Organization of Platform Markets: Coordination, Decision Rights and Information

Krishnan S. Anand
anand@eccles.utah.edu

Nabita Penmetsa
nabita@eccles.utah.edu

David Eccles School of Business, University of Utah

Abstract

Platform markets differ significantly from traditional firms– whether operating through wholly-owned branches or tightly-controlled franchises—in their organization. Unlike traditional firms, platforms rely on decentralized supply of services by diverse, independent agents. Each agent operates with limited information and competes with other agents to maximize her own profits.

Should decision making in a platform market be centralized or decentralized, i.e., controlled by the platform or by agents? To study this question, we define the *coordination structure* (CS) of the platform market as the combination of two elements: its *decision rights structure* ('who decides what') and *information endowment* ('who knows what'). We model a platform providing services through independent, competing agents, and parameterize the degree of competition among agents. Demand uncertainty can be partly alleviated through information. We model different facets of information endowment, including local and aggregate intelligence, the ability to acquire and share intelligence, and non-transferable specific knowledge of both the platform and agents.

We find that centralized coordination of prices, enabled by centralized CS, is a powerful lever, driving performance. Indeed, centralized CS dominates decentralized CS, reversing the results of the extant literature on traditional firms. We show that the value of centralized coordination increases rapidly with the degree of agent competition—a feature less relevant to traditional firms. As the degree of agent competition increases, centralized CS outperforms decentralized CS, even when agents have significant informational advantages over the platform. Our results have important practical implications for the organization of platform markets.

DRAFT ONLY. PLEASE DO NOT CIRCULATE. COMMENTS WELCOME.

1 Introduction

Platform markets are blurring the theoretical distinctions between firms and markets. Platform business models transcend the hierarchies and contractual relationships of traditional firms and yet incorporate a visible hand in the form of centralized controls. Platforms create value by enabling the exchange of goods and services that were once considered personal and non-marketable. Airbnb, considered one of the most important start ups to have fueled the growth of platform markets, had completed more than 3 million bookings worldwide by the end of 2018, surpassing many traditional hotel chains in the number of transactions and market value. Close to 6 million drivers participated and offered rides on ride-hailing platforms Uber and Lyft in 2019, surpassing the 1.5 million workforce of the nation's largest employer, Walmart.

The role of the platform in coordinating the actions of decentralized agents is an important decision that can define how the market functions. For example, the platform can centrally control the terms of exchange between agents and users or allow agents some leeway in setting these terms. In practice, we see both kinds of organization. Ride-sharing platforms such as Uber and Lyft use centralized pricing as a strategic tool to disrupt the taxi industry and scale the platforms effectively. On the other hand, online retailing platforms such as eBay and Etsy allow decentralized pricing to attract sellers and expand product offerings. As we can observe from Table 1, platforms have adopted both centralized and decentralized pricing structures across different categories of services. For example, home rental platforms Airbnb, Homeaway and VRBO use decentralized pricing. On the other hand, leading platforms in ride-hailing services across the world like Uber and Lyft in the US, Didi in China and Grab in South East Asia use centralized pricing. Moreover, we observe that platforms even within a category do not always choose the same pricing structure (Table 1). The choice of pricing structure also seems to be driven by local market conditions. For example, a majority of the rental services platforms in the US use decentralized pricing, but it not always the case in other regions of the world. For example, OYO in India and Tujia in China use Centralized pricing. In certain categories like home services, there does not seem to be any agreement over the pricing structure, with platforms like TaskRabbit and Thumbtack experimenting with different pricing structures.

Category	Decentralized Pricing	Centralized Pricing
Home Rental Services	Airbnb, Homeaway, VRBO	OYO, Tujia
Automobile Rental Services	Turo, Getaround	Coop
Ride Hailing Services	Xooox	Uber, Lyft
Carpooling Services	BlaBlaCar	sRide
Home Services	TaskRabbit	Handy

Table 1: Examples of platforms with alternative pricing structures

In this paper, we develop a theoretical framework for studying the trade-offs between centralization and decentralization in platform markets. To illustrate the tradeooffs that influence the choice of pricing structure, we take a closer look at the peer-to-peer rental platform Airbnb. Airbnb allows users to rent short-term lodging such as spare rooms, apartments, vacation homes etc that are offered by independent hosts. Airbnb allows hosts to list their short term rental on the platform for free and charges a commission fee in the form of percentage of the listing price when the listing is booked. Airbnb allows hosts to price their own listing. Pricing is one of the most important factors for the hosts to be successful on the Airbnb platform. A majority of the hosts on Airbnb are non-professional hosts looking to earn some extra income by renting space in their homes or apartments. It takes several attempts for hosts to get the right price that balances bookings and revenue. Much of the host attrition on Airbnb is related to pricing issues. Hosts either lose revenue on bookings by pricing too low or fail to get a sufficient number of bookings by pricing too high. Further highlighting the importance of pricing, several startups (Beyond Pricing, AirDNA, Wheelhouse etc) offer services to help hosts determine the best price for their listing using data collected from the Airbnb platform. The listing price depends on several factors such as neighborhood characteristics (for example, downtown or suburb, proximity to tourist attractions etc.), type of listing (single room, apartment, whole house etc.), size, design (modern, vintage, colonial etc.), amenities, host characteristics (such as new or experienced host, host rating etc.) and competition (other listings in the area and their prices) etc. The unique nature of each listing creates information asymmetry between what the agents know about the listing and what the platform can observe. However, individual hosts do not have the resources or the expertise to perform a detailed analysis of the market and understand the competition. Platforms may have better information about market characteristics as they have access to the price and booking information of competing listings that an individual host is not be able to observe. To correct for this deficiency, some platforms using decentralized pricing have been investing heavily in information technology to understand market conditions and share the information with agents in the form of pricing suggestions and fair market prices. For example, Airbnb introduced a Smart pricing tool in 2017 that provides pricing suggestions to hosts based on different types of information, including listing characteristics, market demand data (competition, number of guests searching for listings) etc. Similarly, Homeaway introduced a dynamic pricing tool in 2018 that uses machine learning based models to provide pricing suggestions to hosts.

1.1 Platform Markets Vs versus Traditional Firms

As we discuss below, platforms differ from traditional firms in at least three different ways, that might affect the choices of centralized or decentralized controls.

1.1.1 Control of Resources

Traditional firms are largely defined by centralized ownership and control of resources (e.g. capital investments, employees directly under their payroll). This lends naturally to centralized, firm-wide protocols that define various facets of products and services such as quality, reliability standards, delivery schedules and variety of offerings. In contrast, platforms (such as Airbnb and Uber) do not own the means or resources that are responsible for the supply of services. Instead, they are structured around independent agents who control resources. Moreover, these agents are idiosyncratic in their taste, experience and available resources, leading to a high level of individual variation in the services offered by the platform. Thus, platforms can neither directly control the deployment of resources nor standardize their offerings. Platforms have succeeded by transforming this potentially serious liability into an asset, creating viable new markets by tapping into these individual variations and shifting control to agents with unique resources. Indeed, examples abound of platforms that have unleashed the supply of an incredible variety of products and services to customers, in a manner that is simply impossible for traditional organizations. On the car-sharing platform Turo[1], users can rent a wide selection of cars ranging from run-of-the-mill models (Chevrolet Cruze, Toyota Camry) to classics (the 1968 Mercedes Coupe), vintage cars and the latest Tesla models. In fact, the Turo platform offers more than 850 unique makes and models, compared to the 30 models offered by traditional car-rental companies such as Hertz. Airbnb is famous for its unique lodging experiences, including castles, tree-houses, island homes, retro-fitted shipping containers and aircraft[2]. As several specialist travel and lifestyle publications have observed, such other-worldly accommodations are available for rent only through Airbnb; traditional hotels can hardly compete. Similarly, Spinlister is able to offer a wide variety of bicycles for rent, including the Long John cargo bike, the Viking Tarantino tandem, and the Swiss-made ZEM four-person cycle[3]. Flowspace [4] offers warehouse storage for businesses through a network of independent warehouse providers. Across its network, Flowspace can fulfill almost any storage need ranging from heavy equipment to temperature controlled, organic certified and pharmaceutical grade storage. Uber leverages its network of drivers to offer unmatched service availability, overcoming the limitations of geography and time of day that hinder traditional taxi services[5].

