i^P(y, N_i) \mid N_i\right] \end{pmatrix}, \tag{24}$$

and we estimate

$$(\widehat{\theta}_2, \widehat{\psi}, \widehat{\lambda}) = \arg\min_{\tilde{\theta}_2} \overline{g}_2(\tilde{\theta}_2)' W_2 \overline{g}_2(\tilde{\theta}_2). \tag{25}$$

The visit and purchase moments jointly discipline the mean utility parameters. The visit moments involving L_{ij} help distinguish network steering from selection based on mean utility. Purchase choices inform σ through their sensitivity to mean utility differences within the visited set. We estimate this parameter block by two-step GMM, using the inverse of the estimated covariance of the buyer-level moment contributions in the second step.

We next estimate the search cost parameters $\theta_3 = (c_0, c_1, \kappa)$. For each buyer, we use $\widehat{\tilde{\theta}}_2$ to compute the expected benefit $\widehat{EB}_i(n)$ for each feasible search depth. The predicted probability of

visiting n properties is therefore

$$\Pr(n \mid A_i; \theta_3, \widehat{\tilde{\theta}}_2) = \frac{\exp\left[\left(\widehat{EB}_i(n) - c_0 n - c_1 n^2\right)/\kappa\right]}{\sum_{r=0}^{\overline{n}} \exp\left[\left(\widehat{EB}_i(r) - c_0 r - c_1 r^2\right)/\kappa\right]}. \tag{26}$$

Let $m_i^n(n)$ collect moments of visit counts, including indicators for visit-count categories and their interactions with characteristics of the awareness set. We estimate θ_3 by matching their observed values to their model predictions:

$$\overline{g}_3(\theta_3; \widehat{\tilde{\theta}}_2) = \frac{1}{N} \sum_{i=1}^{N} \left[m_i^n(n_i^{\text{obs}}) - \sum_{n \in \mathcal{D}_i} \Pr(n \mid A_i; \theta_3, \widehat{\tilde{\theta}}_2) m_i^n(n) \right],$$

$$\widehat{\theta}_3 = \arg\min_{\theta_3} \overline{g}_3(\theta_3; \widehat{\tilde{\theta}}_2)' W_3 \overline{g}_3(\theta_3; \widehat{\tilde{\theta}}_2).$$

These moments capture both the distribution of visit counts and its variation across awareness sets. The distribution helps distinguish the linear and quadratic cost components, while variation in the response of search depth to predicted benefits helps identify κ.

We conduct inference using a buyer-level bootstrap, repeating the price-control estimation and re-estimating both parameter blocks in each replication. Expected search benefits are recomputed using the replication-specific estimates, thereby propagating uncertainty in $\widehat{\tilde{\theta}}_2$ into the search cost estimates.

6 Results

Table 6 reports the estimates of the awareness stage. The estimated network coefficient is significantly positive, with $\widehat{\eta} = 4.0628$. Conditional on property fixed effects, buyer characteristics, and the measures of buyer–listing fit, a listing represented by a partner of the buyer agent's network is more likely to enter the buyer's awareness after agent contact. Because the specification controls for price distance and location relative to the buyer's pre-contact preference, network-linked listings retain this advantage even when compared with listings that align equally closely with her prior search.

Network linkage shifts the awareness score 1.81 times as much as location in business circles the buyer browsed before agent contact, and it offsets a 1.04 million CNY distance from the buyer's pre-contact median price. Network ties therefore bring listings into the buyer's awareness that lie outside the locations and price range of her independent search. This mechanism helps explain the reorganization of search documented in Section 3: after agent contact, buyers shift their search

across locations and broaden the range of prices they consider. Through the information channel, the agent's collaboration network shapes what listing information flows to the buyer, widening her search toward the locations and price ranges where her agent's network partners hold listings.

