

Who Benefits from Platform Competition?

Evidence from Ride-Hailing*

Nan Wang[†]

October 12, 2026

Job Market Paper

Click [here](#) for the latest version.

Abstract

Competition among digital platforms may enhance efficiency by curtailing platforms' market power or reduce efficiency by diminishing network externalities. Furthermore, because platforms set fees on both sides jointly, the different effects of competition on market power and on network externalities can translate into opposite welfare effects across sides. I study the welfare effects of platform competition with an equilibrium model estimated on data covering the universe of ride-hailing transactions in a major Chinese city. The model incorporates platforms' matching and pricing, drivers' platform adoption and labor supply, and riders' trip requests. I find that concentration can raise efficiency, by up to 3.8% relative to the status quo. However, the impact of concentration differs substantially across the sides of the market. Platforms earn more profit in a more concentrated market, while drivers lose surplus, because concentration not only increases platforms' monopsony power but also raises drivers' productivity, which lowers platforms' demand for labor. Whether riders gain or lose depends crucially on the level of concentration. They gain from mild concentration but lose from severe concentration. Holding market structure fixed, promoting multi-homing improves welfare and benefits all sides of the market.

*I am grateful to Guillermo Marshall, Michael Sullivan, Ekaterina Khmel'nitskaya, and Chunhua Wu for their guidance and support. I also thank Scott Orr, Ron Yang, Erica Moszkowski, and participants at the UBC Sauder Seminar for helpful comments. All errors are my own.

[†]University of British Columbia, Sauder School of Business. Email: nan.wang@sauder.ubc.ca.

1 Introduction

Digital platforms organize markets as centralized matchmakers, connecting users on different sides, such as riders with drivers, buyers with sellers, and workers with employers. These markets exhibit network externalities, in that users benefit from a larger network, which matches them with more users on the other sides. Competition among platforms can therefore raise efficiency, by limiting their power to raise prices, or lower it, by fragmenting the market into separate networks and weakening network externalities. Given this trade-off, whether competition increases efficiency, and whether there is an optimal level of competition, are open questions.

Beyond total welfare, the distributional effect of platform competition has drawn more attention from competition authorities with the rise of gig work. Platforms in these markets intermediate directly between consumers and workers, so platform competition connects platforms' market power in the product market with their monopsony power in the labor market. The two have traditionally been analyzed separately, and the labor market effects of mergers received little attention from competition regulators until recently ([Naidu et al., 2018](#); [U.S. Department of Justice and Federal Trade Commission, 2023](#)). Because competition affects market power, monopsony power, and network externalities differently for consumers and for workers, how consumer surplus and worker surplus respond to the level of competition is a central question for the design of competition policy for platforms.

Regulators have also tried to promote multi-homing, the practice of joining and using multiple platforms simultaneously ([European Commission, 2022](#); [OECD, 2021](#)). This remedy is intended to strengthen network externalities by letting users reach the other side across platforms without market consolidation and thus without granting platforms additional market power. How this remedy performs, and whether it complements or substitutes for merger policy, has not been examined empirically.

This paper fills these gaps by developing and estimating an equilibrium model of platform competition under endogenous multi-homing in the Chinese ride-hailing industry. My empirical analysis uses a novel dataset that contains order-level information on the universe of ride-hailing transactions in a major city in China for three months. Using detailed order-level information, I can observe rich rider-side substitution patterns and driver labor supply patterns. The novelty of this data is that it contains a universal driver identifier, which allows me to back out drivers' working schedules and track their working behavior across platforms.

To disentangle the effect of opposing economic forces, I develop and estimate an equilibrium

model of a ride-hailing market where competing platforms set prices strategically, and use it to compare the status quo to simulated counterfactuals. The model focuses on the equilibrium play among three types of market participants, the platforms, riders, and drivers, who make decisions in sequence. In the first stage, competing platforms set their rider prices and driver payments strategically. Next, drivers decide which platform or platforms to adopt by comparing the expected earnings to a fixed cost of adoption. In the last stage, riders decide which platform to order from, and drivers decide whether to work and on which platform. The model captures network externalities through a matching system in which, on each platform, the number of searching drivers and riders determines the average pickup time and driver search time. Intuitively, with more drivers, riders are picked up faster, and with more riders, drivers receive more orders. The matching system also captures economies of density, in that with more users on both sides of the market, matching is more efficient. Importantly, platforms’ fee-setting internalizes riders’ and drivers’ responses to fees as well as the matching process.

Estimation proceeds in steps. First, I estimate riders’ preferences using a two-step procedure: a micro step recovers each platform’s mean utility in each market, and a market step recovers the parameters governing riders’ value of time. Second, I recover drivers’ costs and labor supply elasticities in a two-step procedure: I estimate platforms’ commission rates and drivers’ opportunity costs by matching observed single-homing drivers’ working probabilities across hours and platforms to what predicted by the model. Then, I estimate drivers’ multi-homing costs by maximum likelihood on the observed search times between orders. Third, I estimate the matching function, which gives an order’s pickup time as a function of the stock of searching drivers, by regressing observed pickup times on that stock. Last, I estimate the weights platforms place on rider and driver surplus relative to short-term profits by matching the observed prices and driver payments with those predicted by the model. These weights should be understood as reduced-form parameters that capture the platforms’ long-run considerations. Riders’ substitution between ride-hailing and alternative modes of transportation and drivers’ long-run occupational substitution elasticities are calibrated to match the experimental estimates of [Castillo \(2025\)](#); [Angrist et al. \(2021\)](#).

Identification of riders’ value of time is complicated by the endogeneity of platforms’ fees and pickup times, which reflect unobserved rider preferences. I address this with Hausman-style instruments, a platform’s price and pickup time on other routes in the same hour. These instruments exploit the across-platform variation in fees and pickup times that is driven by differences across platforms in driver availability, pricing, and matching algorithms, which are exogenous to riders’ preferences on a given route.

To study the welfare implications of platform competition for each side of the market, I use the estimated model to simulate the market equilibrium under changes in market structure that mimic platform mergers. Leveraging the high degree of competition in the observed market, I simulate the equilibrium at different levels of concentration. I find that market concentration can raise total welfare by up to ¥241 million (US\$35 million, 3.8%) annually.¹ However, more consolidation does not always generate higher efficiency. Total welfare peaks in a duopoly of roughly equal-sized rivals, and once the market becomes more concentrated, with one platform dominating, total welfare falls. To decompose the net change, I design two additional counterfactual studies. In the first, I simulate the market equilibrium in which the merged platforms keep matching their own users separately, to isolate the harm of higher market power. In the second, I simulate the equilibrium with fees held fixed, to isolate the gain from higher matching efficiency. I find that higher market power costs up to ¥124 million (US\$18 million, 2%) annually through higher rider prices and lower driver pay, while matching efficiency adds up to ¥364 million (US\$53 million) annually.

I find that there is no optimal market structure for everyone, because the three sides of the market experience substantially different impacts from platform consolidation. I decompose the two effects on platform profit, rider surplus, and driver surplus separately. Greater platform market power benefits the platforms and harms drivers and riders through lower driver payment and higher rider fare. Better matching efficiency alone benefits all sides of the market through the time it saves, but for drivers this gain does not survive the platforms' re-pricing in equilibrium. Intuitively, higher matching efficiency makes each driver more productive, which lowers the platforms' demand for labor. With monopsony power, platforms lower driver pay and ultimately transfer the matching efficiency gain from drivers to their own profit. Combining the two forces, platforms always gain and drivers always lose from market concentration. Whether riders gain or lose depends crucially on the relative strength of the two forces. Compared with the status quo, after a mild merger, consumer surplus rises by up to ¥71 million (US\$10 million) per year, but if the merger results in a duopoly, consumers are worse off by ¥94 million (US\$14 million) annually.

I examine the welfare effects of promoting multi-homing by simulating the market equilibrium under different levels of multi-homing cost. Making multi-homing cheaper raises social welfare. Cutting the hourly multi-homing costs by ¥20 (US\$2.9), roughly a quarter of their estimated level, generates an annual gain of ¥94 million (US\$14 million). A distinguishing feature of this policy is that it benefits all sides of the market, whereas changing market structure always

¹Throughout the paper, values in US dollars are converted at the average exchange rate over the sample period, January to March 2023, of ¥6.84 per dollar.

harms one side. Decomposing the source of the gain, I find that promoting multi-homing does little to improve matching. It raises welfare by inducing drivers to work hours on other platforms that would not be worthwhile at the current multi-homing costs. Two reasons explain why my empirical findings show multi-homing to be less consequential for network externalities than the theory predicts (Teh et al., 2023). First, unlike merchants, drivers can serve only one rider at a time, which limits their gain from multi-homing to the time saved in searching. Second, the estimated multi-homing cost is so large that the use of multi-homing remains limited after a marginal cost reduction.

These results carry several lessons. First, because of the tension between market power and network externalities, there is an intermediate optimal level of concentration, so whether a merger is beneficial depends on the level of concentration. Second, there is no optimal market structure for everyone, so any change in market structure benefits some participants and harms others. Platforms prefer concentration, while drivers prefer competition. The effect on rider surplus is non-monotone, and the net effect depends on the balance between matching efficiency and market power. Third, promoting multi-homing improves welfare and benefits all sides of the market, and is therefore a good complement to merger policy. Merging platforms while making multi-homing cheaper can improve total welfare and benefit all sides.

1.1 Related Literature

This paper contributes to the scarce but growing empirical literature on platform competition. Early work in media markets primarily focused on estimating network effects and demand for two-sided platforms, rather than equilibrium behavior among strategic market participants (Rysman, 2004; Kaiser and Wright, 2006). Among more recent empirical studies, this paper is methodologically most closely related to Rosaia (2025), Sullivan (2026), and Wang (2025). Despite differences in context and modeling details, these papers share a common structure: multiple two-sided platforms set prices strategically, while users on both sides choose participation in response to each other’s presence and to platform choices, with equilibrium outcomes jointly determined by platform and user decisions. Rosaia (2025), in ride-hailing, and Wang (2025), in payment cards, both study the effect of a change in market structure on total welfare. Sullivan (2026) studies online food delivery platforms to quantify inefficiencies in platform fees and the welfare implications of commission regulation. I contribute to this literature in two ways. First, leveraging a more competitive setting, I document empirically that the effect of market structure on user surplus can be non-monotone due to the tension between market power and network efficiency. Second, I incorporate

endogenous and costly multi-homing and quantify empirically the welfare implications of different levels of multi-homing frictions.

This paper also contributes to the literature on competition in labor markets and its implications for antitrust. With the rise of the gig economy, many digital platforms connect consumers directly to workers, but antitrust authorities have historically evaluated mergers by their effect on consumers alone. A growing body of evidence shows that concentration also harms workers. Using cross-market comparisons (Azar et al., 2022; Benmelech et al., 2022) or completed mergers (Prager and Schmitt, 2021; Arnold, 2026), researchers have found that market concentration lowers workers' wages. These findings support the argument that antitrust authorities should also consider the monopsonistic effects of mergers (Naidu et al., 2018; Marinescu and Hovenkamp, 2019). However, these papers use reduced-form methods that focus on the labor market and allow little interaction with the competition effect in the product market, which makes their results hard to translate to platform markets, where the interaction across sides is strong. Careful modeling of the demand side and of the interaction between the sides is thus required to understand the effect of platform competition. My paper fills this gap with a structural model that incorporates the decisions of all market participants, and it is, to my knowledge, the first to evaluate the effect of competition on the labor market jointly with its effect on the product market in a platform setting.

This paper is also related to a growing body of research on spatial models of two-sided on-demand transportation networks, spanning applications such as taxis (Fr chet te et al., 2019; Buchholz, 2022), ride-hailing (Shapiro, 2018; Castillo, 2025), and shipping (Brancaccio et al., 2023). This literature is characterized by search frictions in the matching process and strategic interactions between users on the two sides of the market. Although I do not incorporate explicit spatial heterogeneity in this paper, search frictions are modeled in a closely related way. In particular, I highlight how search frictions shape platform competition by mediating the impact of capacity, prices, and multi-homing on equilibrium matching and welfare.

The welfare implications of platform competition and multi-homing are central topics in the theoretical literature on two-sided platforms. Early work highlights the importance of network effects and of the connection between the two sides in determining the equilibrium distribution of welfare (Rochet and Tirole, 2003; Caillaud and Jullien, 2003; Armstrong, 2006), and later work shows that market concentration can benefit one side and harm the other (Anderson and Peitz, 2020; Tan and Zhou, 2021; Correia-da Silva et al., 2019). A large theoretical literature also shows that multi-homing is a key determinant of equilibrium

prices and welfare under platform competition (e.g., [Armstrong and Wright, 2007](#); [Bakos and Halaburda, 2020](#); [Teh et al., 2023](#)). These papers abstract from the matching technology that generates economies of density, which drives my central finding. I formalize multi-homing costs in a model in the spirit of this literature and quantify empirically their implications for equilibrium outcomes and welfare.