In sum, compared to traditional firms, platforms are characterized by a radical decentralization of the ownership of resources– a potentially serious liability that they transform into a differentiating factor in the marketplace.

1.1.2 Information Endowments: Local and Aggregate Intelligence

The radical decentralization in the ownership of resources under platforms (Section 1.1.1 above), leads to a significant shift in the role of information in decision-making, as discussed next.

The paper closest to ours, albeit in a very different setting, is Anand and Mendelson (1997), who model a traditional firm operating through its branches in multiple horizontal markets. Anand and Mendelson (1997) highlight the importance of *jointly* optimizing the firm's organizational structure ("who decides what") and its informational structure ("who knows what"), and point out that

[1]https://www.businessinsider.com/turo-car-rental-review
[2]https://www.architecturaldigest.com/gallery/most-beautiful-airbnb-in-every-state
[3]https://www.forbes.com/sites/carltonreid/2019/01/03/airbnb-style-bicycle-rental-platform-spinlister-relaunches-via-oprahs-favorite-bike-firm/c015036538f7
[4]https://www.flow.space/warehouse-storage
[5]https://hbr.org/2016/04/the-truth-about-how-ubers-app-manages-drivers

focusing on either of these (to the exclusion of the other) will inevitably lead to sub-optimal outcomes. For example, the efficacy of centralized or decentralized decision-making is critically dependent on the information the decision-makers possess. Hence, AM (1997) define the *coordination structure* of the firm as the combination of its organizational and informational structures. A complete specification of the firm must identify its entire coordination structure, including its organizational and informational structures.

Anand and Mendelson define the firm's *information endowment* as the allocation of potentially useful information among its decision makers. Our modeling of information, in the context of a platform, builds on Anand and Mendelson (1997)'s concepts of IS. In our context, agents have access to information pertaining to their own listings, which we term *local intelligence*. On the other hand, the platform can have access to information on all its listings, which would be the sum total of the local intelligence of each listing. We term this *aggregate intelligence*. We assume (initially) that the platform's information endowment is the aggregate intelligence. Of course, decision makers' information endowments are fluid, and can be altered by information sharing or costly investments in information acquisition, which we also investigate in this research.

1.1.3 Agent Competition

An important structural difference between traditional firms and platforms arises from the nature of competition among their different units. Divisions within a firm compete (implicitly or explicitly) for allocation of the firm's scarce resources, but rarely compete in output markets. This reality is reflected in Anand and Mendelson (1997)'s model, which focuses on how the firm's scarce resource can be allocated among its branches to maximize the firm's overall profits. The situation with agents selling through a platform is quite the reverse. Since resources are largely owned by individual agents (recall Section 1.1.1), their allocation is a *fait accompli*; hence, the problem of allocation of resources doesn't arise for platforms. However, platforms are confronted with the reality that agents are in competition with each other. To study these issues, we model multiple agents selling through a platform. In contrast to Anand and Mendelson (1997), each agent in our model competes with other agents and maximizes her own individual profits. Each agent's demand is affected by both her own price and the prices of competing agents.

1.2 Summary of our Findings

A fundamental question we try to answer in our research is: How should platform markets be organized? Specifically, should decision making in a platform market be centralized or decentralized, i.e., controlled by the platform or by agents? The extant literature has extensively studied the question of optimal organization of traditional firms. However, these insights may not extend to platform markets, which differ from traditional firms in at least the three important dimensions we identified in Section 1.1: control of resources, information endowments of decision makers and competition among agents.

Following Anand and Mendelson (1997), we define the *coordination structure* (CS) of the platform market as the combination of two elements: its *decision rights structure* ('who decides what') and *information endowment* ('who knows what'). The platform provides services through independent, competing agents and faces demand uncertainty that can be partly alleviated through information. The platform and the agents differ in their initial information endowments: Each agent has access only to *local intelligence* (demand information pertaining to her own listing), whereas the platform has *aggregate intelligence*– which is the aggregate of the local intelligence for all listings. We model different facets of information endowment, including local and aggregate intelligence, the ability to acquire and share intelligence, and non-transferable specific knowledge of both the platform and agents.

Anand and Mendelson (1997) find that, for a traditional firm, the decentralized CS dominates the centralized CS in expected profits, in spite of the superior coordination enabled by centralized CS. In contrast, we find that the centralized CS dominates the decentralized CS in platform markets; see Section 4. Perhaps a crucial factor driving this difference is competition among agents. Indeed, we show that the performance gap between centralized and decentralized CS increases sharply with the degree of competition (Theorem 3 of Section 4).

To understand the drivers of this performance gap, we construct a hybrid, decentralized coordination structure termed D^{Team}, wherein agents act as a *team*, with their common objective being to maximize overall profits (Section 4.3.1). However, the decision-rights structure and agents' information endowments are identical to those under decentralized CS. Although the objectives of the platform and the agents are perfectly aligned, we find that C outperforms D^{Team} whenever agents compete (Theorem 4 of Section 4.3.1).

Another factor driving the superior performance of centralized CS over decentralized CS could be the difference in information endowments– recall that each agent has her local intelligence, whereas the center has access to aggregate intelligence. A possible remedy for the agent's lack of aggregate intelligence, observed in practice, is for the platform to enable sharing of aggregate intelligence with agents. Section 5 models such information sharing among agents, and provides several cutting-edge examples of such information sharing being operationalized in practice. We term the decentralized CS with information sharing as D^{IS}. Theorem 6 of Section 5 shows that C outperforms D^{IS}: Even when the platform and agents have identical information endowments, the centralized coordination under C leads to superior performance.

Next, in Section 5.2, we model costs to acquiring information, and allow investments in demand information as a decision variable. We find that C will invest at least as much in demand information as D (Corollary 1). This is because centralized CS uses information more effectively than decentralized CS, to coordinate prices across agents. Thus the marginal benefit of information is greater under C than under D, leading to greater investments for the same marginal costs of information.

Then, in Section 6, we compare the performance of C and D in the presence of extrinsic competition. We find that centralized coordination continues to play a dominant role in driving the performance gap between C and D, even though the scope of centralized coordination is diminished by the presence of the external competitor.

Finally, Section 7 extends our analysis to the case where decision points have access to *non-transferable specific knowledge*. Specifically, we are interested in whether decentralized CS can compensate for the lack of coordination with superior specific knowledge of local conditions. We parameterize both *local* specific knowledge (available exclusively to individual agents, and not transferable to the platform) and *central* specific knowledge (available exclusively to the platform, and not transferable to agents). Once again, we find that centralized CS dominates decentralized CS over a wide range of parameter values, even at low to moderate levels of competition. At higher levels of competition, centralized CS completely dominates decentralized CS– even when agents have perfect local specific knowledge, and the platform has no central specific knowledge.

2 Literature Review

***** TO BE COMPLETED *****

3 Elements of the Model

3.1 A Platform with Multiple Agents

We model a platform that operates an online marketplace enabling transactions between independent agents and customers. An agent offers a service or asset to potential customers as a 'listing' on the platform. For example, Airbnb is a platform on which hosts (agents) offer short-term accommodations (listings) to guests (customers). As the intermediary in the interaction between agents and customers, the platform charges a fee. In practice, the most widely used fee-structure takes the form of revenue sharing between the platform and each agent, which we assume in our analysis (Hagiu and Wright 2019). Revenue-sharing terms are fairly stable within an industry, and revised infrequently relative to operational decisions such as pricing which are the focus of our model. Indeed, Bhargava (2022) point out that "...the practice of all dominant platforms that employ revenue-sharing business models [is]... a uniform non-discriminatory linear revenue-sharing scheme." Hence, as in Bhargava (2022)'s model, we assume that the platform takes a fixed percentage $\gamma \in (0,1)$ of the revenues from each listing, while the agent of that listing retains the rest (which is the fraction $(1-\gamma)$ of revenues).