The remaining coefficients show that awareness also reflects buyers' prior search patterns. The coefficient on price distance is negative ($\hat{\gamma_2} = -3.9016$), indicating that a listing is less likely to enter awareness when its price is farther from the median price of the properties browsed before agent contact. The positive coefficient on the pre-search location indicator ($\hat{\gamma_2} = 2.2392$), so listings in previously browsed business circles remain more likely to enter awareness. In addition, buyers who browsed more intensively before contact continue to acquire awareness at a higher rate after contact with an agent. Female buyers become aware of slightly more listings.

Table 6. Estimates of Awareness Parameters

Variable	Parameter	Estimate
Network linkage	η	4.0628***
Female indicator	γ_1	0.0164***
Prior search density	γ_1	0.1008***
Price distance	γ_2	-3.9016***
Pre-search location indicator	γ_2	2.2392***
Property fixed effects		yes

Table 7 reports the estimates of the search and purchase parameters. The utility estimates have intuitive signs. Buyers prefer lower expected transaction prices, larger properties, and more living rooms, while higher management fees reduce a property's attractiveness. The coefficient on the price-control residual is positive, consistent with unobserved attributes that raise a property's price also increasing its appeal to buyers. Accounting for this component helps separate price sensitivity from preferences for unobserved property quality.

The estimated linear and quadratic search cost coefficients are 0.04942 and 0.00103, respectively, implying that search costs are increasing and convex in the number of visits. For example, the incremental cost of the tenth visit is approximately 36.7 percent higher than that of the first, reflecting the growing burden of undertaking additional property visits.

The estimated network steering coefficient is positive, with $\widehat{\rho} = 2.1659$. Holding mean utility fixed, a network-linked listing receives a showing weight 8.72 times that of a non-linked listing. Network affiliation continues to shape the buyer's search after a listing has entered her awareness, influencing which known properties advance to in-person evaluation. Because buyers do not observe their agents' collaboration networks, they do not anticipate this additional network-based priority when choosing properties to visit. Steering therefore creates a gap between the selection process

buyers expect and the visits they actually undertake, changing the alternatives available for purchase and potentially affecting the quality of the eventual match.

Taken together, η and ρ show that the network operates at both margins of search. It expands the set of listings buyers know through the information channel, and it tilts the listings buyers inspect toward the agent's partners through the steering channel. Section 7 quantifies how each channel affects search depth, match quality and buyer welfare.

Table 7. Estimates of Search and Purchase Parameters

Variable	Parameter	Estimate
Panel A: Preferences and match values		
Expected transaction price (million CNY)	α	−0.01532***
Price-control residual	λ	0.04938***
Property area (m^2)	β	0.16012***
Number of living rooms	β	0.01761***
Management fees	β	−0.01615***
SD of match value	σ	1.7261***
Panel B: Search costs and network steering		
Linear search cost	c_0	0.04942***
Quadratic search cost	c_1	0.00103***
Network steering	ρ	2.1659***

7 Counterfactual Analysis

Having established how agents' collaboration networks shape buyers' awareness and offline search, we next examine the consequences for match quality and buyer welfare. Information shared through network partners may help a buyer discover suitable properties beyond her independent search. A broader awareness set may help her find a better match, but evaluating additional options is costly. Meanwhile, if the added listings are poorly suited to her, they may dilute her search by competing with better alternatives for a limited number of visits. Likewise, recommendations that favor network-linked listings may give buyers access to richer information about these properties, but they may also divert buyers' limited search effort toward partners' listings at the expense of better alternatives. We use the estimated model to quantify the net effect of each channel on search behavior and match quality.

Our counterfactual analysis proceeds in two steps. We first isolate the contribution of each channel by shutting down one at a time while keeping the observed collaboration networks unchanged (CF1, CF2). We then examine how buyers would fare under alternative architectures of agent

cooperation, replacing the recovered networks with shop-based or district-wide cooperation and redefining which listings are network-linked for each buyer (CF3, CF4).