2 Industry Background and Data

2.1 Ride-Hailing in a Major Chinese City

Ride-hailing platforms carry a substantial share of urban passenger transportation in China. In February 2023 alone, 625 million ride-hailing orders were completed nationwide, and roughly one third of trips by car are made through ride-hailing services. These platforms operate much like their U.S. counterparts, such as Uber and Lyft: a rider requests a trip through a mobile application, and the platform dispatches the order to a nearby driver. The fee structure is likewise similar. The rider pays a fare to the platform, the platform retains a share as commission, and the remainder is paid to the driver. Both the rider fare and the commission rate are set by the platform.

The Chinese ride-hailing industry differs from its U.S. counterpart in the intensity of platform competition. Whereas the U.S. market is dominated by two firms, platform competition in China in 2023 was intense. As of February 28, 2023, there were 303 registered ride-hailing platforms nationwide, comprising seven major platforms and a large number of small, local fringe competitors. Three factors account for this degree of competition. First, 2023 was the first year after the COVID-19 pandemic, and the associated surge in travel demand attracted substantial investment into the industry. Second, the previously dominant platform, DiDi, was under investigation by the Chinese government over data privacy and security concerns. New-user registration on DiDi was suspended from July 2021 until January 2023, opening a window for rival platforms to enter and expand. Third, the high population density of major Chinese cities generates enough ride demand and driver supply for several platforms to coexist profitably.

No regulation or platform rule prevents riders or drivers from using multiple platforms. Among drivers, however, it is widely believed that platforms maintain implicit rules that discourage multi-homing. [Chen et al. \(2024\)](#) document one such practice: platforms favor continuously active drivers in order dispatch, which penalizes drivers who divide their time across platforms. These implicit costs help explain why drivers multi-home less than the low

explicit cost of doing so would suggest. I return to this point in Section 2.3.

2.2 Data

My empirical study uses administrative transaction-level data on the universe of completed ride-hailing trips in a major city in China over the eight weeks from January 30 to March 26, 2023. The city’s high demand for transportation and competitive ride-hailing market make it a representative setting in which to study competition in the ride-hailing industry. Platforms typically offer several service tiers, analogous to UberX and Uber Lux. In this study, I focus on the standard tier of each platform and exclude the premium tiers from the sample.

Order-level information The raw order-level data contain rich spatial and temporal information. On the spatial side, I observe the GPS coordinates of the origin and destination of each trip, together with the place names the rider entered when requesting the ride. The place names are highly informative about the purpose of a trip beyond what the coordinates reveal: for two trips ending at the same building, for instance, they distinguish a rider visiting a store from one commuting to a firm located there.

On the temporal side, I observe three timestamps for each order: when the rider placed the request, when the driver picked up the rider, and when the trip ended. From these I compute the pickup time, the time elapsed between the placement of the request and the pickup. One caveat is that I do not observe how long it takes for a request to be matched to a driver, so the matching delay is included in the measured pickup time. This is a potential source of upward bias in pickup time, but it is unlikely to be a serious concern in the Chinese ride-hailing market, where matching typically takes only seconds.

In addition to spatial and temporal information, I observe other trip characteristics, such as the trip distance and the rider fare. A further caveat is that I do not directly observe the commission rate or the driver’s take-home pay. I infer driver payment in the estimation of the model.

Construction of driver schedule A novel feature of the data is a driver identifier that is common across platforms. Because it is an encoded version of the driver’s registration number rather than a platform account, the identifier follows the driver across every platform the driver works on. It is informative in two ways. First, it allows me to assign orders to individual drivers and thereby construct each driver’s working schedule. Although I do not observe when a driver goes online or offline, I observe the gap between consecutive orders, which effectively measures the driver’s idle, or search, time. Second, because the identifier is

common across platforms, I can track a driver’s activity on every platform. This information is essential for understanding drivers’ ability to receive orders from multiple platforms, but it is not a perfect measure of actual multi-homing behavior: since I observe only completed orders, a driver may be searching on several apps while receiving orders from only one platform.

Variation across platforms Table 2.1 presents summary statistics that show the competitive landscape and the variation across platforms. In total, 15 platforms appear in the data, six of which have a market share above 3%. The empirical analysis focuses on these six. Platform names are masked in the data. I refer to the six as platforms A to F, in descending order of market share.

Among the six, platform A is the dominant platform, with the largest market share and the largest driver pool. Platforms B and C are the main challengers, and platforms D, E, and F are fringe competitors. A trip typically costs a passenger ¥21 to ¥24 (US\$3.1 to US\$3.5), with platform E an outlier at ¥29 (US\$4.2). Pickup times and driver search times vary substantially across platforms. The last column reports the number of drivers working on each platform, which varies widely because platforms differ in the extent of multi-homing and in how intensively their drivers work: platform E, for instance, has nearly as many drivers as platform A while carrying an eighth of its orders. Platform B serves most of its orders with exclusive drivers, whereas platforms A and C serve a large share of their orders with multi-homing drivers; Section 2.3 examines these patterns in detail.

Table 2.1: Summary statistics by platform

Platform	Orders (thousands)	Market share (%)	Average fare (¥)	Pickup time (min)	Search time (min)	Drivers (count)
A	6,187	33.7	22.5	4.3	8.7	24,896
B	5,214	28.4	21.3	5.1	5.1	9,159
C	3,247	17.7	24.0	6.3	5.8	20,814
D	886	4.8	23.4	5.8	5.9	4,659
E	794	4.3	28.8	7.3	9.0	22,257
F	664	3.6	23.7	5.3	6.1	3,894
Other platforms	406	2.2	34.1	6.6	7.4	7,912
Market	18,381	100.0	24.1	5.3	6.5	65,471

Notes: Completed ride-hailing trips in the city from January 30 to March 26, 2023. Platform rows are restricted to each platform’s standard service tier; *Other platforms* pools the minor platforms not analyzed in this paper, and *Market* covers all completed orders on all platforms and service tiers. *Drivers* is the number of distinct driver identifiers with at least one order in the row’s set of trips; the market total is below the sum over platforms because drivers multi-home.




Figure 2.1: Weekly pattern in order demand, driver labor supply and price

Notes: This figure shows the average number of orders, working drivers and price over a week. Upper panel shows the hourly pattern of price per minute. Lower panel shows the hourly pattern of orders and the number of working drivers in the hour.

Time pattern across week hours Figure 2.1 shows how market conditions vary over the week. Several patterns stand out. First, order volume follows a standard commuting pattern: demand is high during the day and low at night, with pronounced morning and evening peaks on weekdays. Second, the number of working drivers follows the same pattern as order volume, reflecting the response of drivers' labor supply to demand. Finally, the per-minute rider price moves in the opposite direction from driver supply: it is lowest during the daytime and evening peaks and nearly doubles in the late-night hours before dawn.

2.3 Driver Working Behavior and Multi-Homing

Drivers vary substantially in their working patterns, both in total hours and in schedules. Backing out each driver's working schedule from the order-level records, I report the distribution of weekly working hours in Figure 2.2. As in other gig-economy industries, the ride-hailing market has a large group of occasional drivers who most likely drive as a part-time job: the mean is 22.5 hours per week, while the median is only 12.5.




Figure 2.2: Driver working pattern

Notes: This figure shows drivers' working patterns from two perspectives. Panel (a) shows the distribution of drivers' average working hours per week. Panel (b) shows the working schedules of the different driver types.

Based on drivers' working patterns across weekday hours, I classify drivers into four types using k-means clustering. Types 1, 2, and 3 — which account for 11.4%, 14.0%, and 4.3% of drivers and for 28.6%, 33.8%, and 7.3% of total working hours — are referred to as full-time drivers. Types 1 and 2 work mostly during the day, with type 1 concentrated in the morning and type 2 in the evening; type 3 are the night-owl drivers who work late at night. The remaining 70.3% of drivers, who account for 30.3% of working hours, are referred to as occasional drivers.

Multi-homing is widespread among drivers. Figure 2.3 shows the distribution of drivers across platform combinations, where a driver's adopted set is defined as the platforms carrying more than 1% of the driver's orders. Although the largest groups are exclusive drivers on platforms A, E, and C, one quarter of drivers (24.3%) adopt two or more platforms, and 7.2% adopt three or more. Weighted by activity, multi-homing is even more prevalent, as multi-homing drivers deliver 40.5% of all orders. The most common combinations pair platform A with E or C; the A–C pair alone delivers 6.9% of all orders, more than platforms D, E, or F each deliver through their exclusive drivers.

Despite this widespread adoption, drivers' actual use of multiple platforms is limited. Most multi-homing drivers concentrate their activity on one platform: the median multi-homing driver completes 84% of orders on the primary platform and only 14% on the secondary one;




Figure 2.3: Distribution of driver platform adoption

Notes: This figure shows drivers' platform adoption. Panel (a) reports the share of drivers choosing each option and panel (b) the share of orders accounted for by those drivers. The figure shows the 15 most popular single-, two-, and three-platform options, which together account for 91.8% of drivers and 85.0% of orders; the remaining 48 options (including all options with four or more platforms) account for 8.2% of drivers and 15.0% of orders.

58% of multi-homing drivers complete at least 80% of their orders on the primary platform, and only 16% split their orders more evenly, with a primary share below 60%. In aggregate, orders delivered on a platform other than the driver’s primary one account for 16.9% of multi-homing drivers’ orders and just 6.9% of all orders in the market. Adopting a second platform is thus common, but the second platform typically serves as a marginal supplement rather than an equal alternative. This gap between adoption and use points to substantial frictions in working across platforms simultaneously — the multi-homing cost that the model in Section 3 makes explicit and that the estimation quantifies.

3 Model

3.1 Summary

I develop a model of platform competition to analyze empirically the welfare implications of the economic forces at play in equilibrium. The model focuses on the strategic interaction among riders, drivers, and platforms. Time is discrete and indexed by the hour of the week, $t \in \{1, \dots, 168\}$, from Monday 0:00 to Sunday 23:00. I study a stationary equilibrium of a representative week, in which all variables are constant within an hour and the week repeats identically.

The model captures four features of ride-hailing markets. First, the market is two-sided. Riders and drivers each benefit from greater participation by the other side. Second, matching exhibits network efficiency and economies of density. The more riders and drivers use a platform, the more efficiently it matches them. Third, riders and drivers can choose among platforms, but they have platform-specific preferences, which differentiate the platforms and give each some market power over its users. Fourth, drivers can adopt several platforms and use more than one at a time, at a cost, so the extent of multi-homing is an equilibrium outcome rather than a primitive.

The model has three stages. In the first stage, platforms set rider prices and driver payments strategically to maximize their objectives, modeled as a weighted sum of profit, rider surplus, and driver surplus, which captures platforms’ long-run considerations in a reduced-form way. In the second stage, drivers choose whether to enter ride-hailing and, if so, which platform or combination of platforms to adopt. In the third stage, in each hour of the week, drivers decide whether to work and on which platform(s), and riders decide whether to request a trip and from which platform. Each platform matches its working drivers to its trip requests, which clears the hourly market and pins down drivers’ search time and riders’ pickup time.

An underlying assumption is that adoption is costly for drivers but free for riders. In practice, joining a platform as a driver involves identity verification and a vehicle inspection, whereas riders can install and use an app with minimal hassle. Riders therefore choose among all platforms trip by trip, while drivers’ platform choices are made in advance and constrain their hourly options.

Although the model captures many features of the ride-hailing industry, it abstracts from others for tractability. First, the model is static. It therefore does not capture platforms’ dynamic pricing incentives, such as state-dependent or personalized pricing, which are believed to be used in the industry. Nor does it capture drivers’ dynamic scheduling within a shift, such as expectations of later earnings or cumulative fatigue. The model absorbs the systematic part of these considerations into drivers’ hour-specific outside options. Drivers therefore still respond to changes in hourly earnings in the counterfactuals, but the dynamic considerations behind their schedules are held fixed rather than re-optimized. Second, decisions are aggregated. Drivers make their choices hour by hour without regard to location, and platforms choose the overall level of their fees while the hourly pattern of fees is taken as given. The model therefore does not speak to heterogeneity in policy effects across locations or across hours of the week. Its focus is heterogeneity across platforms and across the three sides of the market, which is the priority of this paper.

The remainder of this section describes the three stages in reverse order, starting from the hourly market.

3.2 The Hourly Market

In each hour of the week, riders request trips, drivers supply labor, and platforms match the two. This subsection describes the three in turn, taking platforms’ fees and drivers’ adoption decisions as given.

3.2.1 Rider Trip Request

Rider i in hour t at origin o contemplates a trip to destination d . I denote a market by $m = (o, d, t)$, a route at an hour of the week, which is the level at which rider demand is defined. This structure assumes that riders substitute across platforms, but not across routes or hours. Each rider demands at most one trip and exits the market after choosing, so riders are in one-to-one correspondence with trips, and I use i to index both. The rider chooses whether to request a trip from a platform k or to take the outside option, $k = 0$, which stands for other modes of transportation or forgoing the trip. The utility of requesting a trip from

platform k , relative to the outside option, is

$$u_{imk} = \alpha_k - \beta^p p_m^k - \beta^w w_m^k + Z_i' \gamma^k + \xi_m^k \quad (1)$$

where α_k is the rider's brand preference for platform k , p_m^k is the per-minute price that platform k posts in market m , and w_m^k is platform k 's pickup time in market m . Pickup time does not depend on the destination, so $w_m^k = w_{ot}^k$. The use of a per-minute price assumes that riders respond to the price level in the same way regardless of trip duration. Z_i is a vector of trip characteristics that shift platform preferences, and ξ_m^k is a market-specific platform preference. The key parameters are the price and pickup-time coefficients β^p and β^w , which determine the price and pickup-time elasticities of demand and, through their ratio, riders' value of time. The utility of the outside option is normalized to 0.