3.2 Agent Competition

In a platform setting, any individual agent has limited demand for his product or service, which is exacerbated by competition from other agents. In Airbnb, for example, agents compete with other agents to attract customers to their own listing. We model agents operating in overlapping markets and hence competing with each other. We also parameterize the degree to which markets overlap, with the case of non-overlapping markets arising as a special, polar case (see Section 3.3 below).

In the interests of expositional simplicity and clarity, we will focus our analysis on the case of a platform operating with two agents indexed by $i \in \{1, 2\}$. However, it is easily shown that our results extend without loss of generality to a platform operating with any arbitrary number of agents $n > 1$.

3.3 Customer Demand as Booking Probability

Consider any arbitrary listing (equivalently, agent) $i \in \{1, 2\}$. We represent the customer demand for listing i as the *booking probability* $b_i \in [0, 1]$, which is the probability that a given listing will be booked at an offered price $p_i \geq 0$. A higher b_i indicates a higher probability of being booked, implying that the listing is more desirable. When markets overlap, agents compete with each other, and the booking probability b_i is affected not just by the offer price p_i but by the competition's price p_j, where $j \in \{1, 2\}$ and $j \neq i$. We define the booking probability $b_i \in [0, 1]$ as

$$b_i = \begin{cases} a_i - p_i + \delta p_j & \text{when} \quad 0 < a_i - p_i + \delta p_j < 1 \\ 0 & \text{when} \quad a_i - p_i + \delta p_j < 0 \\ 1 & \text{when} \quad a_i - p_i + \delta p_j > 1 \end{cases} \tag{1}$$

where $\delta \in [0, 1)$ captures the impact of the competitor's price relative to one's own price on the demand, as in the seminal work of McGuire and Staelin (1983). We term δ as the *competition intensity*. The impact of the competitor's price is most often dwarfed by the effect of one's own price on one's demand, and hence we have $\delta < 1$. At the other end of the scale, when $\delta = 0$, each agent has an exclusive market and his booking probability is unaffected by the competitor's price. In practice, and especially in the context of platforms, where users' search costs are minimal, we would expect that $0 \ll \delta \ll 1$, so that the competition's price p_j has a substantial impact on the booking probability b_i, but clearly less so than the own price p_i.

3.4 Demand Uncertainty

Demand uncertainty is the bane of many markets, and even more so of platform-based markets that are structured around independent, idiosyncratic agents (recall discussion in Section 1.1.1). For a platform such as Airbnb, the customer demand for any listing is influenced by both listing-specific factors such as quality, location, features and agent taste, and exogenous factors such as special events, the weather and economic conditions (in addition, of course, to prices). Further, customers' tastes and requirements are unpredictable. Thus the booking probability b_i of equation (1), which is in essence the demand curve for listing i, is uncertain.

We model the uncertainty in the demand curve for listing i through a random "state" variable s_i. Following Anand and Mendelson (1997), we assume that the demand state is binary, i.e, "High" or "Low", with equal probability. Correspondingly, $s_i \in \{0, 1\}$, where $s_i = 1$ corresponds to "High" demand, and $s_i = 0$ to "Low" demand. As is common in the literature (*cf* Anand and Mendelson 1997), we assume that demand uncertainty affects the booking probability given by equation (1) through the intercept a_i. Specifically, we let

$$a_i = \begin{cases} \frac{1+\theta}{2} & \text{when} \quad s_i = 1 \\ \frac{1-\theta}{2} & \text{when} \quad s_i = 0 \end{cases} \tag{2}$$

where $0 \leq \theta \leq 1$. The intercept defined in (2) has the following useful properties: (i) $0 \leq a_i \leq 1$ for all θ; (ii) the mean intercept is $\frac{1}{2}$ which is independent of θ; and finally (iii) the standard deviation of the intercept is $\frac{\theta}{2}$. Thus θ does not affect the mean, and is purely a measure of the uncertainty in the demand curve. Setting $\theta = 0$ in expression (2), the uncertainty in our model vanishes, in which special case our demand model becomes equivalent to that of McGuire and Staelin (1983). We assume that equations (1) and (2), which together characterize the booking probability under demand uncertainty, are common knowledge. Although demand uncertainty is inevitable in these markets, it can be partially mitigated through information, which is modeled next.

3.5 Information Endowments

Following Anand and Mendelson (1997), we define a *decision point* as the entity (platform or agents in our setting) responsible for a decision, such as the pricing of a specific listing, and its *information endowment* as the aggregate of all potentially useful information relevant to the decision. To characterize information endowments in our setting, we distinguish between local and aggregate intelligence.

3.5.1 Local Intelligence

Local intelligence relevant to a listing includes listing features, information on the historical demand, and crucially, the sensitivity of demand to different combinations of listing features. We model local intelligence pertaining to listing i as an imperfect signal L_i on the demand state s_i (for $i \in \{1, 2\}$). The *precision* of local intelligence L_i is captured by the parameter α. Thus, for $s_i \in \{0, 1\}$,

$$L_i = \begin{cases} s_i & \text{with probability} \quad \alpha \\ 1 - s_i & \text{with probability} \quad 1 - \alpha \end{cases} \tag{3}$$

where, without loss of generality, $\frac{1}{2} \leq \alpha \leq 1$. The signal L_i accurately captures the true state of demand s_i with probability α and is inaccurate with probability $1 - \alpha$.

3.5.2 Aggregate Intelligence

We define *aggregate intelligence* as the sum total of the intelligence available on all listings, which, in the case of two listings, is $\bigcup_{i \in \{1,2\}} \{L_i\}$, where L_i was specified in (3).

3.5.3 Agents' and Platform's Information Endowments

We assume initially that each agent i's information endowment is precisely the local intelligence pertaining to his or her own listing i. Thus, each agent has no intelligence on any other listing. Indeed, in many platform settings, agents (such as Airbnb's individual home-owners or Turo's rental vehicle owners) have little access to historical demand information on other listings (Huang 2022). On the other hand, the platform can obtain such information on all the listings on its site. Thus we assume that the platform's information endowment is the aggregate intelligence. Of course, decision makers' information endowments are fluid, and can be altered by information sharing or costly investments in information acquisition, which we investigate later.

3.6 Coordination Structures

Anand and Mendelson (1997) define a firm's *coordination structure* (abbreviated to *CS*) as the combination of its *decision-rights structure*– which decision point is assigned the power to make which decisions– and its information endowment. They show that the firm's decision-rights structure and information endowment must be co-determined; varying one while leaving the other fixed is likely to lead to suboptimal results. In other words, it is critical for the firm to *jointly* optimize its decision-rights structure and information endowment, i.e., its CS. Our context of a platform with independent, competing agents is different from Anand and Mendelson (1997)'s traditional firm. Nevertheless, as we will show, these constructs encapsulate the critical features that drive platform organization and performance, and are indispensable to our analysis as well.

We model centralized and decentralized decision-rights structures. In the *centralized* decision-rights structure, the platform retains the decision-rights over pricing for all the listings, whereas in the *decentralized* decision-rights structure, individual agents hold the decision-rights for their respective listings. These two decision-rights structures, in combination with the information endowments defined in Section 3.5.3, lead naturally to centralized and decentralized CS, which we analyze in the next section. In later sections, we analyze alternative CS that arise through information sharing or costly investments in information acquisition.