For each scenario, we simulate 20 awareness sets, search decisions, and purchase outcomes per buyer. We average each outcome across simulation draws and buyers, and report the percent change in this mean relative to its baseline value for outcomes that capture buyers' search behavior (expected awareness and expected search depth), the perceived value of search (expected search welfare), and match outcomes (the expected purchase index). Appendix A.5 describes the simulation in detail.

7.1 Connect versus Clique

We begin by separating the network's role in connecting buyers to listings from its role in directing their offline visits. Both experiments retain the recovered collaboration networks. Table 8 reports percent changes relative to the baseline in four outcomes. Expected awareness and expected search depth describe how buyers search the housing market and the value of that search as they perceive it. Expected search welfare is the buyer's ex ante expected net value of search under her perceived showing rule. The expected purchase index is the expected utility of the purchased property, and therefore measures the match quality buyers realize.

In CF1, we shut down the information channel by setting $\eta = 0$. Network affiliation no longer increases a listing's probability of entering the buyer's awareness set, but buyers can still discover listings through the other determinants of awareness. The additional showing weight assigned to network-linked listings remains unchanged. These results indicate that the information transmitted through network partners is valuable on balance. Removing the network-based information advantage leaves buyers aware of 5.64 percent fewer listings and reduces their expected search depth by 1.82 percent. Although buyers consequently make fewer costly visits, expected search welfare falls by 1.78 percent and the expected purchase index by 2.50 percent. The fall in search welfare indicates that the loss of perceived search benefits outweighs the associated savings in search costs. Network-provided information therefore improves the perceived net value of search and helps buyers achieve better eventual matches, despite the additional search effort it induces.

In CF2, we shut down the steering channel by setting $\rho = 0$. Agents continue to share listing information with buyers through their collaboration networks, but no longer prioritize their network partners' listings when recommending properties for offline visits. The selection of listings for offline visits is driven solely by their perceived attractiveness to buyers, rather than agents' network affiliations. Actual selection probabilities therefore coincide with buyers' perceived selection probabilities.

Expected awareness, search depth, and search welfare remain unchanged in CF2. Awareness is unaffected because the information channel is preserved. Search depth and perceived search

welfare are also unaffected because buyers do not anticipate network-based showing priority when evaluating the return to additional visits. Nevertheless, the expected purchase index increases by 2.08 percent. Removing steering thus improves match outcomes through a reallocation of visits at the existing search depth. Under the estimated steering weight, agents' recommendations draw visits away from more attractive alternatives already in buyers' awareness sets and toward listings held by their network partners. This pattern is consistent with steering driven by relational incentives or lower coordination costs with familiar listing agents, either of which could give partners' listings additional showing priority beyond that explained by their expected attractiveness to buyers.

Taken together, the two channels have opposite effects of comparable magnitude. Information transmitted through network partners improves match outcomes, whereas the additional showing priority given to their listings worsens them. Each experiment measures the effect of removing one channel relative to the baseline. These effects need not add up to the effect of removing both channels, because the awareness set determines the alternatives over which steering operates.

Table 8. Effects of Network Channels and Architecture

	Connect versus Clique		Network Architecture	
	CF1 No information	CF2 No steering	CF3 Shop-only	CF4 District-wide
Expected awareness size	−5.64%	0[a]	−4.31%	+155.31%
Expected search depth	−1.82%	0[a]	−1.35%	+5.68%
Expected search welfare	−1.78%	0[a]	−1.29%	+7.98%
Expected purchase index	−2.50%	+2.08%	−0.72%	+21.50%

Notes: Entries are percentage changes in simulated means relative to the baseline. CF1 sets $\eta = 0$ and retains $\rho = \hat{\rho}$; CF2 sets $\rho = 0$ and retains $\eta = \hat{\eta}$. CF3 and CF4 replace network linkage in both the awareness and showing equations. [a]The zero changes in CF2 follow from the unchanged awareness equation and perceived returns to search. Expected search welfare measures the perceived value of search; the purchase index summarizes match outcomes.