Platform k 's weekly rider surplus, measured in ¥ per week, is the area under its demand curve above its price, summed across markets. Formally,

$$CS^k = \sum_m r_m^k \int_{p_m^k}^{\infty} q_m^k(x, p_m^{-k}, w_m^k, w_m^{-k}) dx,$$

where r_m^k is the average trip duration on platform k in market m , which converts the per-minute price into a per-trip price so that the surplus is measured in ¥.

3.2.2 Matching and Dispatch Technology

Driver dispatch. In each hour t , platform k dispatches its orders to the drivers working on it. I assume that a driver's search time between orders follows an exponential distribution with mean s_t^k hours. Equivalently, while a driver is searching, orders arrive according to a Poisson process with rate $(s_t^k)^{-1}$. Two features of this dispatch assumption are worth noting. First, dispatch is aggregated to the hour with no spatial dimension, so the arrival rate is the same for every driver searching on the platform in that hour. Second, drivers take s_t^k as given, since no individual driver can influence it, but it is determined in equilibrium. I return to its determination in Section 3.2.2.

I further assume that dispatch on different platforms is independent. If a driver is searching on two platforms k and l at the same time, orders arrive from the two platforms according to independent Poisson processes with rates $(s_t^k)^{-1}$ and $(s_t^l)^{-1}$. The driver takes whichever order arrives first, so the expected search time is

$$s_t^M = \frac{1}{(s_t^k)^{-1} + (s_t^l)^{-1}} < \min\{s_t^k, s_t^l\}$$

Intuitively, by searching on several platforms at once, a driver shortens the average search time, which is the sole benefit of multi-homing in the model. The underlying assumption is that multi-homing drivers do not select among orders across platforms but accept the first order dispatched to them. This is arguably realistic, for two reasons. First, drivers do not know whether the next order will be better than the current one. Second, most ride-hailing platforms penalize drivers for declining dispatched orders.

Matching. Each hour, platform k matches trip requests to the drivers working on it. The matching process determines drivers' search time s_t^k and riders' pickup time w_t^k , which both sides take as given.

Search time follows from time accounting. In the steady state, the time drivers supply to the platform equals the time required to serve all requests. Therefore,

$$\underbrace{\overbrace{I_t^k}^{\text{search}} + \overbrace{w_t^k Q_t^k}^{\text{pickup}} + \overbrace{r_t^k Q_t^k}^{\text{trip}}}_{\text{time driver supplied}} = \underbrace{(s_t^k + w_t^k + r_t^k) Q_t^k}_{\text{time required by requests}}.$$

where I_t^k is the total time drivers spend searching on platform k in hour t . The left-hand side adds up drivers' time in its three uses, searching, picking up, and driving. The right-hand side counts the same time request by request, as each of the Q_t^k requests occupies a driver for the expected search, pickup, and trip time. Cancelling terms on both sides gives

$$s_t^k = \frac{I_t^k}{Q_t^k}. \quad (2)$$

Search time is the ratio of the searching stock to the request flow. It measures congestion on the driver side, since drivers wait longer for an order when more of them compete for the same requests.

Turning to pickup time, describing it requires a word on what matching friction means on a platform. Matching on a ride-hailing platform differs from matching in traditional transportation markets, such as street-hailed taxis or bulk shipping, because it is centralized. In a decentralized market, agents on the two sides meet stochastically. The matching friction then takes the form of lost matches, as fewer matches are realized in each period than would be feasible. On a ride-hailing platform, by contrast, the platform dispatches every request as it arrives. A rider who requests a trip is matched immediately with the nearest available driver, so the number of realized matches equals the number of requests. The friction therefore takes a different form. It is not a loss in the number of matches but a loss in their quality,

measured by the pickup time, the time the assigned driver needs to reach the rider.

I describe this friction with a negative relationship between the pickup time w_t^k and the stock of idle drivers I_t^k :

$$w_t^k = w_t^k(I_t^k), \quad (3)$$

Intuitively, I_t^k also measures the average number of drivers searching at any moment during the hour. With more drivers searching, the nearest available driver is more likely to be close to a rider who requests a trip, so the pickup takes less time. This reduced-form description of search friction has been used in recent research on centralized matching markets (Castillo et al., 2025; Rosaia, 2025). Theoretically, it is backed by the analytical solution for the expected distance to the nearest driver in a spatial queueing model (Besbes et al., 2022), and there is empirical evidence for the relationship (Castillo et al., 2025).

Two features of this technology matter for the analysis. First, labor supply and demand affect matching through I_t^k . Holding demand fixed, more working drivers spend more time searching and shorten pickups. Holding labor supply fixed, more requests draw drivers into pickups and trips, which leaves less time for searching and lengthens pickups. Second, the two waits respond differently to the scale of the market. By Equation (2), a proportional increase in the number of searching drivers and the number of requests leaves search time unchanged, while by Equation (3) the larger number of searching drivers shortens pickups. Pooling more riders and drivers on one platform therefore makes matching more efficient without lengthening drivers' searches. These are the economies of density that make concentration potentially welfare-improving.

3.2.3 Drivers' Labor Supply

Each hour, drivers decide whether to work and on which platform. By the time they make this decision, platforms have set their per-minute rider fares $\{p_t^k\}$ and driver payments $\{e_t^k\}$, and drivers have made their adoption decisions. For tractability, a driver adopts at most two platforms, which covers nearly all drivers in practice (Section 2.3). I denote driver j 's adopted combination by the ordered pair $\mathcal{G}_j = (k_g, l_g)$, where k_g is the driver's primary platform and l_g the secondary one, with $l_g = 0$ for a driver who adopts a single platform. The order matters, since a driver who uses platform A as the main platform and B as a backup differs from one who uses B as the main platform and A as a backup. To simplify notation, in this section I drop the driver subscript j . Drivers take the market conditions $\{s_t^k, w_t^k, r_t^k\}$ as

given, namely each platform's search time, pickup time, and trip duration, which are constant within the hour. Drivers belong to a finite set of types d , which differ in their opportunity cost of working.

I model drivers' labor supply decision as discrete choice. For a driver of type d who adopted a single platform k , the decision is binary, to work on the platform or not. A driver who stays inactive enjoys the outside option c_{dt} , the type's opportunity cost of working in hour t . A driver who works goes through cycles of search, pickup, and delivery without interruption and is paid only for the delivery period, so the expected number of orders completed in the hour is

$$n_t^k = \frac{1}{s_t^k + w_t^k + r_t^k}.$$

The payoffs of the two choices are

$$v_{gdt} = \begin{cases} e_t^k r_t^k n_t^k + \epsilon_{1t} & \text{if working,} \\ c_{dt} + \epsilon_{0t} & \text{if inactive,} \end{cases} \quad (4)$$

where $e_t^k r_t^k$ is the payment per order, the payment rate times the paid driving time per order, and ϵ_{1t} and ϵ_{0t} are idiosyncratic shocks to the two options, independently drawn from a type-I extreme value (T1EV) distribution with scale parameter σ_1 .

For a driver of type d who adopted a platform pair $\mathcal{G}_g = (k, l)$, the choice is among four options, working on platform k , working on platform l , working on both platforms simultaneously, and staying inactive. The expected earnings from working on a single platform follow the same logic as for a single-homing driver. In addition, working on the secondary platform carries a switching cost γ_k , which is specific to the main platform k rather than to the pair. For example, a driver who adopted (A, B) pays the same cost when working on the secondary platform as a driver who adopted (A, C) . This cost admits two interpretations. It may be a pure brand preference, reflecting for instance the driver's familiarity with the main platform's app. It may also be a cost that the main platform imposes to discourage multi-homing, such as an algorithmic preference for exclusive drivers or bonuses for continuous work that a driver who leaves for another platform forgoes.

When a driver searches on both platforms simultaneously, the next order can come from either platform, which shortens the expected search time to s_t^M as discussed in Section 3.2.2. The driver's earnings depend on which platform the next order comes from. The probability

that it comes from platform k is the platform's share of the combined arrival rate,

$$\omega_k = \frac{(s_t^k)^{-1}}{(s_t^k)^{-1} + (s_t^l)^{-1}},$$

and $\omega_l = 1 - \omega_k$. A driver working on both platforms therefore expects a pickup time of $\omega_k w_t^k + \omega_l w_t^l$ and a trip duration of $\omega_k r_t^k + \omega_l r_t^l$ on each order, and completes

$$n_t^M = \frac{1}{s_t^M + \omega_k w_t^k + \omega_l w_t^l + \omega_k r_t^k + \omega_l r_t^l}$$

orders per hour, each paying $\pi_t^M = \omega_k e_t^k r_t^k + \omega_l e_t^l r_t^l$ in expectation. Working on both platforms also carries a multi-homing cost γ_M , common to all platform pairs, which captures the costs of simultaneous use that are not platform-specific, such as the cognitive cost of managing two devices or accounts at the same time.

Altogether, the payoffs of the four options are

$$v_{gdt} = \begin{cases} e_t^k r_t^k n_t^k + \eta_{kt} & \text{if working on platform } k, \\ e_t^l r_t^l n_t^l - \gamma_k + \eta_{lt} & \text{if working on platform } l, \\ \pi_t^M n_t^M - \gamma_M + \eta_{Mt} & \text{if working on both,} \\ c_{dt} + \eta_{0t} & \text{if inactive,} \end{cases}$$

where the η 's are idiosyncratic shocks to the four options, independently drawn from a T1EV distribution with scale parameter σ_2 . Multi-homing thus benefits a driver only through the shorter search time. Whether a driver multi-homes in a given hour depends on how this benefit compares with the costs γ_k and γ_M , so the extent of multi-homing responds to market conditions and to platforms' fees.

3.3 Platform Adoption

Before the hourly market opens, drivers simultaneously make platform adoption decisions, in a static game of incomplete information in the spirit of [Seim \(2006\)](#) and [Sullivan \(2026\)](#). There are in total N_D drivers. Each driver chooses one of the G platform combinations, a single platform or an ordered pair of platforms, or the outside option $g = 0$ of exiting the industry.

Driver j 's expected weekly payoff from combination g is

$$V_{jg}(N) = \begin{cases} \mathbb{E}[\sum_t \max v_{d(j)gt} \mid N] - \kappa_g + \xi_{jg} & \text{if } g \neq 0, \\ V_0 + \xi_{j0} & \text{if } g = 0. \end{cases}$$

The expectation term is the value of the hourly decisions that combination g makes available, summed over the hours of the week. In each hour the driver, of type $d(j)$, picks the best working mode among those described in Section 3.2.3, and the expectation is over the hourly idiosyncratic shocks. This value depends on the adoption choices of the other drivers, summarized by $N = (N_g)_g$, the number of drivers adopting each combination, because market conditions do. κ_g is the adoption cost of combination g , common to all drivers, and ξ_{jg} is driver j 's idiosyncratic taste for the combination. V_0 is the value of the outside option, such as the payoff from another occupation.

The taste shocks ξ follow a nested logit structure with nest parameter σ_3 , in which the inside combinations share one nest and the outside option stands alone. The two levels of the nest correspond to two margins. Within the nest, drivers who enter sort across platforms and platform pairs according to the value of each combination net of its adoption cost. At the top level, potential drivers decide whether to drive at all, comparing the value of the best available combination with the outside option. The adoption stage therefore also captures drivers' entry into the occupation, and the nest parameter governs the elasticity of this entry margin, that is, how strongly the number of drivers responds to the value of driving relative to other occupations.

3.4 Platform Objectives and Fee Setting

Before the week begins, platforms set their fees. Platform k 's fee schedule consists of a rider price p_t^k and a driver payment e_t^k for each hour of the week, both per unit of trip time, so the platform earns the margin $p_t^k - e_t^k$ on every unit of trip time it intermediates. With trip demand Q_t^k and trip duration r_t^k , platform k 's weekly profit is

$$\Pi^k = \sum_t Q_t^k r_t^k (p_t^k - e_t^k).$$

I abstract from any marginal cost of matching.