4 Analysis

4.1 Centralized CS

Under the centralized CS, the platform controls the pricing for all listings, and its information endowment is the aggregate intelligence $\bigcup_{i\in\{1,2\}}\{L_i\}$. We use the notation $p_i(L_i; \bigcup L_{-i})$ to denote the price p_i of listing i, when the signal on its own demand is L_i and the aggregate intelligence on competing listings is $\bigcup\{L_{-i}\}$. When there are just two listings, indexed by $i \in \{1,2\}$, the notation reduces to $p_1(L_1; L_2)$ for the price of listing 1 and $p_2(L_2; L_1)$ for the price of listing 2.

The expected revenue from listing 1 is given by $p_1(L_1; L_2)b_1(L_1; L_2)$, where the booking probability $b_1(L_1, L_2) = a_1(L_1) - p_1(L_1; L_2) + \delta p_2(L_2; L_1)$ from equation (1). The intercept a_i is determined only by the signal L_i; hence we use the notation $a_i(L_i)$ to denote the expected demand intercept when the local intelligence is L_i. Thus, for example, when $L_i = 0$, $a_i(0) = (1-\alpha)\frac{(1+\theta)}{2} + \alpha\frac{(1-\theta)}{2}$ by equations (2) and (3). Similarly, the expected revenue from listing 2 is given by $p_2(L_2; L_1)b_2(L_2; L_1)$, where $b_2(L_2, L_1) = a_2(L_2) - p_2(L_2; L_1) + \delta p_1(L_1; L_2)$.

Recall from Section 3.1 that γ is the share of revenues accruing to the platform. Hence, under the centralized coordination structure, the platform chooses prices $p_1(L_1; L_2)$ and $p_2(L_2; L_1)$ for listings 1 and 2, respectively, so as to maximize its share of overall profit[6] as follows:

$$\max_{p_1(L_1;L_2),p_2(L_2;L_1)} \gamma\Big(p_1(L_1; L_2)b_1(L_1; L_2) + p_2(L_2; L_1)b_2(L_2; L_1)\Big),$$

which is equivalent to maximizing the overall profit $\Big(p_1(L_1; L_2)b_1(L_1; L_2) + p_2(L_2; L_1)b_2(L_2; L_1)\Big)$. Thus, γ doesn't affect the platform's optimal choices under the centralized CS. Moreover, γ (and $(1-\gamma)$) are merely multipliers to the total profits that determine the split between platform and agent profits. Hence we focus only on total profits in our analysis. (By the same logic, we focus only on total profits in the decentralized CS as well.) Theorem 1 derives the optimal prices and expected total profits from the two listings under the centralized CS.

Theorem 1 *Under the centralized CS, the optimal prices and expected total profits are as follows:*

(A) Prices: The prices $p_i(L_i; L_{-i})$, for $i \in \{1, 2\}$ are:

$$p_i(L_i; L-i) = \begin{cases} \frac{(1-\alpha)(1+\theta)+\alpha(1-\theta)}{2(1-\delta)} & when \quad L_i = 0, L_{-i} = 0 \\ \frac{\alpha(1+\theta)+(1-\alpha)(1-\theta)+\delta((1-\alpha)(1+\theta)+\alpha(1-\theta))}{2(1-\delta^2)} & when \quad L_i = 1, L_{-i} = 0 \\ \frac{(1-\alpha)(1+\theta)+\alpha(1-\theta)+\delta(\alpha(1+\theta)+(1-\alpha)(1-\theta))}{2(1-\delta^2)} & when \quad L_i = 0, L_{-i} = 1 \\ \frac{\alpha(1+\theta)+(1-\alpha)(1-\theta)}{2(1-\delta)} & when \quad L_i = 1, L_{-i} = 1 \end{cases}$$

(B) Expected Total Profits:

$$\pi^C = \frac{\theta^2(2\alpha-1)^2}{8(1-\delta^2)} + \frac{1}{8(1-\delta)} \tag{4}$$

4.2 Decentralized Coordination Structure

Under the decentralized CS, agent i controls the pricing for her own listing i. The agent's information endowment is simply her local intelligence L_i on her own listing; she has no information on the demands for competing listings. We use the notation $p_i(L_i)$ to denote the price of listing i with local intelligence L_i. Agent i chooses price $p_i(L_i)$ to maximize her expected revenues $E((1-\gamma)p_i(L_i)b_i(L_i)) = (1-\gamma)\cdot E(p_i(L_i)b_i(L_i))$. By the same arguments as in centralized CS, this is equivalent to maximizing $E(p_i(L_i)b_i(L_i))$, the overall profit from the listing. From equation (1),

[6] Note that profit and revenue maximization are identical in this model for both centralized and decentralized CS. In Section 5.2, where we explicitly model information acquisition costs, profit and revenue maximization diverge: The contract is still revenue sharing, but the decision points (platform or individual agents) maximize their own profits.

the booking probability for listing i is $b_i(L_i) = a_i(L_i) - p_i(L_i) + \delta p_j(L_j)$, for $i, j \in \{1, 2\}$ and $i \neq j$. Since the booking probabilities depend on *both* prices, which are in turn functions of their respective information signals, the outcome under decentralized CS is the Nash equilibrium arising from the simultaneous solution of two maximization problems:

$$\begin{cases} \max\limits_{p_1(L_1)} E[p_1(L_1)b_1(L_1)] \\ \qquad\qquad \& \\ \max\limits_{p_2(L_2)} E[p_2(L_2)b_2(L_2)] \end{cases}$$

The optimal prices and expected profits from both listings under the decentralized CS are provided in the following theorem.

Theorem 2 *Under the decentralized CS, the optimal prices and expected total profits are as follows:*

(A) Prices: The prices $p_i(L_i)$, for $i \in \{1, 2\}$ are:

$$p_i(L_i) = \begin{cases} \alpha(\frac{1+\theta}{4}) + (1-\alpha)\frac{(1-\theta)}{4} + \frac{\delta}{4(2-\delta)} & when \quad L_i = 1 \\ \alpha(\frac{1-\theta}{4}) + (1-\alpha)\frac{(1+\theta)}{4} + \frac{\delta}{4(2-\delta)} & when \quad L_i = 0 \end{cases}$$

(B) Expected Total Profits:

$$\pi^D = \frac{\theta^2(2\alpha-1)^2}{8} + \frac{1}{2(2-\delta)^2} \tag{5}$$

4.3 Comparison of Profits

Theorem 3 compares the profits of centralized and decentralized CS.

Theorem 3 *The centralized CS dominates the decentralized CS in expected profits. Thus,*

(i) $\pi^C \geq \pi^D$ for $\delta \geq 0$, with strict inequality for $\delta > 0$. Specifically,

$$\pi^C - \pi^D = \frac{(2\alpha-1)^2\delta^2\theta^2}{8(1-\delta^2)} + \frac{\delta^2}{8(2-\delta)^2(1-\delta)} \tag{6}$$

(ii) The performance gap $\pi^C - \pi^D$ is strictly increasing in competition intensity δ, information precision α and uncertainty parameter θ.

Part (i) of Theorem 3 suggests that the competition intensity δ is a key driver of the performance gap between centralized and decentralized CS. Anand and Mendelson (1997) model diseconomies of scale, which induce negative production externalities and favor coordination across a firm's multiple branches. Surprisingly, however, they find that the superior coordination enabled by centralized CS matters little. In fact, "...in a variety of cases, the performance of decentralized CS dominates that of centralized CS in spite of the latter's superior coordination" (Anand and Mendelson 1997). The crucial difference in our model is that agents on the same platform compete with each other. Thus we see that coordination enabled by centralized CS is valuable under agent competition. In fact, the competition intensity δ plays such an important part that the centralized CS strictly outperforms the decentralized CS even for very small values of δ, for any feasible value of α. Indeed, even when $\alpha = \frac{1}{2}$, so that both local and aggregate intelligence vanish, we see that $\pi^C > \pi^D$ for all $\delta > 0$.

Part (ii) of Theorem 3 is an outcome of the superior coordination enabled by centralized CS. Centralized coordination becomes more valuable as the competition intensity, the precision of demand information and the demand uncertainty increase.