7.2 Network Architecture

The channel experiments retain the observed collaboration networks. We now examine how buyers would fare if collaboration were organized around different boundaries. The firm's cooperation system allows agents to collaborate across shops, and the networks we recover reflect how agents use this flexibility in practice. We compare this architecture with two alternatives that bracket the observed networks in terms of reach: cooperation confined to agents within the same shop (CF3) and cooperation among all agents in a district (CF4).

Under shop-only cooperation, all agents affiliated with the same shop form one collaboration

network. This scenario corresponds to a traditional brokerage arrangement in which agents cooperate only with colleagues in their own shop. Because the median collaboration network spans 7 shops, restricting cooperation to the shop reduces the number of listings that receive the network's awareness advantage. Expected awareness falls by 4.31 percent. With smaller awareness sets, the expected benefit of additional visits declines, and expected search depth falls by 1.35 percent. Expected search welfare falls by 1.29 percent and the expected purchase index by 0.72 percent. The decline in match quality is modest because the shop-only definition also reduces the number of network-linked listings subject to steering, which lowers match quality on net. The smaller loss from steering partly offsets the loss of information. On balance, the existing cross-shop relationships provide value to buyers beyond the cooperation available within a buyer agent's own shop.

District-wide cooperation produces substantially larger changes in the opposite direction. We assign all agents in a district to a single network. This scenario corresponds to fully open cooperation, in which every agent can draw on information from every other agent in the district. Expected awareness increases by 155.31 percent, expected search depth by 5.68 percent, expected search welfare by 7.98 percent, and the expected purchase index by 21.50 percent. District-wide cooperation thus substantially expands buyers' awareness and improves purchase outcomes, while inducing only a modest increase in offline visits. Although evaluating additional properties is costly, the increase in perceived search welfare indicates that the expanded search opportunities are valuable to buyers even after accounting for search costs. The counterfactual changes both the alternatives available for evaluation and the priorities governing visits: buyers draw their visits from a much larger awareness set, while network affiliation no longer distinguishes among listings represented by agents within the district. The improvement in purchase outcomes reflects the joint effect of these changes and buyers' adjustment in search depth. These results show that broader cooperation can improve match outcomes without requiring a large increase in the number of visits.

These magnitudes should be interpreted in light of what the counterfactuals hold fixed. All counterfactuals in this section are partial equilibrium exercises on the demand side. We hold fixed the set of listings available to each buyer, list prices, buyer–agent matches, and the listing agent of each property, and we do not model how agents would adjust their effort, cooperation, or recommendation behavior in response to a change in the cooperation regime. This restriction matters most for CF3 and CF4, which change the scope of cooperation itself. CF4 further holds η at its estimated value while extending network linkage to all listings in a district, which assumes that agents can sustain the same per-listing information flow with every agent in the district as with their current network partners. The results therefore describe short-run effects on buyers, and we interpret the CF4 results as an upper bound on the gains from fully open cooperation.

7.3 Implications for Platform Design

Our results point to two principles for platforms that promote agent cooperation: broaden buyers' access to listing information, and make showing recommendations less dependent on network affiliation. The channel experiments show that the first preserves the benefits of network-based information, while the second removes the cost of steering. The architecture experiments show that the two principles are complementary. Open cooperation across a broad area, in which network affiliation no longer distinguishes among listings, achieves both and yields the largest improvement in purchase outcomes. Exclusive cooperation within a narrow area reduces information without removing steering.

These principles suggest several measures within the firm's existing cooperation system. On the information side, the platform could require listing agents to record standardized soft information, such as unit condition, noise exposure, natural light, and the seller's urgency to sell, in listing records accessible to all buyer agents. Information that now circulates within collaboration networks would then become available across them. Commission splits under the cooperation system could also reward cooperation with agents outside an agent's usual partners. Through these rules and incentives, the firm can encourage agents to form more open and wider-reaching cooperation. At the same time, the platform could introduce new tools that make showing opportunities less dependent on agents' network ties. It could publish each listing's available showing times in a shared booking calendar, through which any buyer agent can reserve a showing directly. This gives all buyer agents equal access to showing opportunities and removes the coordination advantage of listings held by familiar partners. Besides, the platform could also provide buyers with fit-based recommendations generated from their browsing histories, as an additional reference for deciding which properties to visit. When these recommendations diverge substantially from their agent's, buyers gain a signal that their agent's recommendations may reflect considerations other than fit and can adjust their search accordingly. Both measures limit the agent's control over which known listings advance to offline visits and thereby reduce steering.