Platforms may care about more than current profit. I model platform k 's objective as a

weighted sum of its profit, the surplus of its riders, and the surplus of its drivers,

$$O^k = \Pi^k + \lambda_R^k CS^k + \lambda_D^k DS^k,$$

where CS^k is the rider surplus defined in Section 3.2.1, DS^k is the expected surplus of the drivers who adopt platform k , alone or as part of a pair, and $\lambda_R^k, \lambda_D^k \geq 0$ are platform-specific weights. Pure profit maximization is the special case $\lambda_R^k = \lambda_D^k = 0$. The weights are a reduced-form way to capture the long-run considerations that a static model cannot represent directly. A platform that keeps fares low or pay high today builds its user base, its reputation, and its network for the future. In a static model, these motives appear as a willingness to give up current profit for the surplus of its users, and the weights measure how much of that surplus the platform internalizes. This device is common in the empirical literature on platform competition (Castillo, 2025; Rosaia, 2025), and the weights are best read as conduct parameters that summarize how far a platform's pricing departs from static profit maximization. Allowing the two weights to differ lets a platform favor one side of the market over the other, for instance by subsidizing drivers to secure supply while charging riders.

Platform k 's fees consist of two parts, a fixed weekly fee schedule and a level shift chosen by the platform,

$$p_t^k = p^k + \bar{p}_t^k, \quad e_t^k = e^k + \bar{e}_t^k.$$

The schedules $\bar{p}_t^k$ and $\bar{e}_t^k$ vary over the hours of the week and capture the hourly variation in riders' and drivers' behavior, and I take them as exogenously determined. The levels p^k and e^k are the platform's strategic choices. This two-part structure simplifies the platform's pricing problem to a level change. In practice, platforms adjust prices at a much finer granularity, down to order-by-order surge pricing. The model keeps the hourly shape of fees fixed and lets platforms compete through the overall level of what they charge riders and what they pay drivers. Each platform chooses its levels to maximize its objective, taking the other platforms' fees as given and anticipating drivers' adoption decisions and the hourly market outcomes that follow. The resulting equilibrium in fee levels is stated in Section 3.5.

3.5 Equilibrium

The equilibrium concept is subgame perfection. Platforms anticipate how drivers adopt and how the hourly market clears, drivers adopt anticipating the hourly market, and every stage is in equilibrium given the outcomes of the stages that follow. I define the equilibrium of each stage in the order the model is described, starting from the hourly market.

Hourly market equilibrium. Fix the fees and the adoption profile N . In hour t , riders' choices generate platform k 's trip demand, $Q_t^k = \sum_m q_m^k$, summing over the markets m in that hour.

Drivers' choices generate the searching stock, and it is worth being explicit about how much search time each driver contributes. A driver who works on platform k alone completes n_t^k order cycles in the hour, each beginning with a search of expected length s_t^k , so the driver spends a fraction $n_t^k s_t^k$ of the hour searching, all of it on platform k . A driver who works on both platforms of a pair spends a fraction $n_t^M s_t^M$ of the hour searching, and this search time counts toward the searching stock of *both* platforms. An idle driver with two apps open is available to be dispatched by either platform, so the same idle minute thickens two matching markets at once. Summing over drivers,

$$I_t^k = \sum_{g \ni k} N_g \left[P_{gt}(k) n_t^k s_t^k + P_{gt}(M) n_t^M s_t^M \right],$$

where the sum runs over the combinations that include platform k , $P_{gt}(k)$ and $P_{gt}(M)$ are the probabilities that a combination- g driver works on platform k alone and on both platforms, implied by the hourly payoffs of Section 3.2.3, and I suppress the average over driver types to ease notation.

It might seem counterintuitive that a multi-homing driver contributes search time to both platforms, but this structure is what reveals the efficiency of multi-homing. Multi-homing not only benefits the driver by shortening the driver's own search time but also reduces the fragmentation of the market. Drivers can be matched with a larger pool of riders, which generates better matches for the whole system and reduces each platform's wait times marginally.

A simple comparison makes the point. Consider a market with two identical platforms with market conditions (s, w, r) and a fixed population of D working drivers, dropping the time subscripts. Under full single-homing, the drivers split evenly and each platform's searching stock is $I_{SH} = (D/2) \cdot s/(s + w + r)$. Under full multi-homing, every driver searches on both platforms with the combined search time $s/2$, so each platform's stock is $I_{MH} = D \cdot (s/2)/(s/2 + w + r)$, larger than I_{SH} by the factor $(s + w + r)/(s/2 + w + r) > 1$. In the new equilibrium, through the pickup-time function, both platforms' riders are picked up faster with the same drivers.

The equilibrium is defined as a fixed point. Let $\mathbf{s}_t = (s_t^k)_k$ and $\mathbf{w}_t = (w_t^k)_k$ collect the hour's market conditions across platforms. In one direction, market conditions determine

the demand $Q_t^k(\mathbf{w}_t)$ and the searching stock $I_t^k(\mathbf{s}_t, \mathbf{w}_t)$ through the rider and driver models, respectively. In the other direction, the demand and the searching stocks pin down the search and pickup times through the dispatch identity Equation (2) and the pickup-time function Equation (3). An hourly market equilibrium is a profile of market conditions that the loop maps back into itself,

$$s_t^k = \frac{I_t^k(\mathbf{s}_t, \mathbf{w}_t)}{Q_t^k(\mathbf{w}_t)}, \quad w_t^k = w_t^k(I_t^k(\mathbf{s}_t, \mathbf{w}_t)), \quad \text{for every platform } k.$$

The fixed point determines both sides' waiting times jointly, which is how the model captures the two-sidedness of the market. More riders shorten drivers' searches, more drivers shorten riders' pickups, and each side congests itself.

Adoption equilibrium. Fix the fees. When drivers evaluate the combinations at the hourly market equilibrium induced by an adoption profile N , an adoption equilibrium is a profile $N = (N_g)_g$ that reproduces itself. The number adopting each combination equals the number of potential drivers for whom that combination delivers the highest adoption value,

$$N_g = N_D \cdot \Pr(V_{jg}(N) \geq V_{jg'}(N) \text{ for all } g') \quad \text{for every combination } g,$$

where the probability is taken over the taste shocks ξ_j and g' runs over all combinations and the outside option. Drivers' expectations about market conditions, and hence about the value of each combination, are thus consistent with the adoption choices they generate.

Fee-setting equilibrium. Anticipating the adoption equilibrium and the hourly market equilibrium that any fees induce, platform k 's objective is a function of the fee levels alone, $O^k(p^k, e^k; p^{-k}, e^{-k})$. An equilibrium of the full model is a profile of fee levels in which every platform best responds to the others,

$$(p^k, e^k) \in \arg \max_{(p, e)} O^k(p, e; p^{-k}, e^{-k}) \quad \text{for every platform } k,$$

together with the adoption profile and hourly market conditions that these fees generate.

4 Estimation

Estimation proceeds in steps.

1) Rider demand is a nested logit: the inner nest is substitution among platforms, the outer

nest substitution between ride-hailing and other modes. The inner nest is estimated in two steps, first recovering each platform’s mean utility in every route market from the platforms’ relative shares, then decomposing the mean utilities into price and pickup time with the platform’s price and pickup time on other routes in the same hour as instruments. The outer nest is calibrated to experimental estimates (Cohen et al., 2016; Castillo, 2025).

2) Drivers’ hourly opportunity costs and the platforms’ commission rates are estimated from the working schedules of drivers who adopt a single platform, with the dominant platform’s commission anchored at 20%.

3) Drivers’ multi-homing costs are estimated by maximum likelihood from the search times between consecutive orders.

4) The matching function is estimated by regressing pickup time on the stock of searching drivers, controlling for travel speed and platform effects.

5) Drivers’ adoption costs are inverted from the model-implied values of each combination and its observed adoption share; the entry margin is calibrated to experimental estimates (Castillo, 2025).

6) Platforms’ objective weights are recovered from each platform’s first-order conditions at the observed fees.

The following subsections describe each step.

4.1 Price Measurement

Platforms price trip by trip, but the model defines substitution at the level of the route market m , so demand estimation requires a price at that level. The index I construct serves a second purpose: it separates the part of a fare that platforms compete on from the part that reflects the trip itself. A longer trip costs more on every platform, and platforms carry different mixes of trips, so raw fares compared within a market would partly compare trips rather than prices.

I decompose the fare of order i into a base price and a surge level,

$$p_i = \underbrace{\left(b_0 + \overbrace{b_\ell (\ell_i - 3)^+}^{\text{mileage fee}} + \overbrace{b_r (r_i - 9)^+}^{\text{duration fee}} \right)}_{\text{base price } b(\ell_i, r_i)} \times \underbrace{\exp(\Omega_{m(i)}^{k(i)} + \Omega_i)}_{\text{surge level}}, \quad (5)$$

where ℓ_i and r_i are the trip’s distance and duration. The base price follows the industry’s

common fare structure: a starting fare, plus mileage and duration fees beyond the first 3 km and 9 minutes. The kink points are not fitted: 3 km and 9 minutes are the thresholds written into the platforms' published fare rules. The schedule is the same for all platforms so any difference in schedule is captured by the surge level. The surge level has two parts: Ω_m^k is platform k 's log surge level over the common schedule in market m , and Ω_i is a mean-zero order-level deviation. Because the schedule is common, every difference in fare schedules across platforms, together with time-of-day multipliers and surge, loads onto Ω_m^k , which is therefore the price difference that riders in market m face. I estimate Equation (5) by nonlinear least squares.

Platform k 's per-minute price in market m is the fare it charges for the market's average trip, per minute of trip time,

$$p_m^k = \frac{b(\bar{\ell}_m, \bar{r}_m)}{\bar{r}_m} \exp(\Omega_m^k), \quad (6)$$

where $\bar{\ell}_m$ and $\bar{r}_m$ are the average distance and duration of trips in market m across all platforms. Evaluating the base price at the market average removes the variation that differences in trip length would otherwise generate, so within a market the price varies across platforms only through Ω_m^k .

4.2 Rider Demand

A rider in market m chooses among the platforms available on the route or the outside option. Utility is given by Equation (1), and riders draw private taste shocks with a nested logit structure: the inner nest is the choice among platforms, the outer nest the choice between ride-hailing and the outside option. I write utility on the within-nest scale, so the nesting parameter $\rho_m \in (0, 1]$ enters only the outer nest, and platform k 's demand in market m is

$$q_m^k = M_m \cdot \mathbb{E}_i \left[\underbrace{\frac{\exp(u_{imk})}{\sum_{l \in J_m} \exp(u_{iml})}}_{\text{inner nest}} \cdot \underbrace{\frac{\left(\sum_{l \in J_m} \exp(u_{iml}) \right)^{\rho_m}}{1 + \left(\sum_{l \in J_m} \exp(u_{iml}) \right)^{\rho_m}}}_{\text{outer nest}} \right], \quad (7)$$

where M_m is the number of potential riders in market m , J_m is the set of platforms available in the market, and the expectation averages over the trip types Z_i of the market's riders.

I estimate the inner nest and calibrate the outer nest. The data record the trips that took place but not the occasions on which a rider saw the prices and chose the outside option, so they identify riders' relative preferences among platforms but carry no information

on substitution to the outside option. The inner nest is estimated in two steps: the first recovers each platform’s mean utility in each route market from the order-level data, and the second decomposes the mean utilities into the contributions of price and pickup time. The nesting parameter is then calibrated so that the aggregate demand elasticity matches experimental estimates. I set ρ_m market by market so that in each market a 1% increase in all platforms’ prices leads to a 0.6% decrease in total order demand, which is consistent with the experimental estimates of [Cohen et al. \(2016\)](#) and [Castillo \(2025\)](#).

Step 1: Mean Utility Inversion

Rewrite the utility as:

$$u_{imk} = \delta_{mk} + Z_i' \gamma^k$$

where δ_{mk} is platform k ’s mean utility in market m and Z_i collects indicators for the types of the trip’s origin and destination. As discussed in Section 2, the place names riders enter in the app reveal what a place is, such as a restaurant or a university, beyond its coordinates. These types shift preferences across platforms: a trip to a tourist attraction is likely requested by a tourist, who may prefer a nationwide platform to a local one, and a trip to a university by a student, who may prefer a platform with student discounts. I match the place names to a point-of-interest database from a leading Chinese map service and classify the points into 16 categories, such as residential, work, shopping, and education.

Conditional on riding, the outer nest cancels and the choice among platforms is a multinomial logit over J_m . Not every platform serves every route: I treat a platform with no orders on a route in an hour as unavailable there. Because the data condition on riding, only differences in mean utility across the platforms in a market are identified. I therefore decompose the mean utility into a market average and a deviation,

$$\tilde{\delta}_{mk} = \delta_{mk} - \bar{\delta}_m, \quad \bar{\delta}_m = \frac{1}{|J_m|} \sum_{l \in J_m} \delta_{ml},$$

and note that the choice probability depends on the mean utilities only through the deviations,

$$\Pr(k \mid i, m) = \frac{\exp(\delta_{mk} + Z_i' \gamma^k)}{\sum_{l \in J_m} \exp(\delta_{ml} + Z_i' \gamma^l)} = \frac{\exp(\tilde{\delta}_{mk} + Z_i' \gamma^k)}{\sum_{l \in J_m} \exp(\tilde{\delta}_{ml} + Z_i' \gamma^l)},$$

because $\bar{\delta}_m$ cancels between numerator and denominator. The same argument applies to γ^k : a taste vector common to all platforms cancels, so γ^k is identified relative to a reference

platform, and I normalize it to zero for platform A. I estimate the deviations and the taste parameters by maximum likelihood on the order-level data,

$$(\hat{\delta}, \hat{\gamma}) = \arg \max_{\tilde{\delta}, \gamma} \sum_i \ln \Pr(k_i | i, m(i)) \quad \text{subject to} \quad \sum_{k \in J_m} \tilde{\delta}_{mk} = 0 \quad \text{for every } m.$$

The deviations $\tilde{\delta}_{mk}$ absorb every market-level determinant of platform shares, so γ^k is identified from how the platforms' shares differ across trip types within a market. The sample pools the eight weeks of data into markets defined by a township pair and an hour of the week within the six central districts, giving 9.2 million orders in 391,122 markets with at least two platforms present.