4.3.1 Parsing the Performance Gap

In order to better understand the performance gap between centralized and decentralized CS, we introduce a hybrid coordination structure, which we call the D^{Team} CS. The defining features of the D^{Team} CS are: *(i)* The decision-rights structure and agents' information endowments are identical to those under decentralized CS. Thus, agent i holds the decision-rights for her own listing i, and her information endowment is only her local intelligence L_i. *(ii)* However, agents' *objective functions* are aligned with centralized CS– their common objective is to coordinate among themselves (within the limitations imposed by their information endowments) to maximize overall profits, as in the centralized CS. In other words, agents act as a *team*, as defined by Anand and Mendelson (1997). The D^{Team} CS is plausible in the context of a firm with multiple branches (as in Anand and Mendelson 1997), but impractical in our context of a platform with independent, competing agents. We use the D^{Team} CS primarily to gain insights into the drivers of the performance gap between centralized and decentralized CS.

As before, we analyze the case of two listings indexed by $i \in \{1, 2\}$. While agents 1 and 2 choose their own prices $p_1(L_1)$ and $p_2(L_2)$ respectively, both agents seek to maximize overall profits, as in centralized CS. The outcome of the D^{Team} CS is the Nash equilibrium arising from the simultaneous solution of two maximization problems:

$$\begin{cases} \max\limits_{p_1(L_1)} E[p_1(L_1)b_1(L_1) + p_2(L_2)b_2(L_2)] \\ \& \\ \max\limits_{p_2(L_2)} E[p_2(L_2)b_2(L_2) + p_1(L_1)b_1(L_1)] \end{cases}$$

Theorem 4 analyzes the profits under the D^{Team} coordination structure.

Theorem 4 *(i) The expected total profits under the D^{Team} CS are:*

$$\pi^{D^{Team}} = \frac{\theta^2(2\alpha - 1)^2}{8} + \frac{1}{8(1-\delta)} \tag{7}$$

(ii) $\pi^C \geq \pi^{D^{Team}}$ with equality holding only when $\alpha = \frac{1}{2}$, $\delta = 0$ or $\theta = 0$.

(iii) $\pi^{D^{Team}} \geq \pi^D$ with equality holding only when $\delta = 0$.

Comparing $\pi^{D^{Team}}$ (given by expression (7) in Part *(*i) of Theorem 4) with π^C (expression (4)) and π^D (expression (5)) reveals the hybrid nature of the D^{Team} coordination structure. Indeed, $\pi^{D^{Team}}$ is the composite of the first term of π^C and the second term of π^D.

Part *(*ii) of Theorem 4 proves that for all non-trivial parameter values, $\pi^C > \pi^{D^{Team}}$. By construction, the agents' objectives under D^{Team} and the platform's objective under the centralized CS are the same: to coordinate prices across listings in order to maximize overall profits. Thus the difference in their profits must arise solely from their different information endowments. Indeed, this is clear from examining the limiting cases when the two profits are identical: (i) When $\alpha = \frac{1}{2}$, the information endowments of the platform and the agents are identical in the sense that neither has any information on the demand of listings. (ii) When $\delta = 0$, the listings are independent monopolies, and the joint optimization problem breaks down into separate optimization problems for each listing. Here, the aggregate intelligence under C serves the same purpose as the local intelligences under D^{Team}. (iii) Finally, when $\theta = 0$, there is no demand uncertainty, and so information endowments are irrelevant.

Part *(*iii) of Theorem 4 proves that for all non-trivial parameter values, $\pi^{D^{Team}} > \pi^D$. Since the allocation of decision-rights and agents' information endowments are identical under decentralized and D^{Team} CS, the superior performance of D^{Team} compared to D, given by expression (5), is solely due to the coordination efforts of agents under D^{Team}. Under D, every agent competes with all other agents to maximize her own profits, which erodes everybody's profits. Only in the limiting case when $\delta = 0$, i.e., agents are monopolists precluding any need for coordination, are the profits identical.

To summarize, Theorem 4 proves a strict performance ranking among the three coordination structures. Barring trivial cases, $\pi^C > \pi^{D^{Team}} > \pi^D$.

5 Sharing Aggregate Intelligence

We showed that the centralized CS outperforms the decentralized CS (Theorem 3), even when agents operate as a team to maximize overall profits (Theorem 4; part *(ii)*). Theorem 4 in particular established that the decentralized CS is handicapped by agents' lack of aggregate intelligence. A possible remedy is for the platform to share its intelligence on all listings with agents, so that all agents have access to aggregate intelligence.

Indeed, several platforms have attempted to do so. For example, Airbnb, VRBO and Turo, that use decentralized CS, share their intelligence with agents through pricing tools. Airbnb's *Smart Pricing* tool collates each individual listing's demand metrics (such as its booking history, listing features and host ratings) with market metrics on its competing listings, including their historical demand patterns and prices, to provide pricing suggestions. Airbnb's competitor VRBO offers a similar tool called *MarketMaker* that aggregates data on average nightly prices and occupancy rates for competing listings, and shows how the prices set by individual agents compare with market rates. Similarly, Turo's automatic pricing tool provides pricing suggestions for renting out cars by showing how the listing price compares with the rental prices of similar listings in the market.

In order to study the impact of such information sharing, we analyze the decentralized CS with aggregate intelligence, which we term D^{IS}. (Analogously, Anand and Mendelson (1997) analyze information sharing between a firm and its branches in what they term the "Fully Distributed" CS.)

5.1 Decentralized Coordination Structure with Information Sharing

Under D^{IS}, agent i has decision rights over the pricing of her own listing i, as in decentralized CS. However, the information endowment is altered by information sharing. Thus, each agent's information endowment under D^{IS} is the aggregate intelligence $\bigcup_{i\in\{1,2\}}\{L_i\}$, which is the same as the platform's information endowment in centralized CS (recall Section 4.1). As in Section 4.1, $p_i(L_i;\bigcup L_{-i})$ denotes the price p_i of listing i, when the signal on its own demand is L_i and the aggregate intelligence on competing listings is $\bigcup\{L_{-i}\}$. In the case of two listings (i.e., $i \in \{1,2\}$), agents 1 and 2 choose prices $p_1(L_1;L_2)$ and $p_2(L_2;L_1)$ respectively, to maximize revenues from their own listings. As before, the outcome under D^{IS} is the Nash equilibrium arising from the simultaneous solution of the maximization problems:

$$\begin{cases} \max\limits_{p_1(L_1;L_2)} p_1(L_1;L_2)b_1(L_1;L_2) \\ \qquad\qquad\& \\ \max\limits_{p_2(L_2;L_1)} p_2(L_2;L_1)b_2(L_2;L_1) \end{cases}$$

Theorem 5 analyzes the profits under the D^{IS} coordination structure.

Theorem 5 *Under the decentralized CS with information sharing (D^{IS}), the optimal prices and expected total profits are as follows:*

(A) Prices: The prices $p_i(L_i;L_{-i})$, for $i \in \{1,2\}$ are:

$$p_i(L_i;L{-}i) = \begin{cases} \frac{\alpha(1-\theta)+(1-\alpha)(1+\theta)}{2(2-\delta)} & when \quad L_i=0, L_{-i}=0 \\ \frac{2(1-\theta+2\alpha\theta)+\delta(1+\theta-2\alpha\theta)}{2(4-\delta^2)} & when \quad L_i=1, L_{-i}=0 \\ \frac{2(1+\theta-2\alpha\theta)+\delta(1-\theta+2\alpha\theta)}{2(4-\delta^2)} & when \quad L_i=0, L_{-i}=1 \\ \frac{\alpha(1+\theta)+(1-\alpha)(1-\theta)}{2(2-\delta)} & when \quad L_i=1, L_{-i}=1 \end{cases}$$

(B) Expected Total Profits:

$$\pi^{D^{IS}} = \frac{(2\alpha-1)^2\theta^2(4+\delta^2)}{2(4-\delta^2)^2} + \frac{1}{2(2-\delta)^2} \tag{8}$$

5.1.1 Comparison of Profits

Theorem 6 *The centralized CS dominates the D^{IS} CS in expected profits. Thus, $\pi^C \geq \pi^{D^{IS}}$. Moreover,*

(i) $\pi^C > \pi^{D^{IS}}$ for $\delta > 0$.