8 Conclusion

Online platforms have made housing listings widely accessible, yet the set of listings an individual buyer learns about, and the way she decides which of them to pursue, remain specific to her own search. In markets where buyers rely on intermediaries, we know little about how they shape the information flow and the decisions that follow. This paper examines how intermediaries exert this influence through the informal collaboration networks among real estate agents. Using linked records of online browsing, offline showings, and transactions in the Beijing housing market, together

with agents' complete working histories, we recover agents' collaboration networks by applying an unsupervised machine learning method. These networks reveal how agents organize their working relationships within a common listing system and allow us to examine how those relationships influence buyers' effective access to the housing market.

Our empirical evidence identifies two distinct margins of network-mediated search. After agent contact, buyers redirect their online browsing toward locations where their agents' networks have greater coverage. Among listings buyers already knew before contact, network-linked properties are also more likely to receive offline visits, conditional on local market and listing characteristics. We develop and estimate a structural model of awareness, costly offline search, and purchase that captures these patterns through separate information and steering channels. The estimates indicate that network affiliation increases a listing's probability of entering the buyer's awareness set and, conditional on awareness, gives it additional showing priority beyond that explained by its expected attractiveness to the buyer. Separating these channels allows us to quantify the benefits of network-based information and the consequences of network-based showing priorities for search behavior and eventual match quality.

The counterfactual results show that the two channels have opposing effects on purchase outcomes. Removing the network-based information advantage reduces the expected purchase index, our measure of match quality, by 2.50 percent. Although buyers undertake fewer costly visits, their perceived net value of search also declines. Eliminating network-based showing priority while preserving information transmission raises the expected purchase index by 2.08 percent. Agent collaboration networks thus connect buyers to useful listings beyond their independent search, but they also act as cliques that divert limited offline visits toward listings held by the agent's partners.

The architecture of cooperation matters as well. Confining cooperation to agents within the same shop lowers match quality, indicating that existing cross-shop relationships provide value to buyers. Extending cooperation to all agents in a district raises the expected purchase index by up to 21.50 percent while increasing search depth by only 5.68 percent. These results suggest that platforms can improve buyer outcomes by broadening access to listing information and making showing recommendations less dependent on selective relationships. Standardized sharing of soft listing information, shared access to showing arrangements, and recommendations based on buyer–property fit are possible ways to pursue these objectives within an existing cooperation system.

More broadly, our findings show that a common digital inventory does not by itself equalize buyers' search opportunities. Buyers who draw on the same listings still differ substantially in the alternatives they learn about and inspect, depending on whom their agent collaborates with. The value of agent cooperation network therefore depends on both the information it transmits and how its boundaries shape access to offline evaluation, which makes the governance of these informal relationships a central question for platform design.

References

Darren Aiello, Mark Garmaise, and Taylor Nadauld. What problem do intermediaries solve? evidence from real estate markets. *The Review of Financial Studies*, 39(2):562–604, 2026.

Jordan M Barry, Will Fried, and John William Hatfield. Et tu, agent? commission-based steering in residential real estate. *Iowa L. Rev.*, 110:1473, 2024.

Panle Jia Barwick, Parag A. Pathak, and Maisy Wong. Conflicts of interest and steering in residential brokerage. *American Economic Journal: Applied Economics*, 9(3):191–222, 2017.

Babur De los Santos, Ali Hortaçsu, and Matthijs R. Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6): 2955–2980, 2012.