Step 2: Parameter Estimation. The second step decomposes the deviations into the contributions of price, pickup time, and brand preference. Substituting $\delta_{mk} = \alpha_k - \beta^p p_m^k - \beta^w w_m^k + \xi_m^k$ from Equation (1) into the definition of $\tilde{\delta}_{mk}$ gives the estimating equation

$$\tilde{\delta}_{mk} = \alpha_k - \beta^p p_m^k - \beta^w w_m^k - \bar{\delta}_m + \xi_m^k, \quad (8)$$

In this specification, the persistent brand preference α_k and the market-average mean utility $\bar{\delta}_m$ are platform and market fixed effects. The platform-market-level preference ξ_m^k is the source of the endogeneity concern. Platforms observe their appeal on a route and set prices and allocate drivers in response to it, so both price and pickup time are correlated with ξ_m^k .

I address this with Hausman instruments for both variables (Hausman, 1996; Nevo, 2001): for each, platform k 's value of the same variable in the same hour on all *other* routes. The logic is that a platform's prices and pickup times across routes share a platform-hour component that no single route's preference drives. This shared component may be a platform-hour supply shock, such as the platform's driver availability, or the platform's own response to a common hour-specific demand shock, which differs across platforms through their pricing and matching algorithms. Averaging over the other routes isolates this component and discards the platform's response to conditions on route m itself, and the variation in the component across platforms within an hour is what identifies the coefficients.

Formally, the instruments are valid if, conditional on the platform and market fixed effects, the preference of each platform in market m , ξ_m^k , is mean independent of the prices p_m^k and pickup times w_m^k that the platforms in the market post on other routes in the same hour. The assumption permits demand shocks that are common to all platforms. Route- or hour-level demand shocks shared by all platforms are absorbed by the market fixed

effects. A citywide shock in the hour, such as rain, is in fact a source of identifying variation: platforms’ algorithms respond to the same shock with different surge and dispatch rules, which moves their relative prices and pickup times on every route without moving their relative appeal. What the assumption rules out is a platform-hour-specific shock, such as a platform’s hour-specific promotion. In ride-hailing this is not a severe concern, because promotions mostly take the form of coupons, which the data record in the paid fare. Even allowing for measurement error, so that the coupon is not fully captured by the fare data, promotions are rarely set at the hour level. They usually last for weeks, and such a shift is absorbed by the fixed effects. A second concern is equilibrium spillover across routes. A platform’s appeal on one route draws its drivers toward that route, which raises its prices and lengthens its pickup times on nearby routes in the same hour, so the instrument would inherit part of route m ’s shock. Because the instrument averages over all other routes and any one route is a negligible share of the platform’s hourly demand, this feedback is of the order of one over the number of routes, and I ignore it.

Table 4.1: Rider demand estimates

	OLS	IV
β^p	0.440 (0.062)	1.229 (0.148)
β^w	0.003 (0.002)	0.129 (0.011)
Value of time (¥ per hour)	6.8	99.1
Route markets	137,696	
Platform $\times$ route market	418,265	

Notes: Estimates of β^p and β^w in Equation (8). Both columns include platform and market fixed effects. The IV column is estimated by GMM with the platform’s price and pickup time on all other routes in the same hour as instruments. The value of time is $60\bar{r}\beta^w/\beta^p$, with $\bar{r} = 15.7$ minutes the average trip duration. Standard errors in parentheses, clustered by hour of the week (168 clusters).

Table 4.1 reports the OLS and GMM estimates of Equation (8). The levels of the coefficients are not informative on their own, since they carry the scale of the calibrated outer nest. Their ratio, the value of time, is what can be compared with the literature. The GMM estimate is ¥99 (US\$14) per hour. The natural benchmark is the nationwide field experiment of Goldszmidt et al. (2026), which puts the value of time of ride-hailing riders in the United States at \$19 per hour, 75% of the mean wage. The average annual wage in the city is ¥159,659 (US\$23,300), or ¥76.5 (US\$11.2) per hour over the statutory 2,088 paid hours, so my estimate is 1.3 times the mean wage, a higher ratio than in the United States. A plausible reason is selection: public transportation in the city is a close substitute, so the riders who

choose ride-hailing are those who value their time most. To give a sense of the substitution across platforms, holding the extensive margin fixed, a 1% increase in price on average leads to a 1.05% decrease in a platform’s own share.

The comparison with OLS shows the direction of the endogeneity. OLS puts the price coefficient at a third of the GMM estimate and the pickup-time coefficient near zero, as expected when a platform’s popularity on a route makes up for a higher price and a slower pickup: where riders like a platform more, it charges more and its pickups are more congested, so both variables are positively correlated with ξ_m^k .

4.3 Driver Labor Supply

This subsection estimates the opportunity costs c_{dt} , the choice scale σ_1 , and the commission rates θ^k from the behavior of single-homing drivers. Two features of the estimation deserve comment. First, I observe the fare of each order but not the driver’s share of it. I assume that each platform takes a constant commission θ^k and infer it from how drivers’ labor supply varies across platforms and hours. Second, I use single-homing drivers because their working mode is known. For a driver who adopted two platforms, I observe which platform each order came from, but not whether the driver was searching on one platform or on both: a multi-homing driver may receive every order in an hour from a single platform. Restricting the sample to drivers who adopted one platform removes this ambiguity. The opportunity costs and commission rates are common to all drivers, and they are carried into the estimation of the multi-homing costs in the next subsection.

Under the payoff model of Equation (4), a driver who adopted a single platform makes a binary hourly decision with logit shocks, so the log odds of working are linear in earnings,

$$\log \frac{P_{dkt}}{1 - P_{dkt}} = \frac{(1 - \theta^k) \bar{e}_t^k r_t^k n_t^k - c_{dh(t)}}{\sigma_1},$$

where P_{dkt} is the probability that a driver of type d works on platform k in hour of the week t , $\bar{e}_t^k$ is the gross fare per minute, θ^k is platform k ’s commission rate, and $h(t)$ maps the hour of the week to its hour of the day on a weekday or a weekend day. Earnings vary across the 168 hours of the week, while the opportunity cost c_{dh} takes one value for each type in each of the 48 day-hour cells.

Identification rests on the assumption that a driver’s opportunity cost varies with type and time but not with the platform. Under it, two drivers of the same type in the same hour who face the same gross hourly earnings on different platforms differ in their probability of working

only through the commission. Their level is not identified, since scaling σ_1 and all take-home rates by the same factor leaves every choice unchanged. I therefore anchor platform A’s commission at 20%, the rate that DiDi, the dominant platform in China, disclosed in its IPO prospectus. The opportunity costs c_{dh} are the type-by-time fixed effects, identified from how a type’s probability of working on a given platform varies across hours beyond what earnings explain.

Table 4.2 reports the estimates. The estimates are $\sigma_1 = 29.5$ ¥ (US\$4.3) per hour and commission rates between 6% and 20% for the other platforms. The implied aggregate elasticity of hours worked with respect to pay is 1.10, close to the experimental estimate of 1.2 (Angrist et al., 2021). Figure 4.1 plots the opportunity costs. For the two daytime types, the cost of working is highest late at night and lowest during the day, mirroring their schedules; the night-owl type’s profile is much flatter, which is what rationalizes its schedule.

Table 4.2: Single-homing labor-supply estimates

	Estimate	SE
σ_1	29.5	(2.49)
θ^A	0.200	—
θ^B	0.070	(0.021)
θ^C	0.126	(0.016)
θ^D	0.060	(0.019)
θ^E	0.203	(0.014)
θ^F	0.170	(0.014)
Cells (type $\times$ time $\times$ platform)	3,024	
R^2	0.97	
Aggregate labor-supply elasticity	1.10	

Notes: Estimates from the single-homing participation model; standard errors in parentheses. σ_1 is the scale of the hourly taste shocks, in ¥ per hour, and θ^k is platform k ’s commission rate. θ^A is calibrated to the dominant platform’s publicly known rate; the remaining rates are estimated relative to it. Cells are working probabilities by driver type, hour of day, weekend, and platform.

4.4 Multi-Homing Costs

The multi-homing primitives, the switching costs γ_k , the multi-homing cost γ_M , and the scale σ_2 , are estimated from the driver-hours of drivers who adopted a pair of platforms. For each hour I observe whether the driver was active, the platform of each order served, and the search times between consecutive orders. If the driver received orders from different platforms in an hour, then I directly observe the simultaneous search. However, if the driver received orders from only one platform, it may be single-platform work or simultaneous search that




Figure 4.1: Estimated hourly opportunity costs by driver type

Notes: Estimated opportunity costs c_{dh} of the three full-time driver types, by hour of day for weekdays (panel a) and weekends (panel b), in ¥ per hour. The costs are recovered from the type-by-time fixed effects of the single-homing participation model.

happened to draw every order from one side. In that case I do not directly observe the working mode.

Search times are informative about the working mode. Under the Poisson dispatch model, the search time of a driver searching on platform k alone is exponentially distributed with mean s_t^k , and that of a driver searching on both platforms with the shorter mean s_t^M , whichever platform the orders come from. The mode is never observed, but each mode assigns a likelihood to the search times realized in an hour: an hour whose orders all come from platform k is more likely to be simultaneous search the shorter its search times are relative to s_t^k . The estimation combines these likelihoods with the choice probabilities of the modes, as follows.

The likelihood of a driver-hour depends on its order composition. An hour with no orders contributes the probability of staying inactive, $P_{jt}(0)$, from the four-option logit of Section 3.2.3 at the hour's payoffs, with the opportunity cost c_{dt} from the single-homing estimation. An active hour contributes the probability of working times the likelihood of its orders, each described by its platform k_o and the search time τ_o that preceded it. The density of an order under single-platform search on k and under simultaneous search is

$$g_k(\tau_o, k_o) = \frac{1}{s_t^k} \exp\left(-\frac{\tau_o}{s_t^k}\right), \quad g_M(\tau_o, k_o) = \frac{1}{s_t^{k_o}} \exp\left(-\frac{\tau_o}{s_t^M}\right),$$

where the first is the exponential density with mean s_t^k and every order comes from k , and the second is the joint density of the search time and the platform under simultaneous search: the search time is exponential with mean s_t^M , and the order comes from platform k with probability s_t^M/s_t^k . An hour with orders from both platforms can only be simultaneous search, so its likelihood is the conditional probability of that mode times the product of g_M over its orders. An hour whose orders all come from platform k may be single-platform search on k or simultaneous search, so its likelihood mixes the two, each weighted by its choice probability conditional on working. Collecting the cases, the log-likelihood is

$$\ell(\gamma_k, \gamma_M, \sigma_2) = \sum_{j,t} (1-f_{jt}) \log P_{jt}(0) + f_{jt} \left[\log (1 - P_{jt}(0)) + \log \sum_{a \in A_{jt}} \Pr_{jt}(a \mid \text{work}) \prod_o g_a(\tau_o, k_o) \right],$$

where f_{jt} indicates an active hour, A_{jt} is the set of modes consistent with the platforms of the hour's orders, and $\Pr_{jt}(a \mid \text{work})$ is the choice probability of mode a conditional on working. The parameters enter through these choice probabilities: γ_M lowers the probability of simultaneous search, γ_k lowers the probability of working on the secondary platform for pairs whose primary is k , and σ_2 scales all of them.

Table 4.3 reports the estimates. The multi-homing cost is $\gamma_M = 77.3$ ¥ (US\$11.3) per hour, and the switching costs range from ¥61 to ¥93 (US\$8.9 to US\$13.6). The costs are on the order of a full hour's earnings, which rationalizes the pattern of Section 2.3: adopting a second platform is common, but using both at once is rare. The switching costs also vary substantially with the pair's primary platform. Drivers whose primary platform is A face the highest cost of working on their secondary platform, ¥93 (US\$14), while drivers primarily on C, D, or E face about ¥60 (US\$8.8). These costs measure drivers' attachment to their primary platform in hourly labor supply, and they are a source of monopsony power: the higher the cost, the more firmly a platform's drivers are locked in to it. By this reading, the dominant platform holds the tightest grip on its own drivers.