(ii) The performance gap $\pi^C - \pi^{D^{IS}}$ is strictly increasing in competition intensity δ, information precision α and uncertainty parameter θ.

To understand the drivers of Theorem 6, recall that the decision points under C and D^{IS} have identical information endowments. Thus, the difference in profits $\pi^C - \pi^{D^{IS}}$ must arise only from the divergent incentives of the platform and the agents. Anand and Mendelson (1997) observe that "...agency costs sometimes undermine the advantages of distributed systems [D^{IS} in our context]..." Theorem 6 shows that this insight applies to our context of a platform with independent, competing agents as well. In fact, the effect of divergent incentives, and hence the *value of centralized coordination*, is even greater in our model because of agent competition. Whereas, in Anand and Mendelson (1997), decentralized CS with information sharing (corresponding to our D^{IS}) often outperformed centralized CS even with the problem of divergent incentives, the centralized CS always outperforms D^{IS} in our setting.

Part (i) of Theorem 6 establishes that δ continues to play a crucial role in the comparative performance of C and D^{IS}, even with their identical information endowments, as it did in the case of C and D (recall Theorem 3). Part (ii) of Theorem 6 is a direct outcome of the superior coordination enabled by centralized CS. First, when competition among agents intensifies, i.e., δ increases, centralization is all the more valuable to mitigate the harmful aspects of competition. Second, centralization enables a coordinated response to extract the most from demand information; hence $\pi^C - \pi^{D^{IS}}$ increases with information precision α. Third, coordination is more valuable as demand uncertainty (measured by θ) increases, and hence, $\pi^C - \pi^{D^{IS}}$ increases with θ.

Theorem 7 *The D^{IS} CS dominates the decentralized CS in expected profits. Thus, $\pi^{D^{IS}} \geq \pi^D$. Moreover,*

(i) $\pi^{D^{IS}} > \pi^D$ provided $\delta, \theta > 0$ and $\alpha > \frac{1}{2}$.

(ii) The performance gap $\pi^{D^{IS}} - \pi^D$ is strictly increasing in competition intensity δ, information precision α and uncertainty parameter θ.

Agents are the decision points under both D and D^{IS}, and under both CS, their objective is to maximize their own profits. Thus, both D and D^{IS} have identical incentive structures, and differ only in their information endowments. While each agent has only her local intelligence under D, she has access to aggregate intelligence under D^{IS}. Thus, the difference in profits $\pi^{D^{IS}} - \pi^D$ must arise purely from the additional value of aggregate intelligence over local intelligence. Part (i) of Theorem 7 shows that D^{IS} always outperforms D, barring the extreme cases $\delta = 0$, $\theta = 0$ and $\alpha = \frac{1}{2}$. In these cases, the incremental value of aggregate intelligence is zero. When $\delta = 0$, each agent is an independent monopolist, and information about other agents' demand (that aggregate intelligence provides) is irrelevant to her profits. When $\theta = 0$, there is no demand uncertainty, and hence, no role for either aggregate or local intelligence. Finally, when $\alpha = \frac{1}{2}$, neither aggregate nor local intelligence provides any demand information.

Part (ii) of Theorem 7 is driven by the incremental value of aggregate intelligence over local intelligence. First, as competition among agents intensifies, i.e., δ increases, the information on competitors' demand, which is available to agents only under D^{IS}, is more valuable. Second, as information precision α increases, π^D which depends only on information on one's own demand increases. However, $\pi^{D^{IS}}$ which depends on both own and competitor information increases at a faster rate. Hence, $\pi^{D^{IS}} - \pi^D$ increases with information precision α. Third, information is more valuable as demand uncertainty (measured by θ) increases, and since agents have more intelligence available to them under D^{IS}, $\pi^{D^{IS}} - \pi^D$ increases with θ.

To summarize, Theorems 6 and 7 together provide a strict performance ranking among the three coordination structures C, D and D^{IS}. In general, $\pi^C > \pi^{D^{IS}} > \pi^D$, with equality holding only in a few extreme cases.

In the analysis thus far, we assumed that information was available to the decision points as part of their information endowment, and moreover, at a fixed precision α. Hence we did not consider the cost of information acquisition. In the following Section, we relax both assumptions. First, we assume that the information precision α is selected endogenously (for each CS). Second, the cost of information acquisition is assumed to be an increasing function of the desired level of α. Comparing the optimal α (and corresponding investments) under each CS, we are able to study the marginal value of information in each CS, i.e., the ability of each CS to exploit information.

5.2 Investing in Information

As discussed previously, several platforms such as Uber, Lyft and Prosper employ centralized CS (with prices set by the platform), while others such as Airbnb, VRBO and Turo employ decentralized CS with agents free to set their own prices. In either case, platforms have increasingly recognized the value of demand information and made significant investments in acquiring information. For example, Uber, Lyft and Prosper have invested heavily in developing algorithms to predict demand, to help them set prices centrally. Uber employs more than a thousand leading scholars in the areas of Artificial Intelligence (AI) and Machine Learning (ML), and have developed several state of the art AI and ML tools and architectures to improve their demand prediction and pricing models [7]. Among platforms with decentralized CS, we discussed examples of demand prediction and pricing tools such as Airbnb's *Smart Pricing*, VRBO's *MarketMaker* and Turo's automatic pricing tool. Early versions of Airbnb's *Smart Pricing* tool considered a limited set of key attributes, but failed to account for the geographical granularity and diversity of listing features (ranging from a yurt to a castle) that greatly impacted prices. Subsequently, Airbnb invested in a machine learning, open source library

[7] https://www.forbes.com/sites/forbestechcouncil/2019/01/08/dynamic-pricing-the-secret-weapon-used-by-the-worlds-most-successful-companies/?sh=52721877168b

termed *Aerosolve* that considers thousands of factors, including geographic location and diversity of listing features at a very detailed level (Hill 2015).

In this Section, we assume that the investment in information acquisition is also a decision variable for the platform. We model the cost of information acquisition as increasing and convex in information precision α; specifically, the cost to acquire aggregate intelligence with precision $\alpha \in [\frac{1}{2}, 1]$ is $C(\alpha) = nc(\alpha - \frac{1}{2})^2$, where n is the number of listings and c is a constant. Thus, $C(\frac{1}{2}) = 0$, corresponding to no investments, and $C(1) = \frac{nc}{4}$, which is the cost of perfect demand information. The cost of information acquisition with imperfect precision $\alpha \in (\frac{1}{2}, 1)$ falls between these two extremes. We analyze and compare the two CS (C and D^{IS}) under which decision points possess aggregate intelligence.

Under the centralized CS, total expected revenue from two listings, for any α, is given by expression (4). The platform's problem is to choose the optimal level of information precision $\alpha_C^* \in [\frac{1}{2}, 1]$ that maximizes its expected profit (revenue net of cost). Thus,

$$\alpha_C^* = \arg\max_{\alpha}[\frac{\theta^2(2\alpha - 1)^2}{8(1-\delta^2)} + \frac{1}{8(1-\delta)} - 2c(\alpha - \frac{1}{2})^2]$$

Similarly, under the D^{IS} CS, total expected revenue from two listings is given by expression (8). The platform's problem is to choose the optimal level of information precision $\alpha_D^* \in [\frac{1}{2}, 1]$ that maximizes its expected profit. Thus,

$$\alpha_D^* = \arg\max_{\alpha}[\frac{(2\alpha - 1)^2\theta^2(4+\delta^2)}{2(4-\delta^2)^2} + \frac{1}{2(2-\delta)^2} - 2c(\alpha - \frac{1}{2})^2]$$

The results of the optimization are summarized in the next theorem.