Guoying Deng and Xuyuan Zhang. Digital platform consolidation and offline expansion: Strategic convergence and market welfare in china's second-hand real estate market. Working paper, arXiv:2409.04326, 2025.

Marco Di Maggio, Francesco Franzoni, Amir Kermani, and Carlo Sommavilla. The relevance of broker networks for information diffusion in the stock market. *Journal of Financial Economics*, 134(2):419–446, 2019.

Ying Fan, Ziying Fan, and Yiyi Zhou. Dynamic search in a non-stationary search environment: An application to the beijing housing market. Technical report, National Bureau of Economic Research, 2025.

Lu Fang, Darren K Hayunga, and Henry J Munneke. Real estate agents' collaboration and housing outcomes. *Real Estate Economics*, 2026.

Amit Gandhi and Jean-François Houde. Measuring substitution patterns in differentiated-products industries. *NBER Working paper*, (w26375), 2019.

David Genesove and Lu Han. Search and matching in the housing market. *Journal of Urban Economics*, 72(1):31–45, 2012.

Lu Han and Seung-Hyun Hong. Understanding in-house transactions in the real estate brokerage industry. *The RAND Journal of Economics*, 47(4):1057–1086, 2016.

Igal Hendel, Aviv Nevo, and François Ortalo-Magné. The relative performance of real estate marketing platforms: Mls versus fsbomadison. com. *American Economic Review*, 99(5):1878–1898, 2009.

Elisabeth Honka. Quantifying search and switching costs in the US auto insurance industry. *RAND Journal of Economics*, 45(4):847–884, 2014.

Elisabeth Honka, Ali Hortaçsu, and Maria Ana Vitorino. Advertising, consumer awareness, and choice: Evidence from the U.S. banking industry. *RAND Journal of Economics*, 48(3):611–646, 2017.

Ali Hortaçsu and Chad Syverson. Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly journal of economics*, 119(2):403–456, 2004.

Kwok Hao Lee and Leon Musolff. Entry into two-sided markets shaped by platform-guided search. *Job Market Paper, Princeton University*, 2021.

Keyang Li, Jing Wu, Jianwei Xing, and Jubo Yan. Steering in the housing market: Incentive induced by the tax scheme. *Regional Science and Urban Economics*, 103:103953, 2023. doi: 10.1016/j.regsciurbeco.2023.103953.

Luis A. Lopez. Asymmetric information and personal affiliations in brokered housing transactions. *Real Estate Economics*, 49(2):459–492, 2021.

José Luis Moraga-González, Zsolt Sándor, and Matthijs R. Wildenbeest. Consumer search and prices in the automobile market. *Review of Economic Studies*, 90(3):1394–1440, 2023.

Charles Murry and Yiyi Zhou. Consumer search and automobile dealer colocation. *Management Science*, 66(5):1909–1934, 2020.

Amil Petrin and Kenneth Train. A control function approach to endogeneity in consumer choice models. *Journal of marketing research*, 47(1):3–13, 2010.

David Scofield and Jia Xie. Network formation and effects: Observations from us commercial real estate markets: Network formation and effects: Observations from us commercial real estate markets. *The Journal of Real Estate Finance and Economics*, 66(2):487–504, 2023.

Lily Shen and Xiaojin Sun. The more, the merrier: Agent connectedness in real estate trading networks. *Journal of Real Estate Research*, 47(2):149–169, 2025.

Patrick S Smith, Velma Zahirovic-Herbert, and Karen M Gibler. Building social capital at the expense of principals: Evidence from residential real estate agent trading networks. *Journal of real estate research*, 41(1):147–182, 2019.

Jia Xie. Importance of professional networks in trade: Evidence from real estate market. *Real Estate Economics*, 50(3):882–900, 2022.

Dayin Zhang, Panle Jia Barwick, Lu Han, and Jonathan Kroah. Market power in mortgage pricing: the role of referral lending. Technical report, National Bureau of Economic Research, 2026.