4.5 Matching Function

I estimate the pickup-time function Equation (3) on a citywide panel of platform-by-hour-of-week cells in log-linear form,

$$\log w_t^k = a^k + a^N \mathbb{1}\{t \in \text{night}\} - b \log I_t^k + c \log v_t^k + \varepsilon_t^k, \quad (9)$$

where the platform intercepts a^k capture differences in matching efficiency across platforms and v_t^k is the average travel speed, which controls for the traffic condition. The night indicator

Table 4.3: Multi-homing cost estimates

	Estimate	SE
γ_M	77.3	(0.17)
γ_A	93.3	(0.25)
γ_B	71.9	(0.28)
γ_C	61.9	(0.22)
γ_D	62.0	(0.46)
γ_E	60.8	(0.49)
γ_F	71.8	(0.44)
σ_2	30.4	(0.04)

Notes: Maximum-likelihood estimates on pair adopters' driver-hours, in ¥ per hour; standard errors in parentheses. γ_M is the multi-homing cost, γ_A – γ_F the switching costs by the pair's primary platform, and σ_2 the scale of the taste shocks. The opportunity costs c_{dt} are taken from the single-homing estimation. Active driver-hours are those with at least one order.

(hours 1–6) deserves a word. The model abstracts from space, so density is captured by the aggregate searching stock I_t^k , as if searching drivers were spread evenly over the city. This approximation is reasonable only while the active area of the market is stable, and between day and night it is not. At night, riders and drivers concentrate in a small part of the city, so the citywide searching stock falls roughly tenfold while the driver density a rider actually faces falls much less. The night indicator absorbs this systematic contraction of the active area, so b is identified from density variation within the day and night regimes rather than from the contrast between them.

I estimate Equation (9) using weighted least squares with order volume as the weight for each platform-by-hour-of-week cell. The estimates are $b = 0.357$ (SE = 0.024) and $c = -0.50$, with $R^2 = 0.91$. A ten percent increase in searching drivers shortens pickups by about 3.5 percent.

4.6 Platform Adoption

With the hourly model estimated, the weekly value $\bar{V}_{gd}$ of each combination to each driver type is computed from the model at the observed market conditions, and two parameters remain in the adoption stage, the adoption costs κ_g and the nest parameter σ_3 . Given σ_3 , the costs are just identified: I invert them by a fixed-point contraction so that the model's combination shares match the observed shares exactly, normalizing the cost of adopting the dominant platform alone to zero. The nest parameter is then estimated by minimum distance from the variation across driver types. With the costs common to all types, a type adopts a

combination more often the higher its value, and how strongly adoption responds to value pins down the scale of the taste shocks. The estimate is $\sigma_3 = 532$ ¥ (US\$78) per week.

The entry margin cannot be estimated from within the market, because potential drivers who do not enter are not observed. I calibrate it: the pool of potential drivers is set to twice the number of active drivers, and the scale of the top-level shock is set so that the long-run elasticity of the number of drivers with respect to pay equals 0.6, in line with experimental estimates (Castillo, 2025).

Table 4.4 reports the estimated adoption costs of all combinations. The costs line up with the adoption patterns of Section 2.3: the three large platforms are the cheapest to adopt, the popular pairs built on primary A come next, and the fringe platforms cost as much as a pair. The cost also depends on the order within a pair: adopting C as a secondary to A costs less than half of adopting A as a secondary to C.

Table 4.4: Estimated adoption costs of platform combinations

Secondary	Primary platform					
	A	B	C	D	E	F
SH	0	607	906	1,726	1,851	2,146
A	×	2,885	2,279	3,466	2,641	2,759
B	2,423	×	2,652	3,709	4,149	3,841
C	1,151	2,405	×	3,151	2,672	2,779
D	2,517	2,698	2,362	×	4,018	3,678
E	1,218	2,532	2,352	3,401	×	3,569
F	1,990	3,835	3,059	4,409	4,176	×

Notes: Estimated adoption costs κ_g in ¥ per week. Columns are the combination's primary platform and rows its secondary platform; the "None" row is single-platform adoption. The cost of adopting platform A alone is normalized to zero, and × marks a platform paired with itself, which is not a combination.

4.7 Objective Weights

The weights $(\lambda_R^k, \lambda_D^k)$ are recovered from the condition that the observed fee levels are optimal for each platform. At the observed equilibrium, platform k 's objective is stationary with respect to its two levers, the level of its rider price and the level of its driver payment,

$$\frac{\partial \Pi^k}{\partial p^k} + \lambda_R^k \frac{\partial CS^k}{\partial p^k} + \lambda_D^k \frac{\partial DS^k}{\partial p^k} = 0, \quad \frac{\partial \Pi^k}{\partial e^k} + \lambda_R^k \frac{\partial CS^k}{\partial e^k} + \lambda_D^k \frac{\partial DS^k}{\partial e^k} = 0.$$

I compute the derivatives numerically, re-solving the adoption and hourly market equilibrium at perturbed fee levels, so each derivative carries the full equilibrium response of both sides of the market. The two conditions are linear in the two weights, so $(\lambda_R^k, \lambda_D^k)$ is just identified platform by platform. The price condition reveals how much rider surplus the platform must value for its price not to be higher, and the payment condition how much driver surplus it must value for its pay not to be lower.

Table 4.5: Recovered objective weights

	λ_R^k	λ_D^k
A	0.039	0.854
B	0.875	0.078
C	0.691	0.180
D	0.546	0.273
E	0.529	0.366
F	0.353	0.319

Notes: The weights on rider surplus (λ_R^k) and driver surplus (λ_D^k) that solve each platform’s two first-order conditions at the observed fee levels, with the derivatives computed through the full equilibrium re-solve. The weights are just identified, one pair per platform.

Table 4.5 reports the recovered weights. All twelve are positive: every platform sets its fees more favorably to users than static profit maximization would imply, charging riders less and paying drivers more than a pure profit maximizer would. Positive conduct weights of this kind are also found in ride-hailing by [Rosaia \(2025\)](#) and in other platform industries by [Sullivan \(2026\)](#). The two user-surplus weights sum to between roughly 0.7 and 0.95, so under a normalization in which the weights sum to one, platforms place slightly more than half of the total weight on their own profit, a level broadly in line with [Rosaia \(2025\)](#).

5 Counterfactuals

I design two counterfactual studies to explore separately the welfare implications of platform competition and of multi-homing. First, I simulate the market equilibrium under platform mergers that lead to different levels of market concentration. By comparing the equilibrium across these structures, I quantify the welfare impact of mergers, separate the roles of matching efficiency and market power, and quantify the distributional effect on each side of the market. Second, I assess the impact of changing the level of multi-homing, by banning it or by making it cheaper for drivers.

I do not consider the monopoly case. With competing platforms, fee-setting is disciplined

by riders' and drivers' substitution across platforms, which the data identify. However, the fee-setting of a monopolist is governed by the substitution to the outside options, riders' substitution to other modes of transportation and drivers' entry into the industry, and both of these margins are calibrated to experimental estimates rather than estimated. The monopoly outcome would therefore reflect the calibration more than the estimated model.

5.1 Market Structure

In this counterfactual study, I simulate mergers between platforms that generate different levels of concentration, to study the trade-off between market power and network efficiency as the market concentrates. A merger is implemented as a change in the platforms' objective functions, with every platform continuing to operate. The merged platforms coordinate on rider prices p^k and driver payments e^k to maximize their combined objective function, evaluated at the conduct weights $(\lambda_R^k, \lambda_D^k)$ of the largest member. The merger can therefore be viewed as an acquisition: the largest platform absorbs the smaller ones, keeps their brands in operation, and imposes its own objective on the combined entity.

Formally, let $K = \{k_1, \dots, k_n\}$ be a set of merged platforms, ordered so that k_1 is the member with the largest order volume at the status quo. The merged block chooses the fee levels of all its members to maximize

$$O^K = \sum_{k \in K} \Pi^k + \lambda_R^{k_1} \sum_{k \in K} CS^k + \lambda_D^{k_1} \sum_{k \in K} DS^k, \quad (10)$$

taking the fees of the platforms outside K as given.

For a given number of platforms, such as a duopoly or a triopoly, there are multiple ways to merge the six platforms into that structure. To keep the analysis tractable and to make each case representative, in the main analysis I focus on the merger that yields the most evenly sized platforms at each level of concentration.² Holding the number of platforms fixed, the market can be effectively more concentrated when one platform has a much larger market share, so the selected consolidation path can also be read as the least concentrated case for each number of platforms.

When platforms merge, there is an efficiency gain, because drivers and riders can be matched across the merged platforms. As discussed in Section 3.2.2, the matching process captures

²For the duopoly, I merge platform A with D, E, and F and platform B with C. At the observed equilibrium the two blocks' market shares are 50.9% and 49.1%. For the triopoly, I merge platform C with D, E, and F, which gives market shares of 36.9% for platform A, 31.1% for platform B, and 32.0% for the merged block. For the four-platform structure, I merge D, E, and F, yielding shares of 36.9%, 31.1%, 18.0%, and 14.0%.

the economies of density that underlie network efficiency. After a merger, drivers' search time, which is governed by the balance of order demand and driver supply, need not change, but pickup time unambiguously shrinks. Intuitively, in a market fragmented by competing platforms, riders may miss potential matches with nearby drivers who are active on a different platform, and vice versa. A merger eliminates this inefficiency by making all idle drivers simultaneously available to riders on all merged platforms, improving matching efficiency. This channel is what makes concentration potentially welfare-improving.

Yet the welfare impact of a merger is ambiguous, because a merger also gives the merged platforms greater market power, which enables them to raise fees inefficiently. To separate the effects of the two forces in equilibrium, I compute two additional equilibrium outcomes, each holding one force at its status quo level. In the *fee-only* equilibrium, the merged platforms coordinate fees but keep separate fleets and dispatch systems, so concentration brings greater market power without any matching-efficiency gain. In the *matching-only* equilibrium, the merged platforms match drivers and riders across the merged platforms while fees are held at their observed levels, so this case reveals only the gain from higher matching efficiency.

5.1.1 Total Welfare

Figure 5.1 shows total welfare and annual trips at different levels of market concentration. The most notable result on total welfare is that reducing the number of platforms in the market can be welfare-improving. Intuitively, as the market becomes more concentrated, the matching-efficiency gain outweighs the loss from market power, and total welfare rises with concentration at every structure. Relative to the status quo, concentration creates up to 1.5 million additional trips annually (at the triopoly) and up to ¥241 million (US\$35 million) of additional total welfare (at the duopoly). The fact that reduced platform competition can raise welfare has been documented by recent empirical studies of platform competition in the ride-hailing market (Rosaia, 2025) and the payment-card market (Wang, 2025).

However, it is not guaranteed that welfare keeps rising as the market consolidates. Figure 5.2 shows total welfare along more consolidation paths. To keep the analysis tractable, I first merge the fringe platforms D, E, and F into one mid-sized platform, and then consider every way of consolidating the main platforms A, B, and C and this mid-sized platform into a triopoly or a duopoly. The selected baseline structures are the least concentrated ones, in terms of HHI, for each number of platforms.

The level of concentration matters for the efficiency impact of market consolidation. Every merger in Figure 5.2 raises total welfare above the status quo, but the gain does not grow




Figure 5.1: Total welfare across market structures

Notes: Total welfare and annual trips at the counterfactual equilibrium of each market structure. Fee-only and matching-only are partial-equilibrium outcomes that hold, respectively, the matching method or the fees at their status quo levels, to isolate the effects of market power and matching efficiency. The number on each connector is the gap between the counterfactual and the status quo.

monotonically with concentration. Welfare peaks at the least concentrated duopolies, where network consolidation raises matching efficiency and each platform’s market power is kept in check by a rival of comparable size. Once one platform becomes dominant, the additional harm of market power outweighs the additional gain from matching, and total welfare falls.




Figure 5.2: Total welfare across other merger cases

Notes: Each point is a counterfactual equilibrium with different merges. The horizontal axis is the Herfindahl-Hirschman index of the platforms’ order shares at the counterfactual equilibrium. The dashed line marks the status quo. Crosses mark the three baseline market structures. Total welfare is the sum of platform profit, rider surplus, and driver surplus.

5.1.2 Market Power and Matching Efficiency

Comparing the status quo with the *fee-only* counterfactuals along the consolidation path isolates the harm of market power. As Figure 5.1 shows, market power alone costs 0.3 to 8.6 million trips and ¥13 to ¥124 million (US\$1.9 to US\$18.1 million) of total surplus annually as the market consolidates from the four-platform structure to the duopoly.

This cost is borne by drivers and riders through fees rather than wait times. As Figure 5.3 shows, in the *fee-only* equilibriums riders pay more and drivers earn less as the market becomes more concentrated, while the market-level pickup time remains essentially unchanged across these structures. To give a sense of the magnitude of the fee changes, moving from the status quo to the duopoly raises the rider price per order by ¥3.0 (US\$0.44, 13%) and lowers driver

earnings per order by ¥0.7 (US\$0.10, 3.6%), so that the platforms' margin per order rises by ¥3.7 (US\$0.54) — from ¥3.1 to ¥6.8 (US\$0.45 to US\$0.99), more than double its status-quo level.




Figure 5.3: Fee structure and matching times across market structures

Notes: Equilibrium fees and wait times at each market structure. Fees and wait times are order-weighted means across platforms and hours. Driver pay is the pay received by the serving driver, which under centralized dispatch follows the driver's home platform.