Theorem 8 *When information acquisition is an endogenous decision for the platform, the optimal information precisions under both centralized and D^{IS} coordination structures are always either $\frac{1}{2}$ or 1. That is, the platform's optimal choice is to invest in maximum aggregate intelligence or not invest at all. Specifically, let $c^C = \frac{\theta^2}{4(1-\delta^2)}$ and $c^D = \frac{(4+\delta^2)\theta^2}{(4-\delta^2)^2}$. Then*

$$\alpha_C^* = \begin{cases} 1 & \text{if} \quad c \le c^C \\ \frac{1}{2} & \text{otherwise} \end{cases}$$

$$\alpha_D^* = \begin{cases} 1 & \text{if} \quad c \le c^D \\ \frac{1}{2} & \text{otherwise} \end{cases}$$

Corollary 1 *$\alpha_C^* > \alpha_D^*$ and $C(\alpha_C^*) > C(\alpha_D^*)$ when $c \in (c^D, c^C)$; $\alpha_C^* = \alpha_D^*$ and $C(\alpha_C^*) = C(\alpha_D^*)$ otherwise.*

Theorem 8 derives the optimal information precision (and *inter alia* the platform's optimal investment in information acquisition), under both centralized and decentralized CS. In each case, it is optimal for the platform to either invest in perfect information (when the cost coefficient c is below the respective thresholds c^C and c^D) or not invest at all (when c is above the respective thresholds).

Moreover, the thresholds under C and D are rank-ordered as $c^C \geq c^D$. This leads to Corollary 1, which establishes that the investment in information acquisition under centralized CS is always at least as much as (and, for some parameter values, strictly greater than) the investment under decentralized CS. Due to the superiority of centralized coordination, centralized CS better exploits demand information than decentralized CS, and hence, the marginal benefit of additional information is greater. Since the marginal cost functions are identical under both C and D, the platform invests more in information acquisition under centralized CS than under decentralized CS. To summarize, Theorem 8 and Corollary 1 together demonstrate the value of centralized coordination in the context of information acquisition, providing more insight into why C performs better than D^{IS} (Theorem 6). The platform's ability to coordinate prices across listings and moderate the deleterious effects of agent competition emerges as the dominant factor in the comparison of alternative CS.

6 Coordination under Extrinsic Competition

In addition to intra-platform competition from other agents (listings) within the platform, agents often face competition from external entities. For example, listings on AirBnB compete, not just with other listings on AirBnB, but also with listings extrinsic to the platform, such as traditional hotel chains, resorts and competing platforms. The crucial difference between extrinsic competition and intra-platform competition among agents is that the platform's choice of CS has no direct bearing on the external competitor. In this Section, we model extrinsic competition in addition to agent competition, and study its impact on the performance of alternative CS.

We extend the booking probability defined in expression (1) to incorporate competition from an external source. Let p_e denote the external competitor's price, where the subscript e denotes the $e$xternal (or $e$xtrinsic) competitor. We modify expression (1) to define the booking probabilities under extrinsic competition (for agents within the platform and the external competitor) as:

$$b_i = \begin{cases} a_i - p_i + \delta\left(\frac{p_j + k_e p_e}{2}\right) & \text{when} \quad 0 < a_i - p_i + \delta\left(\frac{p_j + k_e p_e}{2}\right) < 1 \\ 0 & \text{when} \quad a_i - p_i + \delta\left(\frac{p_j + k_e p_e}{2}\right) < 0 \\ 1 & \text{when} \quad a_i - p_i + \delta\left(\frac{p_j + k_e p_e}{2}\right) > 1 \end{cases} \tag{9}$$

and

$$b_e = \begin{cases} a_i - p_e + k_e \delta \bar{p} & \text{when} \quad 0 < a_i - p_e + k_e \delta \bar{p} < 1 \\ 0 & \text{when} \quad a_i - p_e + k_e \delta \bar{p} < 0 \\ 1 & \text{when} \quad a_i - p_e + k_e \delta \bar{p} > 1 \end{cases} \tag{10}$$

where $k_e \in (0, 1]$ is the scaling factor that captures the differential impact of extrinsic and intrinsic competition on each other's booking probability, and, in expression (10), $\bar{p} = \frac{p_i + p_j}{2}$ is the platform's average price. Thus, comparing the booking probabilities of (9) and (1), we see that the competing agent's price p_j in expression (1) is replaced in expression (9) by the average of the competing agent's price p_j and the external competitor's price p_e weighted by the scaling factor k_e. Similarly, in expression (10), the platform's prices affect the extrinsic competitor's booking probability in proportion to the scaling factor k_e. In general, we would expect that a listing's demand is affected less by the price of an external competitor than by the price of a listing competing within the same platform. For example, once logged on to a platform such as AirBnB, a customer's search costs within the platform would be lower than outside the platform. Moreover, a competing listing on the same platform is likely to be a closer substitute to a listing than an external competitor, and therefore influence the listing's booking probability more heavily. In all such cases, $k_e < 1$. In the limiting case, when there is no differential impact between extrinsic and intrinsic competition, $k_e = 1$. (Of course, when $k_e = 0$, the impact of extrinsic competition vanishes, and only intra-platform agent competition remains, as in the main model.) Theorem 9 derives the expected profits in equilibrium under the different coordination structures.

Theorem 9 *The platform's expected total profits under the C, D and D^{IS} coordination structures with extrinsic competition are:*

$$\pi_e^C = \frac{(1-2\alpha)^2\theta^2}{2(4-\delta^2)} + \frac{((2-\delta)(4+k_e\delta)^2}{4(8-4\delta-k_e^2\delta^2)^2}$$
$$\pi_e^D = \frac{\theta^2(2\alpha-1)^2}{8} + \frac{((4-\delta)(4+k_e\delta)-2k_e^2\delta^2)^2}{16(4-\delta)^2(4-\delta-k_e^2\delta^2)^2}$$
$$\pi_e^{D^{IS}} = \frac{2(16+\delta^2)(2\alpha-1)^2\theta^2}{(16-\delta^2)^2} + \frac{((4-\delta)(4+k_e\delta)-2k_e^2\delta^2)^2}{16(4-\delta)^2(4-\delta-k_e^2\delta^2)^2}$$

Recall that, in the absence of extrinsic competition, C outperforms D^{IS} and D (Section 5.1.1), driven by its superior ability to exploit available information through centralized coordination (Section 5.2). We analyze whether this result holds under extrinsic competition in the following theorem.

Theorem 10 *Centralized CS dominates decentralized CS even with extrinsic competition for any k_e. Specifically, $\pi_e^C \geq \pi_e^{D^{IS}} \geq \pi_e^D$ for $\delta \geq 0$, with strict inequalities when $\delta > 0$.*

Theorem 10 shows the superiority of centralized coordination of agents even under extrinsic competition, leading to $\pi_e^C \geq \pi_e^{D^{IS}}$. Indeed, the superior coordination of agents' prices within the platform (under C) improves the profits all parties– the platform, individual agents and even the external competitor (see the Technical Appendix).

Finally, we find that extrinsic competition differs from competition among agents *within* the platform in its impact on profits, as discussed below. A "fair" comparison requires setting $k_e = 1$ in Theorem 9, so that the *total* competitive pressure on a listing is the same with and without extrinsic competition; *only* the sources of competition differ in the two cases. Thus, competitive pressure is either solely due to competing agents on the same platform (as in the main model) or distributed between the external competitor and agents within the same platform (as in the model of this section). Indeed, when $k_e = 1$, we see that (by construction) $\pi_e^D = \pi^D$– after all, each agent is only concerned with her own profits, and is indifferent to the source– whether other agents or external– of the competition. However, we can show that $\pi_e^C < \pi^C$, since the platform cannot directly coordinate on prices with the external competitor, and its ability to control prices is thus limited to within the platform. Furthermore, $\pi_e^{D^{IS}} < \pi^{D^{IS}}$ for different (informational) reasons: Agents' information endowments under D^{IS} are limited to demand information on listings within the same platform, and excludes information on the external competitor's demand. Hence the scope of the demand information available to agents under $D^{IS}(e)$ is partial relative to D^{IS}.