Comparing the status quo with the *matching-only* counterfactual isolates the gain from matching efficiency, and the gain is substantial. As Figure 5.1 shows, pooling the merged platforms' dispatch at the observed fees raises total welfare by ¥117 million (US\$17 million) per year (1.9%) in the four-platform structure, ¥219 million (US\$32 million, 3.5%) in the triopoly, and ¥364 million (US\$53 million, 5.8%) in the duopoly, and raises annual trips by 1.6, 3.7, and 6.5 million. At every structure the gain is a multiple of the market-power cost, nine times in the four-platform structure and still nearly three times in the duopoly. The gain grows with concentration because a larger merged block pools more riders and drivers into one network.

The source of the gain is shorter pickups. Panel (b) of Figure 5.3 shows that pooling dispatch cuts the market-level pickup time from 5.2 minutes at the status quo to 4.8 minutes in the four-platform structure, 4.3 minutes in the triopoly, and 3.7 minutes in the duopoly, a reduction of 8% to 30%. This is the economy of density of Section 3.2.2 at work. Pooling the merged platforms' idle drivers puts every rider closer to the nearest available driver, so




Figure 5.4: Welfare Decomposition across market structures

Notes: Each panel shows the change in one side’s annual surplus relative to the observed six-platform status quo, by market structure, for the three counterfactual equilibrium outcomes. The fee-only equilibrium re-optimizes fees under the new ownership with separate dispatch and isolates the market-power effect of concentration. The matching-only point pools dispatch within each merged block while holding every fee at its observed level and isolates the matching-efficiency gain. The full equilibrium combines the two. Driver surplus is that of full-time drivers.

pickups shorten with essentially the same driver supply.

5.1.3 Distribution Across Sides

Finally, to understand who benefits from market concentration, I decompose total welfare into platform profit, rider surplus, and driver surplus, and compare the three across levels of market concentration in Figure 5.4. For each of the three, I also show the *fee-only* and *matching-only* equilibrium, which separate the impact of market power from that of matching efficiency.

There are two remarkable takeaways from the full equilibrium in Figure 5.4. First, the three sides of the market experience substantially different impacts from market concentration. Second, the level of concentration matters for whether riders gain or lose. Platforms earn more profit and drivers lose surplus as the market concentrates. This pattern, in which merging competing platforms benefits the platforms while harming their users, has been documented in recent studies (Rosaia, 2025; Wang, 2025; Sullivan, 2026). What is novel and has not been documented is that the level of concentration determines whether riders gain or lose from a merger. Relative to the status quo, riders gain from a mild merger but lose from a severe one.

The decomposition into matching efficiency and market power shows why the sides are affected so differently. Comparing the status quo with the *matching-only* equilibriums shows that the matching-efficiency gain benefits all sides, but predominantly riders. The reason is that, although the time saved on a pickup is the same for everyone, riders value that time much more than platforms and drivers do. Riders can use the time for things they value more, so every minute saved is a direct utility gain. Drivers and platforms, in contrast, gain from shorter pickups only if they receive more orders. As Figure 5.1 shows, merging to a duopoly raises trips by only 6% in the matching-only equilibrium, and platforms receive a modest gain in profit because their margin is currently low. Drivers’ gain is also limited, because the pickup time saved is not fully converted into more orders. Riders’ demand responds only modestly to pickup time, so as drivers save time on pickups, they spend more time searching for the next order in equilibrium, as Figure 5.3 shows. Comparing the status quo with the *fee-only* equilibriums shows the impact of market power, which benefits only the platforms and harms drivers and riders, as discussed in Section 5.1.2.

For drivers, the gain from matching efficiency does not survive the platforms’ fee response. Although drivers benefit from higher matching efficiency when fees are held fixed, their surplus in the full equilibrium is lower than in the *fee-only* equilibrium, where there is no matching-efficiency gain. The reason is a fall in the platforms’ demand for labor. The matching-efficiency gain effectively makes drivers more productive, because they can complete an order faster. The higher productivity is only partly converted into additional rider demand, through the shorter pickups that riders value, so at the same driver pay drivers sit idle for longer. Platforms therefore need less labor, because they can deliver the same quality of service, in terms of pickup time, with fewer drivers. As Figure 5.3 shows, platforms cut driver pay in the full equilibrium by even more than in the *fee-only* equilibrium. Although drivers spend less time searching for orders, their earnings per hour are still lower. This pattern, in which an increase in workers’ productivity lowers the demand for labor, has been discussed as the displacement effect in the recent labor literature (Bessen, 2019; Acemoglu and Restrepo, 2019).

5.2 Multi-homing

Having examined the trade-off between market power and network efficiency, I turn to the second policy lever, the ease of multi-homing itself. I simulate counterfactuals with separate changes in the two driver working frictions, the switching cost γ_k and the multi-homing cost γ_M , under different market structures. These counterfactuals thus demonstrate how each friction impacts welfare and how the impact of multi-homing is related to changes in market

structure.

In the main analysis, I focus the counterfactuals on marginal changes in these costs, so that the costs remain positive after the reduction. I run these experiments at every market structure along the consolidation path of Section 5.1, which reveals how the value of multi-homing varies with concentration.

5.2.1 Total welfare

Promoting multi-homing is welfare-improving. Banning it lowers both total welfare and the number of trips, and making it cheaper raises both, at every market structure along the consolidation path (Figure 5.5). The magnitudes are modest but not negligible: at the current market structure, lowering both frictions by ¥20 (US\$2.9) per hour raises annual welfare by about as much as merging the three fringe platforms into a single firm with centralized dispatch.

Figure 5.6 separates the two frictions at the observed market structure, reducing the switching cost alone, the multi-homing cost alone, and both together, over a wider range of cuts. Three regularities stand out. First, the two frictions are close substitutes in welfare terms. A given reduction in either one raises total welfare by nearly the same amount, and drivers take the largest share on both axes through the participation channel described above. Second, the multi-homing cost delivers somewhat more to riders as the cut deepens, and it does so through increased matching efficiency. Yet a cut in the switching cost makes drivers more elastic to pay competition, so the shift is smaller and benefits riders less. Third, the effect of the joint reduction is approximately the sum of the two single-cost effects, with a modest complementarity that appears only at the largest cuts.

The source of the gain, however, is not the one usually invoked for multi-homing. Multi-homing is often expected to improve network efficiency. A driver who is online on several platforms is exposed to more orders, so search and pickup times should fall as multi-homing spreads. At the frictions I estimate this channel is small, and the welfare gain comes almost entirely from drivers' higher willingness to work.

Figure 5.7 isolates the matching channel by tracing the response of pickup time to the reallocation of drivers' working modes, holding fees and total labor supply fixed at their observed levels. When drivers reallocate single-platform working time to multi-homing, there is a matching-efficiency gain, as the market average pickup time falls with the same labor supply. This channel, however, is small. In money, at the deepest cut it generates additional rider welfare of only ¥4.7 million (US\$0.7 million) per year. The rest of the gain comes




Figure 5.5: Total welfare across multi-homing level

Notes: Total welfare and annual trips at the equilibrium of each market structure under the multi-homing ban, the observed costs, and joint reductions of the switching and multi-homing costs by ¥10 and ¥20 (US\$1.5 and US\$2.9) per hour. The number next to each marker is the change relative to that structure's equilibrium at the observed costs. Merged structures are under centralized dispatch.




Figure 5.6: Welfare incidence along the three multi-homing cost axes

Notes: Change in annual platform profit, rider surplus, and full-time driver surplus relative to the observed six-platform equilibrium, after platforms reset fees. Each panel changes the named cost and holds the other at its estimate. The point at infinity is a ban of the working mode.

from added working time. With the two costs lower, more drivers find it worthwhile to work additional hours, whether by running both apps at once or by moving to the better-paying platform, in hours they would otherwise have stayed off the road. At the observed equilibrium a full-time driver works 52.5 hours per week, of which 1.4 are spent online on two platforms at once. A ¥20 (US\$2.9) reduction in both frictions raises weekly hours to 53.6 and simultaneous multi-homing hours to 3.2. Multi-homing hours more than double, but the actual use of multi-homing remains limited by the high remaining costs and thus generates only a limited gain in matching efficiency.

5.2.2 Distribution across sides

Figure 5.8 decomposes the welfare change into platform profit, rider surplus, and full-time driver surplus, by side of the market and by market structure. The general pattern is that the easier it is to switch and to multi-home, the more each side gains. At the six-platform structure, moving from the ban to the observed costs and on to the ¥10 and ¥20 (US\$1.5 and US\$2.9) cuts raises annual platform profit from ¥10 million (US\$1.5 million) below the status quo to ¥5 million (US\$0.7 million) and ¥11 million (US\$1.6 million) above it, rider surplus from ¥31 million (US\$4.5 million) below to ¥17 million (US\$2.5 million) and ¥42 million (US\$6.1 million) above, and driver surplus from ¥46 million (US\$6.7 million) below to ¥23 million (US\$3.4 million) and ¥59 million (US\$8.6 million) above. Drivers gain the most and platforms the least.




Figure 5.7: Pickup time from the reallocation of working modes alone

Notes: Direct response of market average pickup time to drivers' reallocation of working mode due to change of frictions at the status quo market structure. Rider demand, fees, adoption, and each driver's total working time are held at their observed values, and the resulting idle stock is passed once through the matching technology, so the points are not equilibrium outcomes. Each curve changes one cost and holds the other at its estimate.

The crucial difference from the market-structure counterfactuals of Section 5.1 is that cheaper multi-homing benefits all three sides of the market, whereas concentration can harm drivers and riders. At the duopoly, where concentration has left riders and drivers worse off, the ¥10 (US\$1.5) cut still raises platform profit by ¥10 million (US\$1.5 million), rider surplus by ¥5 million (US\$0.7 million), and driver surplus by ¥9 million (US\$1.3 million). Drivers gain in aggregate despite the lower pay per order because the lower frictions raise the value of an hour of work by more than the pay cut takes away. Within the driver population the cost falls on the drivers who remain single-homers, who face the lower pay without the offsetting gain, and occasional drivers, whose surplus is omitted from the figure, lose modestly as the additional hours lengthen search times slightly.

The gains from cheaper multi-homing are larger the less concentrated the market. As a result, the level of multi-homing can change riders' preference over market structures but not that of drivers or platforms. For riders, driver multi-homing is a substitute for concentration as a source of gains. Concentration delivers shorter pickups, cheaper multi-homing delivers lower prices due to the shift of competition. Riders may therefore prefer a less concentrated market with a high level of multi-homing. In Figure 5.8, riders prefer the triopoly to the four-platform structure when multi-homing is banned or at its observed cost, the two structures




Figure 5.8: Welfare decomposition across multi-homing levels

Notes: Annual platform profit, rider surplus, and full-time driver surplus at the equilibrium of each market structure under the multi-homing ban, the observed costs, and joint reductions of the switching and multi-homing costs by ¥10 and ¥20 (US\$1.5 and US\$2.9) per hour. Merged structures are under centralized dispatch.

are about equal at the ¥10 (US\$1.5) cut, and riders prefer the four-platform structure at the ¥20 (US\$2.9) cut. As the cut deepens, the rider gap between the four-platform and the six-platform structures also narrows, from ¥66 million (US\$9.6 million) at the observed costs to ¥43 million (US\$6.3 million) at the ¥40 (US\$5.8) cut.

5.2.3 Fee Structure

A large theoretical literature, across a range of market structures, predicts that an increase in multi-homing on one side of a platform market raises the price charged to that side and lowers the price charged to the other side (Rochet and Tirole, 2003; Armstrong, 2006; Teh et al., 2023). The intuition, broadly, is that multi-homing makes users on that side less exclusive and therefore less valuable to compete for, so platforms shift their competitive effort to the other side. My results support this prediction.