7 Local and Central Specific Knowledge

In our analysis thus far, we assumed that the decision points (platform or agents) have the same accuracy of signal, α on the demand of individual listings.

Specifically, we assumed that the only information available to decision points (platform or agents) is the local intelligence on a listing. Since this information can be shared, all decision points with access to this information have the same accuracy of demand signal α. However, in addition to the local intelligence, decision points (platform or agents) may also have information that cannot be shared or communicated. Decision points can tap this information to form a more accurate demand signal than that based only on local intelligence. Anand and Mendelson (1997) classify a decision maker's informational endowment into two broad categories: *(i)* transferable *data* and *(ii)* nontransferable *specific knowledge*. Specific knowledge refers to specialized information possessed by decision points which cannot be transferred in a cost-effective or timely manner to other decision points (such as the firm's headquarters or the platform). Platform and agents can have specific knowledge based on their access to information on specific factors influencing demand for listings. We use the terms platform specific knowledge and agent specific knowledge to denote the information available only to the platform or to the agents, respectively.

An agent's specific knowledge, in combination with local intelligence on the listing can be more informative in predicting the demand for a listing, than just the local intelligence alone. [8] To understand why, AirBnB provides an excellent context. Unlike the commoditized accommodations offered by the hotel industry, every listing on Airbnb is unique, and spans a broad spectrum of location, quality and features (recall discussion in Section 1.1). Agents have specific knowledge of these attributes with respect to their own listings, not easily accessible to AirBnB. Factors like local events (football, conventions, music festivals etc), seasonal trends (like coastal vs ski towns), listing features (like river and ocean front properties, views, architecture, castles, tree houses) etc that can substantially impact the demand for listing can be difficult to capture in data.

Platform, on the other hand, has access to information on market wide demand factors that can impact the overall demand for services offered on the platform. In the context of AirBnB again, macro economic factors (like inflation, jobs, financial markets etc) and global factors (like pandemic

[8] Technically, specific knowledge in combination with local intelligence is more informative than just local intelligence alone in the Blackwellian sense; see Anand (2017).

restrictions, gas prices, travel trends etc) can influence the overall demand for short term rentals or create new demand patterns. Platform has the expertise (through experts and investments) to assess the impact of such factors on demand. In addition, platform can gather signals on the influence of such factors on demand by observing the aggregate number of users looking for rentals and analysing user actions like search patterns and cancellations. Platform's specific knowledge about market-wide demand factors combined with the aggregate intelligence on listings can be more informative than aggregate intelligence alone in predicting demand for listings.

To capture agent and platform specific knowledge, we allow the accuracy of demand signals received by the agents and platform to differ. Specifically, we let $\alpha_D \geq \alpha$ denote the accuracy of the agent's demand signal based on the combination of the agent's specific knowledge and local intelligence on the listing. Similarly, we let $\alpha_C \geq \alpha$ denote the accuracy of the platform's demand signal based on the platform's specific knowledge and aggregate intelligence on all listings. To derive the expected total profits under the centralized CS with platform specific knowledge, we replace α in the profit expression of (4) with α_C. Similarly, we derive the expected total profits under decentralized CS by replacing α in the profit expression of (5) with α_D. We compare the expected total profits under centralized and decentralized CS and present the results of the comparison in the following picture.

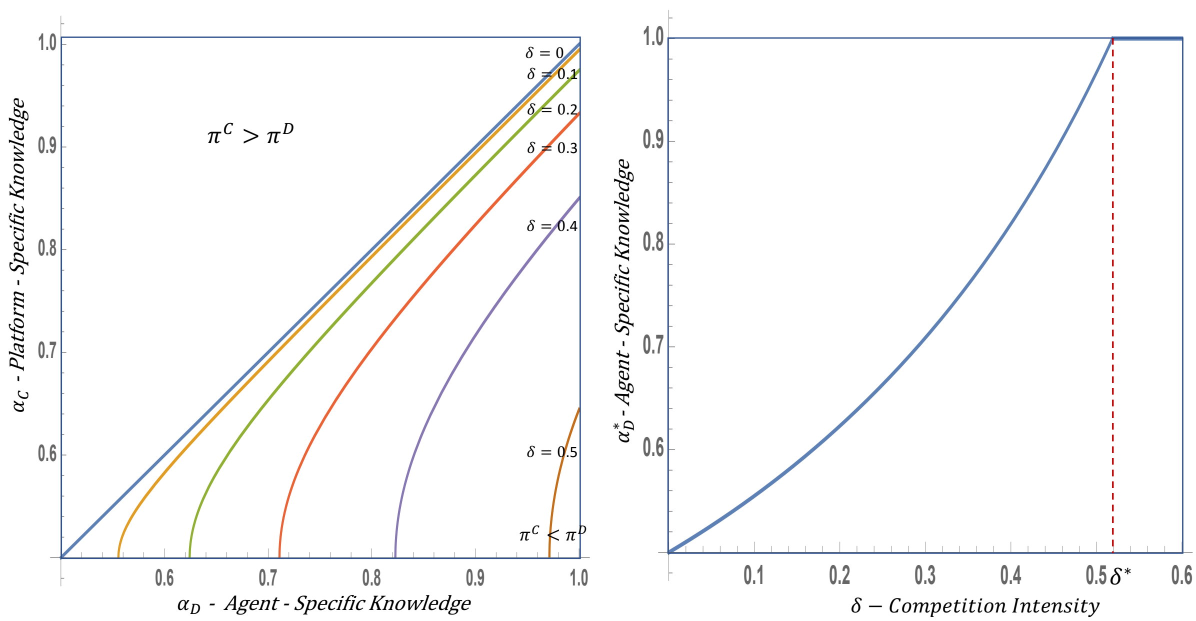

a) Indifference Curves (α_C, α_D) between centralized and decentralized CS.

b) Threshold Information Precision (α_D^*) under decentralized CS.

Figure 1: Impact of Platform and Agent-Specific Knowledge on the Performance of Centralized and Decentralized CS

When the precision of demand signals under centralized (α_C) and decentralized CS (α_D) differ due to the different information endowments of the platform and agents, centralized CS continues to dominate decentralized CS. When competition intensity is very low ($\delta = 0.1$ in Fig 1a), centralized CS dominates decentralized CS when the informational advantage of agents is not too high. Fig (1a) shows that as the intensity of competition (δ) increases, centralized CS rapidly dominates decentralized CS. Even for low levels of competition ($\delta = 0.3, 0.4$), the agent's informational advantage over the platform needs to be sufficiently high for decentralized CS to perform better than centralized CS. When competition is moderately high (for $\delta > 0.5$ in Fig 1a), decentralized CS completely cedes dominance to centralized.

Fig (1b) shows the threshold information precision (α_D^*) required for decentralized CS to dominate centralized CS. The figure shows that α_D^*, is increasing in δ. As the intensity of competition increases, decentralized CS requires higher information precision to compensate for the lack of coordination. When δ is sufficiently high (shown by δ^* in Fig (1b)), even with perfect information, decentralized CS underperforms compared to centralized CS, bolstering our result that information cannot match the superior coordination under centralized CS.

References

Hagiu, Andrei and Julian Wright (2019). "The optimality of ad valorem contracts". In: *Management Science* 65.11, pp. 5219–5233.

Hill, Dan (2015). "How much is your spare room worth?" In: *Ieee Spectrum* 52.9, pp. 32–58.

Huang, Yufeng (2022). "Pricing Frictions and Platform Remedies: The Case of Airbnb". In: *Available at SSRN 3767103*.