Table 5.1 reports the equilibrium fee response to a joint reduction in both frictions along the consolidation path. As multi-homing becomes easier, platforms lower driver pay and rider price together, at every market structure. At the six-platform structure, moving from the multi-homing ban to the ¥20 (US\$2.9) joint cost reduction lowers driver pay per order from ¥19.97 to ¥19.34 (US\$2.92 to US\$2.83) and rider price per order from ¥22.99 to ¥22.50 (US\$3.36 to US\$3.29). The pass-through of the pay cut into price is less than one-for-one: riders receive about three quarters of the saving at the six-platform structure and about a

Table 5.1: Fee structure and matching times across multi-homing levels

	Rider price (% change)	Driver pay (% change)	Margin (% change)	Pickup time (% change)	Search time (% change)
<i>Panel A: Six platforms (status quo)</i>					
Multi-homing banned	+0.9%	+1.4%	-2.5%	+0.3%	-0.8%
Cheaper multi-homing (−¥10/hr)	-0.5%	-0.8%	+1.1%	+0.0%	+0.2%
Cheaper multi-homing (−¥20/hr)	-1.2%	-1.8%	+2.2%	-0.3%	+1.1%
<i>Panel B: Four platforms</i>					
Multi-homing banned	+0.9%	+2.1%	-5.8%	+0.3%	-0.4%
Cheaper multi-homing (−¥10/hr)	-0.5%	-1.0%	+2.5%	-0.2%	+0.2%
Cheaper multi-homing (−¥20/hr)	-1.2%	-2.4%	+5.8%	-0.4%	+0.6%
<i>Panel C: Triopoly</i>					
Multi-homing banned	+0.6%	+1.7%	-4.1%	+0.4%	-0.8%
Cheaper multi-homing (−¥10/hr)	-0.3%	-0.9%	+2.3%	+0.0%	+0.1%
Cheaper multi-homing (−¥20/hr)	-0.9%	-2.2%	+4.8%	-0.3%	+0.6%
<i>Panel D: Duopoly</i>					
Multi-homing banned	+0.4%	+1.8%	-2.8%	+0.1%	+0.0%
Cheaper multi-homing (−¥10/hr)	-0.2%	-0.9%	+1.4%	+0.0%	+0.0%
Cheaper multi-homing (−¥20/hr)	-0.5%	-2.4%	+3.7%	+0.0%	+0.0%

Notes: Percentage change in equilibrium outcomes relative to each structure’s equilibrium at the observed multi-homing costs. Rider price and driver pay are order-weighted means per order, and pay is the pay received by the serving driver, which under centralized dispatch follows the driver’s home platform. The margin is the difference between the two. Pickup and search times are order-weighted means.

third at the duopoly, where the merged blocks face less price competition. Platform margins therefore widen as multi-homing becomes easier, from ¥3.01 to ¥3.16 (US\$0.44 to US\$0.46) per order at the six-platform structure and from ¥7.54 to ¥8.04 (US\$1.10 to US\$1.18) at the duopoly, and the widening is larger the more concentrated the market.

Figure 5.9 separates the effect of the two frictions on the fee structure. The two frictions work in the same direction: lowering either one reduces both driver pay and rider price. Lowering the multi-homing cost, however, moves both fees by more. The logic is the one above. A reduction in either friction encourages drivers to work more, but a lower switching cost also makes drivers more responsive to pay differences across platforms, which intensifies the competition for drivers and offsets part of the fee change.

5.2.4 Distribution across platforms

Lowering the frictions also promotes competition by shifting usage toward the small platforms. Table 5.2 reports the change in each platform’s orders when the frictions are reduced. When




Figure 5.9: Equilibrium fees along the two multi-homing cost axes

Notes: Rider price and driver pay per order at the six-platform equilibrium after platforms reset fees, order-weighted across platforms and hours. Each series changes the named cost and holds the other at its estimate. The point at infinity is the ban of that cost.

the costs are lowered, the small platforms gain disproportionately relative to the dominant platforms A and B: at the ¥20 (US\$2.9) joint reduction, orders rise by 10.1% on E, 3.8% on D, and 3.6% on C and F, against 0.9% on A and a 3.0% fall on B. Lowering the switching cost is the more effective of the two, in that it lets the small platforms gain orders while leaving the dominant platform A unchanged or worse off, whereas lowering the multi-homing cost also benefits A.

Platform B is a net loser under either cut. It relies almost exclusively on single-homing drivers, so when the frictions fall and adopting a second platform becomes more valuable, B's drivers shift part of their work to other platforms.

Table 5.2: Order volume across multi-homing cost levels

Platform	Cost reduction (¥ per hour)				
	10	15	20	30	40
<i>Both costs reduced</i>					
A	+0.5	+0.7	+0.9	+1.5	+2.7
B	−2.1	−2.5	−3.0	−4.4	−6.3
C	+2.5	+3.0	+3.6	+5.5	+8.4
D	+1.1	+1.9	+3.8	+7.3	+13.0
E	+4.7	+7.2	+10.1	+17.9	+29.6
F	+0.8	+2.2	+3.6	+8.3	+14.3
<i>Switching cost γ_k reduced</i>					
A	+0.3	+0.3	+0.2	+0.2	+0.2
B	−1.8	−2.1	−2.3	−3.0	−4.2
C	+2.2	+2.4	+2.7	+3.5	+4.8
D	+0.9	+1.3	+2.0	+4.8	+7.9
E	+3.1	+4.5	+6.3	+10.8	+17.8
F	−0.4	+0.2	+0.9	+2.6	+5.7
<i>Multi-homing cost γ_M reduced</i>					
A	+0.6	+0.7	+0.9	+1.5	+2.5
B	−1.8	−1.9	−2.2	−2.6	−3.5
C	+2.2	+2.4	+2.7	+3.3	+4.5
D	+0.5	+0.7	+0.9	+1.5	+3.0
E	+2.3	+3.3	+4.5	+7.6	+12.1
F	−0.2	+0.6	+1.5	+3.4	+6.9

Notes: Percentage change in each platform's weekly orders relative to the observed six-platform equilibrium, after platforms reset fees. Each panel reduces the named cost by the stated amount and holds the other at its estimate. In the first panel both costs are reduced together.

To isolate the reallocation of driver-hours from driver entry and from the equilibrium responses of fees and market conditions, I compute the direct effect of the cost reduction holding fees




Figure 5.10: Reallocation of driver-hours across platforms

Notes: Change in each platform's weekly driver-hours relative to the observed level when both costs are reduced by the same amount and only drivers' choice of working mode responds. The number of drivers, adoption, fees, pickup and search times, and each driver's hourly participation are held at their observed values.

and market conditions fixed. I also hold each driver's hourly participation fixed, so the only change in Figure 5.10 comes from the reallocation of working hours across modes. After the joint cut, platforms A and B lose driver-hours while every fringe platform gains.

Figure 5.11 separates the contributions of the two frictions. The pattern is the same, the dominant platforms lose hours to the fringe, but the size of the effect on each platform depends on its exposure to multi-homing.

6 Conclusion

These results carry several lessons. First, market concentration can be desirable in the ride-hailing market because integration brings substantial matching-efficiency gains. This does not mean that the tension between market power and network effects disappears. As the market concentrates, greater market power raises rider prices and lowers driver pay, but the matching-efficiency gain outweighs this harm. The tension matters because it determines the market structure that each side prefers. Platforms benefit from both market power and network efficiency and thus prefer the most concentrated market. Drivers are harmed by both and thus prefer the least concentrated market. Riders benefit from network efficiency and




Figure 5.11: Decomposition of the working-hour effect: switching cost and multi-homing cost
Notes: The joint reduction of Figure 5.10 separated into its two components under the same accounting. Each panel reduces the named cost and holds the other at its estimate.

are harmed by market power, and thus prefer a market structure in between. For regulators, no market structure is optimal for everyone.

Promoting multi-homing also improves total welfare, and it benefits all sides of the market at once. The source of the gain, however, is not matching efficiency. It is that drivers work more hours, hours that would not be worthwhile at the higher multi-homing costs. Promoting multi-homing also promotes competition by giving the smaller platforms a competitive advantage.

References

- ACEMOGLU, D. AND P. RESTREPO (2019): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33, 3–30.
- ANDERSON, S. P. AND M. PEITZ (2020): “Media See-saws: Winners and Losers in Platform Markets,” *Journal of Economic Theory*, 186, 104990.
- ANGRIST, J. D., S. CALDWELL, AND J. V. HALL (2021): “Uber versus taxi: A driver’s eye view,” *American Economic Journal: Applied Economics*, 13, 272–308.
- ARMSTRONG, M. (2006): “Competition in two-sided markets,” *The RAND journal of economics*, 37, 668–691.
- ARMSTRONG, M. AND J. WRIGHT (2007): “Two-sided markets, competitive bottlenecks and exclusive contracts,” *Economic Theory*, 32, 353–380.
- ARNOLD, D. (2026): “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes,” *American Economic Review*, forthcoming.
- AZAR, J., I. MARINESCU, AND M. STEINBAUM (2022): “Labor Market Concentration,” *Journal of Human Resources*, 57, S167–S199.
- BAKOS, Y. AND H. HALABURDA (2020): “Platform competition with multihoming on both sides: Subsidize or not?” *Management Science*, 66, 5599–5607.
- BENMELECH, E., N. K. BERGMAN, AND H. KIM (2022): “Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?” *Journal of Human Resources*, 57, S200–S250.
- BESBES, O., F. CASTRO, AND I. LOBEL (2022): “Spatial capacity planning,” *Operations Research*, 70, 1271–1291.
- BESSEN, J. (2019): “Automation and Jobs: When Technology Boosts Employment,” *Economic Policy*, 34, 589–626.
- BRANCACCIO, G., M. KALOUPSTIDI, T. PAPAGEORGIOU, AND N. ROSAIA (2023): “Search frictions and efficiency in decentralized transport markets,” *The Quarterly Journal of Economics*, 138, 2451–2503.
- BUCHHOLZ, N. (2022): “Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry,” *The Review of Economic Studies*, 89, 556–591.

- CAILLAUD, B. AND B. JULLIEN (2003): “Chicken & egg: Competition among intermediation service providers,” *RAND journal of Economics*, 309–328.
- CASTILLO, J. C. (2025): “Who benefits from surge pricing?” *Econometrica*, 93, 1811–1854.
- CASTILLO, J. C., D. KNOEPFLE, AND E. G. WEYL (2025): “Matching and pricing in ride hailing: Wild goose chases and how to solve them,” *Management Science*, 71, 4377–4395.
- CHEN, Y., Y. LUO, AND Z. YUAN (2024): “Driving the Drivers: Algorithmic Assignment in Ride-Hailing,” *Available at SSRN 4299499*.
- COHEN, P., R. HAHN, J. HALL, S. LEVITT, AND R. METCALFE (2016): “Using Big Data to Estimate Consumer Surplus: The Case of Uber,” Working Paper 22627, National Bureau of Economic Research.
- CORREIA-DA SILVA, J., B. JULLIEN, Y. LEFOUILI, AND J. PINHO (2019): “Horizontal Mergers Between Multisided Platforms: Insights from Cournot Competition,” *Journal of Economics & Management Strategy*, 28, 109–124.
- EUROPEAN COMMISSION (2022): “Regulation (EU) 2022/1925 on Contestable and Fair Markets in the Digital Sector (Digital Markets Act),” Official Journal of the European Union, oJ L 265, 12 October 2022.
- FRÉCHETTE, G. R., A. LIZZERI, AND T. SALZ (2019): “Frictions in a Competitive, Regulated Market: Evidence from Taxis,” *American Economic Review*, 109, 2954–2992.
- GOLDSZMIDT, A., J. A. LIST, R. D. METCALFE, I. MUIR, V. K. SMITH, AND J. WANG (2026): “The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments,” *American Economic Journal: Economic Policy*, forthcoming.
- HAUSMAN, J. A. (1996): “Valuation of New Goods under Perfect and Imperfect Competition,” in *The Economics of New Goods*, ed. by T. F. Bresnahan and R. J. Gordon, University of Chicago Press, 207–248.
- KAISER, U. AND J. WRIGHT (2006): “Price structure in two-sided markets: Evidence from the magazine industry,” *International Journal of Industrial Organization*, 24, 1–28.
- MARINESCU, I. AND H. HOVENKAMP (2019): “Anticompetitive Mergers in Labor Markets,” *Indiana Law Journal*, 94, 1031–1063.
- NAIDU, S., E. A. POSNER, AND E. G. WEYL (2018): “Antitrust Remedies for Labor Market Power,” *Harvard Law Review*, 132, 536–601.

- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- OECD (2021): “Data Portability, Interoperability and Competition,” OECD Competition Policy Roundtable Background Paper 260, Organisation for Economic Co-operation and Development, dAF/COMP(2021)15.
- PRAGER, E. AND M. SCHMITT (2021): “Employer Consolidation and Wages: Evidence from Hospitals,” *American Economic Review*, 111, 397–427.
- ROCHET, J.-C. AND J. TIROLE (2003): “Platform competition in two-sided markets,” *Journal of the european economic association*, 1, 990–1029.
- ROSAIA, N. (2025): “Competing Platforms and Transport Equilibrium,” *forthcoming in Econometrica*.
- RYSMAN, M. (2004): “Competition between networks: A study of the market for yellow pages,” *The Review of Economic Studies*, 71, 483–512.
- SEIM, K. (2006): “An empirical model of firm entry with endogenous product-type choices,” *The RAND Journal of Economics*, 37, 619–640.
- SHAPIRO, M. H. (2018): “Density of Demand and the Benefit of Uber,” *Unpublished working paper*.
- SULLIVAN, M. (2026): “Fee Optimality in a Two-Sided Market,” *Journal of Political Economy*, forthcoming.
- TAN, G. AND J. ZHOU (2021): “The Effects of Competition and Entry in Multi-sided Markets,” *Review of Economic Studies*, 88, 1002–1030.
- TEH, T.-H., C. LIU, J. WRIGHT, AND J. ZHOU (2023): “Multihoming and oligopolistic platform competition,” *American Economic Journal: Microeconomics*, 15, 68–113.
- U.S. DEPARTMENT OF JUSTICE AND FEDERAL TRADE COMMISSION (2023): “Merger Guidelines,” Washington, DC, December 18, 2023.
- WANG, L. (2025): “Regulating Competing Payment Networks,” *Unpublished Working Paper*.